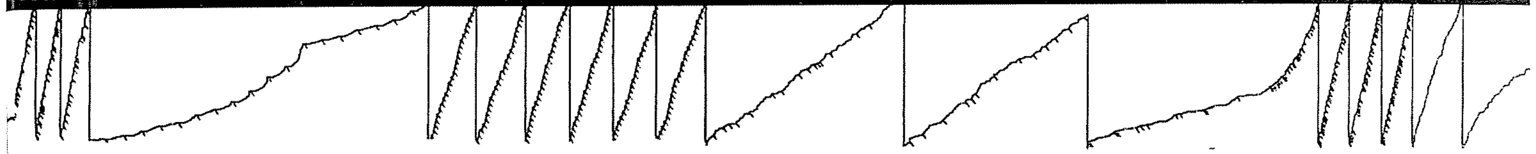

LEARNING

Fifth Edition

Part I: INTRODUCTION
Part II: BEHAVIOR WITHOUT LEARNING
Part III: LEARNING WITHOUT WORDS
Part IV: LEARNING WITH WORDS
Part V: CONCLUSION

A. CHARLES CATANIA

THE CAMBRIDGE CENTER–SLOAN CENTURY SERIES
IN BEHAVIOR ANALYSIS

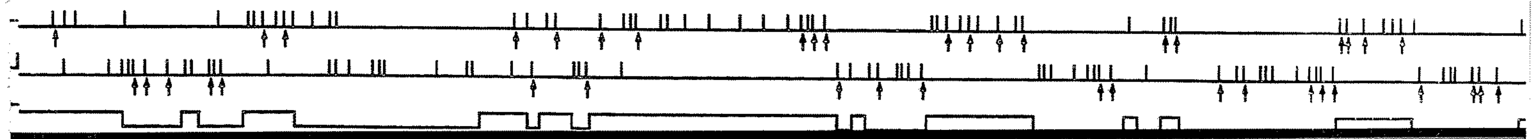

LEARNING

Fifth Edition

A. Charles Catania

2013
Sloan Publishing
Cornwall on Hudson, NY 12520

Library of Congress Cataloging-in-Publication Data

Catania, A. Charles.
Learning / A. Charles Catania. -- 5th ed.
p. cm. -- (The Cambridge Center-Sloan century series in behavior analysis)
ISBN 978-1-59738-023-2
1. Learning, Psychology of. 2. Conditioned response. 3. Psychology, Comparative. I. Title.
BF318.C37 2013
153.1'5--dc23
2012029537

Cover design by Amy Rosen, K&M Design, Inc., based on top and bottom cumulative and event-record friezes provided by the author.

Four previous editions were published by Pearson Education, Inc.

Sloan Publishing, LLC
220 Maple Road
Cornwall-on-Hudson, NY 12520

Printed in the United States of America

10 9 8

ISBN 13: 978-1-59738-023-2
ISBN 10: 1-59738-023-7

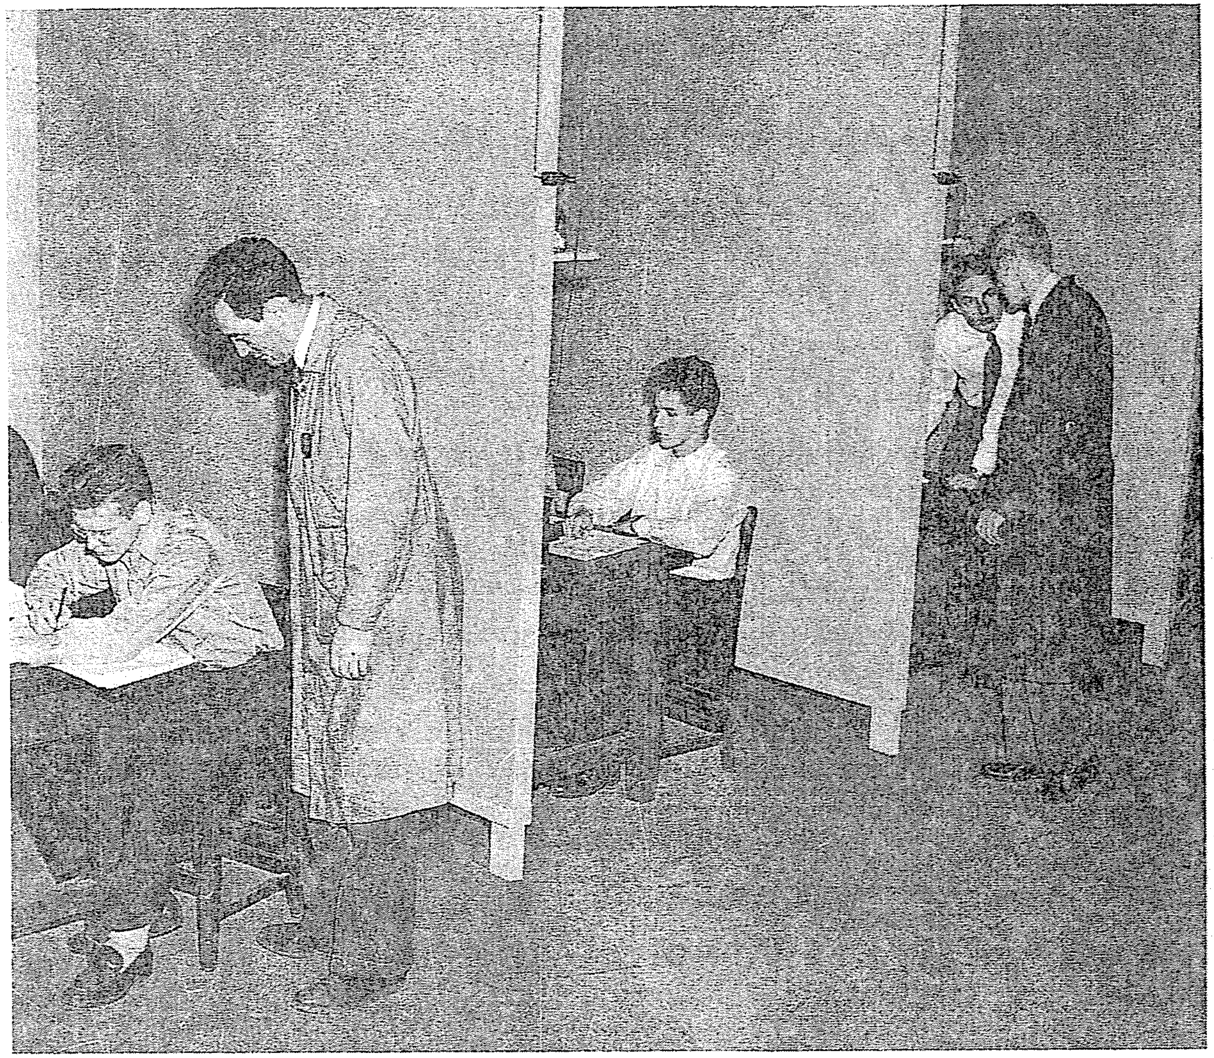

The undergraduate laboratory in *PSYCH 1-2: Introductory Psychology* at Columbia College, Spring 1954. Nat Schoenfeld is standing at the left looking at a student's work. Fred Keller is standing at the right talking to another student, the author.

The Cambridge Center—Sloan Century Series in Behavior Analysis

Fred Keller
At My Own Pace
(1-59738-017-2)

G. Roy Mayer, Beth Sulzer-Azaroff, and Michele Wallace
Behavior Analysis for Lasting Change, 2nd Edition
(1-59738-032-6)

Beth Sulzer-Azaroff, Kathleen Dyer, Susan Dupont, and Dianne Soucy
Applying Behavior Analysis Across the Autism Spectrum: A Field Guide for Practitioners, 2nd Edition
(1-59738-035-5)

Erik Mayville and James A. Mulick, editors
Behavioral Foundations of Effective Autism Treatment
(1-59738-031-8)

Jay Moore
Conceptual Foundations of Radical Behaviorism
(1-59738-011-3)

William Uttal
Time, Space, and Number in Physics and Psychology
(1-59738-015-6)

A. Charles Catania
Learning: 5th Edition
(1-59738-023-7)

A. Charles Catania, *Series Advisor*

For more information about these titles, visit our website at
http://www.sloanpublishing.com/behavior

Brief Contents

Part 1 Introduction

Part II Behavior Without Learning

Part III Learning Without Words

Part IV Learning With Words

Part V Conclusion

Contents

PART III LEARNING WITHOUT WORDS

PART IV LEARNING WITH WORDS

To Connie, Bill and Ken, and to Chessie, too.

Preface

A little learning is a dang'rous thing;
Drink deep, or taste not...
— Alexander Pope

I've no doubt been presumptuous in choosing my title. I've heard that the faculty of an eminent university refused to allow the offering of a course entitled *Learning* on the grounds that learning is the domain of all scholarship. The course was accepted only after its name was changed to *The Psychology of Learning*. The trouble is that covering the topic now involves more than one discipline and is relevant to more than one academic department. But a more thoroughly descriptive title such as *The Psychology of Learning in Relation to Behavior Analysis, Biology, Cognitive Science, Linguistics, Social Psychology and Various Other Disciplines* is just a bit too unwieldy.

To ask what an organism can learn is to ask how much of its behavior depends on its evolutionary history and how much depends on what it has experienced during its own lifetime. Studies of learning have ranged from relatively simple animal procedures to the complexities of human language and problem solving. Research in these areas is so different and the literatures are so extensive that it is tempting to restrict attention solely to issues of nonhuman behavior and learning or solely to issues of human learning and memory. Many learning texts yield to this temptation. In turn, the study of learning has become more and more divided, with each approach developing its own languages and research methodologies.

In this book, part of my purpose is to bring these literatures together and to explore some continuities between human learning and learning by other organisms. Humans are undoubtedly unique, but they share an evolutionary heritage with other species. The properties of nonhuman learning are therefore likely to be relevant to learning by humans. If we do nothing more than show that some types of human learning aren't reducible to types of learning known to occur with other organisms, we've at least begun to define what is peculiarly human.

THE CONTENTS

This book surveys major areas in learning from a consistent behavioral point of view. I won't attempt to outline the nature of a behavioral orientation here. That view has evolved considerably from its parochial origins and is better treated in the context of specific psychological issues, many of which are taken up in the text. I'll only note that a behavioral position needn't exclude aspects of human behaving such as feelings and imagining and emotion. For those who like to think in terms of scientific paradigms and paradigm shifts, this text illustrates a behavioral paradigm that emerged among behavior analysts during the first half of the 20th century and consolidated itself during the second half. With selection at its core, it encompasses all of the phenomena of behavior. Thus, topics often regarded as the province of cognitive psychology or social psychology or linguistics will be treated here along with those more traditionally regarded as behavioral.

In its overall structure, this book has three major parts (II through IV) framed by an introduction (I) and a conclusion (V). Part II covers behavior without learning in an evolutionary context. Part III deals with learning without words, surveying basic topics in nonhuman behavior and learning. Part IV extends the treatment to words and examines human learning and memory. These parts are in some places fairly independent, but more often the concepts developed earlier are prerequisites for the treatment of more complex issues later.

Some advanced topics are covered in addenda that appear at the end of some chapters. The material in later chapters does not build upon material in the addenda of earlier ones, so instructors can include none or some or all of the addenda depending on course level. The chapter organization may also allow instructors conveniently to devote a course exclusively to the basic topics of nonhuman learning by using Chapters 1 through 19, perhaps along with the concluding Chapter 30.

I've made a special effort to include examples of human behavior in discussing the relevance of nonhuman studies of learning and to refer to appropriate concepts from nonhuman behavior in discussing human learning. The underlying behavior taxonomy of the field has been somewhat expanded, now that it is more respectable than once it was to speak of taxonomies for categories other than plants or animals. A philosophy of science has emerged as behavior analysis turned its attention to the behavior of the scientist; I have made this philosophy more explicit in this edition. As in the last revision, examples of human applications of basic processes have been added and enhanced throughout. Some show how behavior analysis may contribute to education in significant ways. It would be surprising if the topic of learning did not yield such implications.

ON THEORY

Throughout its history, the psychology of learning has been concerned with theories. Particular learning theories were developed, elaborated, and then displaced by others. Many remain with us, typically more restricted in their scope than when they were first introduced. Whatever their current status, the data that gave rise to these theories are still to be dealt with. For that reason, and in the tradition of Fred Keller and Nat Schoenfeld, I've emphasizes research findings rather than learning theories. Their usages also included the earliest consistent constraints on speaking of reinforcing responses but not organisms (Keller & Schoenfeld, 1950).

This approach is theoretical mainly in that it adheres to a consistent behavioral language and attempts a systematic organization that accommodates a broad range of research on learning. Nevertheless, I've tried to include enough detail about experimental procedures, terminology and data to provide an effective starting point for students, instructors or general readers who wish to pursue particular theories. The emphasis of the book isn't so much on the interpretation of particular findings as on the relations among the varied phenomena included within the analysis of behavior. Whenever I encountered new and interesting findings in the research literature, I found that one of the most crucial questions to resolve was where to put it were I to include it somewhere in this book. Whether there was a place for it at all was a test of the adequacy of the taxonomy that drove the organization of the book. My intent has been to make this book useful not only to those who are already behaviorally inclined but also to those who are decidedly not of the behavioral persuasion.

SUPPLEMENTARY MATERIAL

A *Test-Item File* of objective questions based on the content of previous editions remains available, and it will gradually be expanded to include items based on new material in this edition. Early editions were once supplemented by a set of computer programs called *Behavior on a Disk* that ran on the MS-DOS operating system. They included shaping simulations, demonstrations of cumulative records and exercises on reinforcement schedules, among others. With the advent of PCs and Windows, those programs became obsolete, but with recent MS-DOS emulators that run on Macintosh computers (e.g., *DOSBox* and *Boxer*) some of them may again become usable.

Like the previous editions, this fifth edition of *Learning* includes etymological notes at the beginning of each chapter. These capsule word histories are important reminders of how easily our language changes. Consistencies in vocabulary are essential to technical treatments, but the language must also grow and adapt to new findings and new perspectives. We must use our language of behavior with care, but perhaps we'll be less likely to become rigid about it if we know something of its origins. The first page of each chapter also includes a chapter outline and a list of key words.

This edition, again like the earlier ones, includes a glossary that summarizes the major terminology of the field and may provide a convenient organizer for study and review. The preparation of a glossary forces an author to attend to potential contradictions and ambiguities in basic concepts and may serve that function for readers as well. The glossary includes an introduction with some comments on its scope and its special features. In the reference section, entries are keyed to the chapters in which they are cited; they've been chosen as useful starting points for exploring the literature of learning as well as for documenting specific points.

ORIGINS AND APPRECIATIONS

The content of this book grew over successive offerings of an introductory lecture course in the Psychology of Learning, first at the University College of Arts and Science of New York University and then at UMBC, the University of Maryland, Baltimore County. I'm indebted to my students and colleagues at both campuses, and especially to Eliot H. Shimoff, the most proximal and thorough and long-standing commentator on and critic of this book over its several editions. I hope he would have approved of its latest incarnation. I'm exceedingly grateful to those students who not only studied their text critically but were undeterred from calling something they found amiss to my attention. I thank Hank Pennypacker and Mike Cataldo in their roles as Board members of the Cambridge Center for Behavioral Studies in helping to make this book possible.

As my teachers and colleagues and friends, numerous others contributed by their comments, discussions and encouragement. I must particularly single out for mention Philip N. Hineline, Peter Killeen, Allen Neuringer, and John A. Nevin, and one outside of behavior analysis but with a keen appreciation for this approach, Norbert Hirschhorn. But any proper list is bound to be far more extensive. Knowing that some omissions are inevitable, I here name a few others, in many cases because I can still identify particular contributions of each: Nathan H. Azrin, Beatrice H. Barrett, John Borrero, Joseph V. Brady, Thomas A. Brigham, Guy Bruce, Joseph Cautelli, Daniel Cerutti, Leonard Cook, Willard F. Day, Iser DeLeon, Deisy de Souza, Peter B. Dews, James A. Dinsmoor, Charles B. Ferster, Lanny Fields, Israel Goldiamond, Lewis R. Gollub, Ernest S. Graham, Joel Greenspoon, Stevan Harnad, Eliot Hearst, Ralph F. Hefferline, Per Holth, Koji Hori, Pauline Horne, Herbert M. Jenkins, Kenneth J. Keller, Tatiana Lary, Victor G. Laties, Kennon A.

Lattal, Ogden Lindsley, Richard A. Littman, C. Fergus Lowe, Byron A. Matthews, Ernst L. Moerk, William H. Morse, Koichi Ono, David Palmer, Giovambattista Presti, Robert R. Provine, Howard Rachlin, Mark Reilly, Bob Remington, George S. Reynolds, Marc Richelle, Terje Sagvolden, Masaya Sato, Charles P. Shimp, Murray Sidman, B. F. Skinner, Deisy de Souza, William C. Stebbins, S. S. Stevens, Mark Sundberg, Zera Taylor, João Todorov, Vicci Tucci, William S. Verplanck, and Ray Weitzman, plus all of those other scholars who contributed to the literature upon which I drew. Fortunately the contributions of some others I may have missed are likely to have been acknowledged by their inclusion among the references. The list includes too many who are no longer with us; they have left us with a rich intellectual legacy.

Some things seem to go without saying, but I must say nevertheless that I recognize now more than ever how much I've learned from my family: first and foremost my wife, Connie, who contributed wisdom and patience and so much more; my sons, Bill and Ken Catania; my brother, Robert Catania; my parents, Charles John and Elizabeth Catania; among my extended family my uncle, Ray Benza, and my cousin, Albert D'Angelo, along with many other cousins and aunts and uncles and nieces and nephews; and my new-found cousin Ivano Zappulla, who helped me to discover how much larger my extended family was than I had once thought.

I must especially acknowledge Bill Webber, the kind of publisher about whom all authors dream: available, knowledgeable, efficient, constructive, patient. He was there every step of the way with encouragement, advice and technical expertise. I could not have brought this project to fruition without his unflagging and enthusiastic support.

Above all, and as before, another acknowledgment remains. I owe more than I can say to Nat and Fred: W. N. Schoenfeld and Fred S. Keller. Their courses and their *Principles of Psychology* introduced me to the analysis of behavior and irrevocably committed me to exploring its wonders. I marvel still at what a trip it has been. I hope that this book is true enough to their teaching that readers who knew them and learned from them will recognize something of them in these pages.

A. Charles Catania
Columbia, Maryland

Student Preface

You'll encounter some topics in this text that may appeal to you; others may disturb you. As you meet them in this text, what will you think of electric shock in punishment and escape, evolution, social aspects of discrimination learning, failures of memory, and how verbal behavior affects political and religious institutions, to name just a few? In understanding something, it shouldn't matter whether you like it or not but it usually does. For example, punishment occurs often in human settings and we may not like it, but still we should know whether it works and if so how it works. You can react to a subject matter in terms of whether you like it or in terms of whether you get it. If you think something is hard to understand, maybe the problem is instead that you don't like its implications. That may be worth considering. Understanding something about how behavior works doesn't imply endorsing it; on the other hand, if you don't master it that doesn't mean it will go away.

The approach we'll take here is often unconventional. For example, we'll question everyday practices like attributing our behavior to our emotions. Are you ready for the argument that it is misleading to say that emotions are causes of behavior? We learn names for emotions from those around us when we see people acting angrily or happily or sadly. Later we see others acting similarly and are inclined to say they did so because they were angry or happy or sad. This is a case of circular reasoning, not explanation. Be alert for such problems; we'll encounter many examples.

A course in learning ought to offer a lot about studying and mastery and taking notes, but you might get to some relevant material only long after you need it most. Among topics that may be of practical value are metamemory (managing your remembering), mnemonics (memory techniques), and verbal learning (list organization). Browse the index for terms like *study habits*, *education* and *teaching* and you may sometimes be surprised where the entries take you, but you might also find it useful to figure out why they are where they are.

One thing to learn about learning is that learners learn what they do, so what you learn in any course depends far more on what you do than on what your instructor does. Your doing includes reading, taking exams or quizzes, figuring out why you were able to answer some questions but not others, writing papers, memorizing, making judgments about what you've mastered and what needs more study, taking notes, and so on. Any time you think about asking your instructor to do something more for you, like prepare more outlines, consider the possibility that you will learn far more by doing it yourself. Your instructor is not there to pour knowledge into you; your instructor is there to get you to engage with this material and you cannot be passive and still learn. The more you do, the more you'll learn.

But choose your study plans carefully. Some things you do might make minimal or even negative contributions. For example, does highlighting do anything but save you time later when you skip what isn't highlighted? Probably not. And wasn't highlighting the first time through the worst possible time to tell important from unimportant passages? Probably. If later while studying you skip all but what you highlighted, you might be skipping what you most need to review. Recommendation: don't highlight the first time around, and if you must highlight don't do it in an irreversible way; use pencil in the margins rather than permanent markers.

This is just one example. The text provides others. It also offers some tools for evaluating your study practices and for creating new ones of your own. You don't have to wait until you get to them in your syllabus to start finding out about them. There is a lot to learn, but heed the warning from Alexander Pope with which the main preface began. My hope is that you'll come away from your time with this text not only having learned a lot but also better able to judge for yourself what works for you as you study and what doesn't.

A. Charles Catania
Columbia, Maryland

PART I
INTRODUCTION

Chapter 1
Learning and Behavior: Conditioning

> *The English word* learning *probably comes from an Indo European root, leis, which meant a track or furrow. Before reaching its present form, it went through many changes:* læstan, leornian, lernen. *At various times in the evolution of our language it might have been understood as following a track, continuing, coming to know, or perhaps even getting into a rut. The verb* last, *to endure, came from the same root.*

The Language of Learning and Behavior
  Behavioral and Cognitive Languages
  The World and the Laboratory
Antecedents, Behavior, Consequences
  Stimuli and Responses
  Behavior Hierarchies

What is learning? Any satisfactory definition must include learning where things are, learning what has happened in the news, learning how to get along with people, learning definitions, and learning how to play a game or ride a bicycle. But how is learning what to say related to learning what to do and how to do it? How is learning how something works related to learning how to use it? Do these examples have anything in common? There is at least one thing we can say: whatever else goes on, you're different after having learned something than you were before. Something about you is new.

It's usually easy to see what's new. We can say some things we couldn't say before or do some things we couldn't do before. But how did that happen? What are the sources of our new behavior? This book deals with what we know about the answers to questions like these. Before we can get down to the details, however, we need to cover some preliminaries. We need to worry about the behavior already there before learning begins. Before we can even do that we need to consider what behavior is—not just ours but also the behavior of the many varieties of other organisms with which we humans share our planet. Throughout all of this, behavior will be paramount. We are who we are because of what we can do, and everything about us evolved in the service of our behavior.

We'll start by seeing what behavior is like without learning. We'll examine its origins and its evolution and its development. With that foundation, we'll next consider learning when there are no words. We humans are special particularly because we talk. We talk so much that we find it hard to look at the behavior of the nonverbal creatures with which we share our world without projecting our

KEY TERMS: Antecedents, Behavior and Consequences (ABC) in a Three-term Contingency; Behavior Analysis; Behaviorism and Mentalism; Behavior Hierarchy; Function and Structure; Response, Response Class, Sponse; Stimulus (plural = Stimuli).

words onto them. In many ways they are as alien to us as extraterrestrials. We have to approach them without assuming that they think like us. To make sense of their behavior we must not assume that they do what they do by first talking to themselves about it. Only after we have seen how nonverbal learning works will we be able to appreciate the new kinds of learning that words make possible. We will also see that nonverbal learning provides the underpinnings of verbal learning: we cannot do anything with words unless they are built on what was there before words existed. Once we get that far, we'll be ready to talk more clearly about topics especially dear to us, like human language and memory and cognition.

Now let's get back to learning. Suppose I'm unfamiliar with the word *phenomenon* and then I notice it in a sentence. I may decide from the context that it means something like an event worth noticing. If I look it up I might find it defined as an event that can be observed and probably also as a remarkable or unusual person or thing. The definition might show the word as a singular noun with *phenomena* as its plural, and this form of plural, so different from the usual final -*s*, might suggest that this is a pretty old word with Latin origins. What I learn by checking its definition could be useful the next time I come across the word even if I hardly ever use it myself.

But what about the definition of the subject matter of this book? What is this phenomenon called *learning*? The word doesn't give us trouble in everyday talk, but a dictionary definition telling us that it means getting to know something or gaining knowledge and skill isn't very helpful. The word *learning* is more familiar than *phenomenon* and yet is much harder to define. Sure, we can usually say whether we've learned something and we can usually agree on what counts as learning. Even so, we run into problems when we try to frame a definition. For example, a textbook might define learning as a relatively permanent change in behavior resulting from experience (e.g., Kimble, 1961 pp. 1–13). But what is meant by *behavior* and by *experience*, and how permanent is *relatively permanent*? Staring at an eclipse of the sun is an experience and if it damages your eyes it will certainly change your behavior. Yet if I told you this damage was an example of learning I hope you'd disagree.

THE LANGUAGE OF LEARNING AND BEHAVIOR

From the start we must face the fact that we won't be able to define learning. There are no satisfactory definitions. But we won't let that stop us. We can look at how organisms come to behave in new ways. In our study of learning, we'll ask two very different types of questions: (1) What is the nature of these events we call learning? and (2) What is the best way to talk about them? Our main concern will be to show how learning works. What are the conditions under which organisms learn and what happens as they do so? This will give us plenty to do, so we'll only occasionally consider theories of learning or explanations of learning in terms of changes in the brain or other physiological events, though those topics are of interest in their own right.

Words like *learning* or *knowledge* seem obviously important, but they function in different ways in different contexts. For example, sometimes we speak of learning *about* something; at other times we speak of learning *how to do* something. I could learn how a car works without knowing how to drive one; conversely, I could learn to drive a car without being able to say how it works. Some kinds of learning involve deeds and others involve words. Should we treat these two kinds together or separately? Philosophers make this kind of distinction when they debate the difference between "knowing how" and "knowing that" (e.g., Ryle, 1949). Psychologists make it by contrasting *procedural* knowledge or memory with *declarative* knowledge or memory. The distinction is so fundamental that, as the table of contents shows, this book includes two major parts, one concerned with learning that doesn't involve words and the other with learning that does. Learning means different things at different times to different people.

Consider some examples. A pigeon discovers food in its travels and returns to the same place later when hungry. A child becomes able to read a story or to spell simple words. A dog is taught to sit or lie down on command. A patient who once had a bad experience in a dentist's office feels uneasy in the waiting room. A young cat, after its early hunting expeditions, now avoids skunks and porcupines. A shopper sees an announcement for a sale that hasn't begun yet and several days later returns to the store to take advantage of bargain prices. An author who encounters an unfamiliar word later uses it in a short story. You read a chapter in a mathematics text and find a way to solve a problem that had baffled you. I need to check a point in a chapter of this book and find a relevant paper through an internet search. What do these examples have in common? They involve dogs and cats, children and adults, and we'd probably agree that they are all instances of learning. But is it reasonable to group a pigeon who learns a route to food with a human who discovers a solution to a mathematical problem? Are these examples as distantly related as chipmunks and spiders and clams are distantly related in the realm of animal life?

Can we resolve our problem by adding that learning has to come about through some change in the brain? We might think so, but do we look at an organism's brain to decide whether it has learned? Even if we could watch a brain doing something, how would we know that learning was what it was doing? No doubt the brain does a lot, but except with the aid of sophisticated instruments we see only its products. Ordinarily we see only behavior.

This isn't to say that learning has no physiological basis. Of course it does, and it would be fascinating to find out what neurological changes accompany learning. Yet we'd have trouble figuring out what to look for in the nervous system if we didn't know much about learning. In fact, we can't have an adequate neuroscience of learning unless we understand its behavioral properties. Those properties determine what neuroscientists must look for in the nervous system if they want to know what happens during learning. That's why our main concern will be with the behavioral properties of learning rather than with its physiological basis.

Furthermore, the brain changes as a function of changes in the interaction of an organism with its environment. For example, when a stroke has partially paralyzed someone's arm, one therapy is to restrain the other arm so that use of the impaired arm will help it to regain full function. This is said to change the brain, but the recovery of the arm and not the change in the brain is the objective of the therapy. Of course the changes in the brain matter, but if the therapy is successful it is important to remember that the behavior, the use of the arm, is what drove those changes. That is why, though we will sometimes consider the nervous system, behavior will always be primary. **Behavior will always be our starting point.**

So far we've hardly worried about the facts of learning; we've mostly worried about how we talk about it. Languages are variable; their vocabularies reflect what is currently important to their speakers. One trouble is that the language that has evolved in our everyday interactions with others isn't necessarily well suited to be a language of learning (that is one reason for the etymologies, or word histories, included at the beginning of each chapter).

We're usually more interested in what other people know and in what they're likely to do than in how they came to be that way. For example, a parent might worry about a child who fights with other children and rarely plays cooperatively. If the child begins to play cooperatively, the parent might not care whether this happened because of the natural rewards of cooperative play or because cooperative play was explicitly taught or because fighting was punished. The child's play might look the same in each of these cases, but it might make a difference how the child got there.

Our language for describing what people do is useful. It is important to know what to expect of others, and that's probably why we describe people in terms of how they are likely to behave. We speak of each other as *outgoing* or *reserved*, *easygoing* or *compulsive*, *trustworthy* or *unreliable*. Describing people with words like *artistic*, *athletic*, *social*, *intellectual* or *musical* specifies their preferred activities. Yet this kind of vocabulary isn't suitable for discussing how someone's particular interests or traits arose or changed.

Consider the difference between lying and telling the truth. If one child learns to stay out of trouble by telling lies and another by telling the truth, should we be surprised if the first child grows up to be less truthful than the second? Yet the behavior of each child is shaped by its consequences: Each child behaves so as to keep out of trouble. This shaping of behavior should concern us, whichever way it leads each child, but our everyday vocabulary doesn't equip us well for discussing it.

Similar problems exist in other fields. When physicists look at events in the world, they don't find the everyday vocabulary adequate. Sometimes it even gets in the way. So they coin new terms or take over existing ones. The latter tactic can create trouble. Words like *work*, *force* and *energy*, for example, mean different things to physicists in their technical talk than they do to most people in casual conversation. Fortunately for physicists, much of what they now study is remote enough from our daily experience that we don't confuse their technical language with our everyday talk.

This isn't so for behavior. We are all inescapably involved with it. We speak of how people grow and change and we speculate about why they do things. If we want to talk about these events in new ways, we must take care not to confuse our new ways of talking with the old ways. We've all spent most of our lives talking about what we do, but those familiar ways may interfere with our new ways of talking, so we must beware of language traps. Some parts of this book will introduce a language of behavior. That language will not be a paraphrase of everyday usages. **It will be fundamentally different.** It will demand new ways of looking at familiar phenomena.

Behavioral and Cognitive Languages

Sometimes we talk about what people do and sometimes we talk about what they know. On the one hand, what someone does is all we have to go by. We have nothing else to study but behavior. Someone in a learning experiment may describe thoughts or feelings, but these descriptions are still only behavior; verbal behavior may be special, but it is behavior nonetheless. Thoughts and feelings are supposed to be inside of us, but if nobody else can detect them how can anyone teach us how to talk about them? Our words and our theories must ultimately be derived from behavior, from what organisms do. This is another sense in which behavior must come first.

On the other hand, there is more to an organism than shows in its behavior. Two students may sit silently through my class. They aren't behaving differently right now and yet based on their past coursework I may know that one can answer certain questions and solve certain problems while the other cannot. The difference is in what each potentially can do; I might simply say that one student knows more than the other. When we study this knowledge, it is tempting to say that we study the mind.

The debate between those who call themselves behaviorists and those who call themselves cognitivists or mentalists has been long-standing. To a large extent it has been about how we talk. The behaviorist argues that, because behavior is all that is available to measure, the language of mental events can be misleading, especially when a mentalistic account is accepted as explanatory and therefore discourages further inquiry.

For example, we sometimes casually say that an idea, a feeling or a hunch led someone to do something. The behaviorist doesn't dispute the existence of ideas, feelings and hunches, but rather criticizes appeals to them as causes of behavior. It is too easy to be satisfied with an explanation in these terms. For a behaviorist, it isn't enough to say that someone did something because of an idea, a feeling or a hunch. Ideas, feelings and hunches are about the world, and therefore must have their origins in our experiences with the world. We must look further, to these past experiences or, in other words, to past behavior, to account for what we do. If we are successful, we may also have something useful to say about the origins of our ideas, feelings and hunches.

The cognitivist maintains that such a view is unnecessarily narrow. Things must occur in our dealings with the world that aren't observable in our behavior. When we try to recall a word that is "on the tip of our tongue" or try to solve a problem by "sleeping on it," things are happening that don't show directly in our behavior and we may not even be able to report them. If we can find out something about such events, it is bound to be relevant to our study of learning.

This dispute stems as much from different ways of talking about behavior as from differences in research findings. Some difficulties arise because behaviorists and cognitivists are often interested in very different types of questions. Behaviorists tend to deal with questions of function and cognitivists with questions of structure.

Suppose I'd like to teach a child to read. Where do I start? On the one hand, I could worry about how to involve the child in reading. What will keep the child alert, what will help the child to attend to the words presented, and what will help the child remember what the various words are? Will I be more successful rewarding the child's right answers or penalizing the child's wrong ones? When I arrange different consequences for the child's different answers, I determine the functions of these answers or, more technically, the functional relations between behavior and its consequences.

On the other hand, no amount of worry about the effects of reward and punishment on the child's mastery of reading will tell me the most efficient way to present reading materials to the child. What is the best way to order them? Should I start with single letters, with syllables, or with whole words? When I present these materials in different orders, I'm structuring the subject matter. Are words best taught as units or as structures built up from simpler parts such as letters or syllables? Problems of structure are concerned with the organization of behavior and of the world within which it takes place.

Educators concerned mainly with function might try to improve a school system by changing what happens in the classroom, without worrying much about how the curriculum is structured. Those concerned mainly with structure might try to improve the school system by changing the curriculum, without worrying much about what happens in the classroom. But obviously both are important. Any attempt to improve how children learn to read that ignores either is likely to prove deficient. Any that ignores both, as in concentrating mainly on instilling a vaguely defined trait such as self-esteem, has no hope at all.

Historically, some controversies about learning arose because those interested in functional problems tended to speak a behavioral language whereas those interested in structural problems tended to speak a cognitive or mental language. The problems in which behaviorists and cognitivists were most interested interacted with how they each talked about them. If I'm concerned with function, I study the consequences of particular relations between environmental events and actions; I can conveniently express these in the behavioral language of stimuli and responses. If I'm concerned with structure, I study the properties of particular capacities or abilities; I can conveniently express these in the cognitive language of knowledge and mind. This issue is not unique to psychology. A parallel distinction between structure and function, the separation of anatomy and physiology, was significant in the history of biology (e.g., Russell, 1916).

Let's not be sidetracked by this controversy. We'll consider both function and structure in learning and we'll therefore examine both types of research. In either case, it will often be useful to describe situations in terms of *antecedents*, or the circumstances that set the occasion for behavior, the *behavior* that occurs in those circumstances, and the *consequences* of the behavior (these three terms are conveniently abbreviated as ABC). We can consider either *function*, the relations among the terms (given certain antecedents, what consequences are produced by responses?) or *structure*, the properties of behavior (how are environments and responses organized?).

Though I will usually emphasize the behavioral approach, this book deals with both function and structure and therefore encompasses both behavioral and cognitive concepts. These two approaches differ in their languages and in what they study, but they have in common the reliance on experimental method, the anchoring of concepts to experimental observations, and the assumption that our subject matter, however complex, is orderly and not capricious.

The World and the Laboratory

How then do we find out about behavior? Our language is not the only problem. We live in a complex world. The events that influence our behavior don't occur in isolation. Thus, to understand a situation we must strip away the unessentials. We must analyze it. To analyze something is simply to break it down into its component parts. To do this we start in the laboratory, studying organisms simpler than ourselves, in simple environments, in a science called behavior analysis. We must face the objection, of course, that a laboratory experiment is artificial and so may be inappropriate for establishing generalizations about learning outside of the laboratory. But starting with simple events helps us to develop techniques and vocabularies applicable to complex ones. Biologists could not have mapped the human genome if they hadn't started with the genetics of simpler creatures, such as the fruit fly *Drosophila melanogaster* or, simpler still (but nevertheless complex), the bacterium *E. coli.*

In the laboratory environment we look at one thing at a time. We arrange circumstances so that we know what goes into a situation; if we're careful enough, we'll exclude distractions that might otherwise obscure what we wish to study. The simplicity of our laboratory environment may also help us to see things more clearly. We must be able to identify events before we can study their properties. Yet even after we've studied behavior inside the laboratory, we can't expect to successfully interpret every instance of behavior outside. There are limits to what we can know. It is tempting to ask why someone did this or that, what led to a certain incident, how someone came to have certain interests, fears or attachments. But usually we have so little information that giving a plausible interpretation is the best we can do.

In this respect, behavior analysis is not much different from other sciences. If I see a leaf blow across cars on a busy street and land at the foot of someone sitting on a sidewalk bench, I couldn't say how or why it got there. But a failure to account for every twist and turn in the path of that falling leaf doesn't invalidate the principles of aerodynamics. We can't possibly measure the details of air currents, leaf surface, and so on in enough detail. Similarly, the principles of behavior aren't invalidated when we can't account for what someone did on some occasion. Again, we simply can't measure personal history and other factors in enough detail. In our study of learning, we must recognize what remains out of our reach. In what follows, we'll find that the most profitable course is one that stays close to the data. We'll worry less about theory than about properly describing our findings. It will usually be more useful to *describe* what an organism has learned or remembered than to try to *explain* its learning or its remembering.

ANTECEDENTS, BEHAVIOR, CONSEQUENCES

Behavior is no easier to define than learning. We may say glibly that behavior is anything we do, but this definition is too global. Should we count respiration or metabolism along with muscle movements and glandular secretions? We describe behavior with verbs: people walk, talk, think, do things. But we also distinguish between active and passive actions. We may say that someone breathes, but are we likely to say that someone heartbeats? People bleed when cut but do we want to call their bleeding behavior? Let's not try to resolve this problem. The phenomena of behavior are varied, and though they sometimes share common names we'll probably do better by considering examples than by attempting definitions. We can deal with specific examples without much risk of misunderstanding.

Stimuli and Responses

When we observe an organism, we see properties of its environment, *stimuli*, and properties of its behavior, *responses*. In the singular, we speak of a *stimulus* and a *response*. Neither is of interest by itself. An experimental analysis determines what kinds of relations exist between stimuli and responses and how they can be changed. It must also consider the situations or broader contexts within which these stimuli and responses are embedded.

Imagine a pigeon in an experimental chamber. On one wall is an opening to a feeder that can dispense food. Above the feeder opening is a recessed translucent disk or key that can be lit from behind. The pigeon has learned to earn food by pecking the key whenever it is lit. Now suppose the pigeon hasn't eaten for a while, the key is lit, and a peck on the key immediately makes some food available. We need to know the context before we can guess what the pigeon will do. It is one thing if the alternative, not pecking, is never followed by food; it is another if not pecking is followed by a somewhat delayed but much larger amount of food. In each case a response, the key peck, is followed by a stimulus, food. But the contexts are very different. We would expect the pigeon to peck the key in the first case, but what about the second? If the pigeon doesn't peck, we might want to say that it shows self-control, forgoing the small amount of immediate food in favor of the larger but delayed amount. We'll discuss this type of situation in more detail later. For now, the point is that we must look at not only the moment-to-moment details of events but also their contexts over extended times.

Let's examine relations between environment and behavior further by observing a human infant. We might want to start by asking what the infant feels, but that won't work. She isn't yet verbal and can't tell us. Even if she were an older child who could tell us, we'd have to wonder how she learned the words and whether they'd mean the same thing to us as to those who taught them to her. When we get to language, in Part IV, we'll examine the role it plays in molding our knowledge of ourselves and others, but that won't help us here.

We know the infant is active, learning from the environment and interacting with it. But how do we find out what's going on? We watch for a while and see her move her hands or arms or legs. Perhaps at some point she begins to cry. If the crying stops without our intervention, she may sleep or may lie quietly with open eyes. If we look closely, we may see her eyes moving, although it might be difficult to judge what she is looking at, if anything. We could catalogue what she does and we might discover that some movements usually occur in particular sequences. But if we only watch, we can't say much more than that she does different things more or less often and more or less in certain orders.

We needn't be restricted to watching. We might touch or rock her, move objects in or out of her view, make sounds, or offer a pacifier. We'd expect her to respond to each event in a characteristic way. If we touched her palm, for example, she would most likely clench that fist, grasping whatever touched it. The vocabulary for these events is already familiar: we call the touch to the palm a *stimulus*, and the grasping a *response*. We may notice that this stimulus and this response are correlated (Skinner, 1931); they tend to occur together in sequence.

Besides the grasping produced by a touch to the palm, we could catalogue other examples of stimulus-response correlations: crying caused by a sudden loud noise; sucking produced by a nipple in the mouth; blinking triggered by a flash of light. We see the environment act on this infant when stimuli produce responses, but things may go in the other direction. She can act on the environment. Her crying, for example, often brings her mother's attention. Crying, then, is a response that often produces a consequence: mother's presence. This case involves stimuli and responses, but here the responses come first, not the stimuli; here behavior has consequences.

The relations can get more complicated. If the infant's eyes move while the lights are on, those eye movements change what she sees. Eye movements can't have this effect with the lights off. Thus, she may come to look around in the light but not in the dark. In the presence of one stimulus, the light, moving the eyes produces other stimuli, some new things seen. Eye movements can't have this consequence in the dark. The relation involves three terms: an antecedent stimulus, the light; a response, eye movement, in its presence; and a consequence, what is newly seen given this response in its presence. This three-term relation, stimulus-response-consequence, is sometimes called a *three-term contingency*. It is important because an organism's behavior depends on both antecedents and consequences.

An *antecedent* is simply something that comes before. A *consequence* is simply what is caused by or what happens as a result of some event. Thus, everyday usage corresponds pretty well to the senses of these terms when we use them technically. It is important to note that **consequences should not be identified with stimuli**. Responses can have many types of consequences. They sometimes produce stimuli that would otherwise have been absent, but they can also prevent things from happening or change the consequences of other responses. For example, food produced by a response is both a stimulus and a consequence, but food presented independently of behavior is a stimulus only; shock prevented by a response is a stimulus, but the consequence of the response is the absence of shock, which isn't a stimulus. Sometimes the consequence of one response is a change in the consequences of some other response, as when a light switch stops working and changing the light bulb restores the usual consequences of operating the switch.

For *stimulus* and *response*, the relations between technical and everyday usages are not so simple. Stimuli are events in the world and responses are instances of behavior. The term *stimulus* is often restricted to specific physical events such as lights or sounds or touches. But organisms respond to varied features of the environment, including relations (e.g., to the left of, on top of), complex behavior (e.g., facial expressions, tones of voice), functional properties (e.g., edible, comfortable), and so on (cf. Gibson, 1979). We'll often speak of such environmental features as stimuli even when we cannot specify their physical dimensions.

As for the term *response*, everyday usage often implies that it is *to* something (typically a stimulus). The term won't function that way here, however, because an account of what causes responses typically includes other factors along with or instead of the stimuli that precede them. We'll be especially interested in responses that are not elicited by stimuli but are caused in other ways. Such responses are said to be *emitted*. Unfortunately *sponse*, a useful word for such cases, is not a well-established term (but see Provine, 1976).

A stimulus is an environmental event, but such events vary in complexity. When the infant's crying produced the mother's attention, we regarded the mother as a stimulus. An infant's environment is certainly different when the mother is present than when she is absent. Yet what sort of stimulus is the mother? We don't know which aspects of her looks, her voice or her touch are important to the infant early in life. We might guess that the infant wouldn't react to her as usual if she approached wearing a surgical mask, but we couldn't be sure unless we did the experiment. Despite our ignorance with respect to these questions, we don't doubt that the mother is an important part of the infant's environment. It is useful to speak of the effects she has as she comes and goes in the infant's world.

This example again illustrates structure and function. When we try to analyze which of mother's visual, auditory and tactile features are important to the infant, we deal with the structure of this complex stimulus, mother. We might ask how the infant learns to respond to a particular individual as mother despite changes in her dress or hair style, her facial expression or posture. If we concentrate instead on how mother interacts with the infant's responses, we're concerned with the functional significance of mother in the infant's environment. For example, if an infant cries we may not care whether the infant recognizes mother by her face, her hair or her voice, as long as her presence makes a difference; it would be enough to see that when mother goes to the infant the crying stops. Later we'll often be interested in simpler stimuli: lights, sounds, food in the mouth. But even with simpler stimuli we'll have to distinguish between structural problems, as in analyzing stimulus properties, and functional problems, as in analyzing the interactions between stimuli and responses.

And what about responses? In describing responses, we encounter at least two difficulties. The first is that behavior isn't repeated exactly from one instance to the next. If the infant grasps an object on two different occasions, the grasping won't be the same each time. The difference may be small, in the force of the grasp, for example, or in the exact placement of the fingers. But if there is any difference at all, we must worry whether the two grasps should be regarded as two instances of the same response or as two different responses. We must speak not of individual responses but of classes of responses having common properties.

The second difficulty is that responses are sometimes adequately described in terms of movements, but at other times the description must include the environment in which the responses occur. For example, suppose we want to compare the infant's grasping of an object with clenching a fist. If we look just at muscles, grasping an object with the right hand and clenching that fist have more in common than grasping an object with the right hand and grasping an object with the left hand. Yet sometimes it is more important to speak of the act of grasping an object, no matter which hand is used, than to speak of the movement of closing a particular hand.

An account of behavior must distinguish between *movements*, responses defined by their form or the musculature used, and *actions*, responses defined by their relations to the environment. We'll find that actions are more important for our purposes. Consider how often we speak of doing things, going places or manipulating objects, without regard to the details of how we do them.

Even in the absence of movement we sometimes conclude that behavior has occurred. We do many things that involve no obvious movement. For example, while listening to a song I may shift my attention back and forth between the vocalist and the accompaniment. Those shifts of attention are behavior even though we can't record them as movements. Many aspects of thinking and imagining involve no movement, but as things we do they count as varieties of behavior. Thus, **not all instances of behavior need be movements.** In fact, we will later find that a useful criterion for whether something counts as behavior is not its form but whether it varies depending on how it affects the environment.

Whether behavior involves movement or not, it typically has consequences, and one of the most significant consequences of behavior is that it provides opportunities for other behavior. For instance, if a child is given a cookie, the cookie affords the child an opportunity to eat. The significance of the cookie is based on the child's eating, its behavior with respect to that stimulus. As we will see again and again, we cannot characterize stimuli independently of an organism's behavior, nor can we characterize responses independently of an organism's environment. **Behavior is the interaction between an organism and its environment.**

Behavior Hierarchies

One way to classify an organism's behavior is to rank its responses according to the relative frequencies with which the organism engages in them. Such a ranking has been called a behavior hierarchy (Hull, 1943). Let's construct an environment for a rat that gives it access to different stimuli in different compartments. Figure 1-1 shows some possibilities: compartments containing a food hopper, a drinking tube, an activity wheel, or nesting material; a small dark compartment and a large bright one; compartments in which the rat's entry turns on quiet music, a changing display of dim lights, loud noises, a cold shower or a blast of hot air; and, finally, a compartment with an electrified grid floor.

If the rat lives here for a while, it will learn what is in each compartment. From time to time it will enter the food compartment and eat, or the water compartment and drink, or the activity wheel compartment and run. After a few days it might always sleep in the compartment with nesting material, or it might prefer the small dark compartment and carry some nesting material there. It might spend some time looking or listening in the light show and music compartments, but would seldom stay long in the large bright compartment. And after only a few entries it would rarely enter the compartments with loud noise, a cold shower, hot air, or the electrified grid.

If we wondered about responses relatively low on the hierarchy, such as looking at the light show or listening to music, we could rank just those two by restricting the rat's access to all of the other compartments. If we were interested in the behavior occasioned by the compartments the rat rarely enters, such as those with noise or showers or hot air or shock, we could place the rat in these compartments and rank them based on how quickly the rat leaves, or we could set things up so it has to be in one or the other and see which one it stays with.

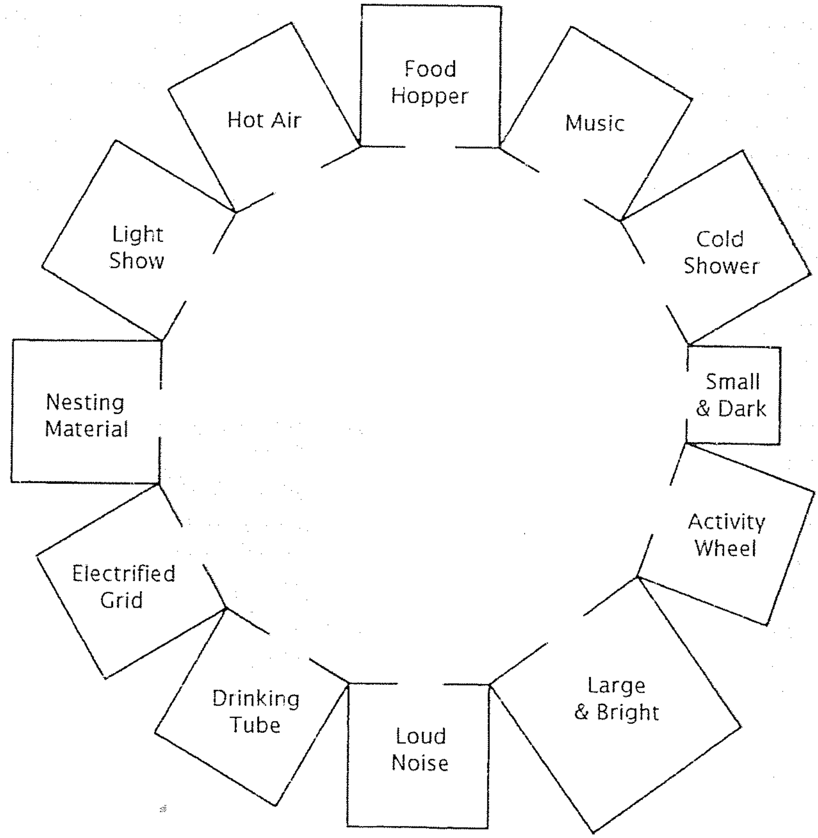

Figure 1-1 A hypothetical behavior space for a rat. The order and frequency of the rat's visits to each compartment and the rapidity of its departure if it is placed in one can be described in terms of a behavior hierarchy. Eating and drinking and running in the wheel are likely to be high on this hierarchy, though either might outrank the other at any particular time; looking at lights or listening to music might occur at an intermediate level; and spending time with loud noises or in a cold shower or on an electrified grid will probably be low in the hierarchy.

Our description of the rat's hierarchies will of course depend on what we make available to it. The compartments in Figure 1-1 will tell us nothing about the rat's social or sexual behavior. If we wanted to know where interaction with other rats stood in the hierarchy, we'd have to add more compartments, with a male rat, a female rat, and different sized groups of rats of one or both sexes.

Now let's look at hierarchies with a child instead of a rat as our example. If we gave a child an opportunity to eat, to play with toys or to take a bath, we might find that the child plays a lot, eats occasionally, and hardly ever volunteers for a bath. Playing, as the most likely or most probable behavior, comes first in this ranking, followed by eating and then by taking a bath. An equivalent way of describing the ranking is in the language of preference: we could say that the child prefers playing to eating and prefers either of these to taking a bath.

Behavior hierarchies are changeable. For example, if we wait until the child's usual mealtime and provide a choice between eating and playing, we might find that eating has become more probable than playing. While eating, the child is neither playing nor taking a bath, but we could find out about the relative rankings of those responses by giving the child a choice between the toys and the bathtub. We might discover that this child almost always prefers the toys. We therefore conclude that right now playing with toys ranks above taking a bath in this child's hierarchy. We might even find that this child always leaves the bathtub area even when there isn't much to do elsewhere. Maybe the child recently had a bad experience there. For any kind of behavior, we must consider when it stops as well as when it starts.

It is often convenient to speak of stimuli rather than of opportunities for responding. Thus, for this child we might describe food as an appetitive stimulus or event and taking a bath as an aversive one, with events that are neither appetitive nor aversive categorized as neutral. Unfortunately, even though we may be able to use such terms in specific situations, stimuli in general cannot be sorted out so neatly. Context makes too much of a difference. We cannot just divide the environment into three simple classes of events called appetitive, neutral and aversive. Instead, we must evaluate each stimulus relative to the others available.

With changes in the behavior hierarchy come changes in the significance of stimuli. For example, consider how food may change from appetitive to aversive over the course of an unusually large holiday dinner. In the bathtub example, if the parents handle things carefully this child may begin to tolerate baths and eventually even come to prefer toys in the tub to toys in other places. We've now surveyed some general properties of stimuli and responses as they enter into the relations among antecedents, behavior and consequences. With these preliminaries behind us, we are ready to move on to some classic experiments and findings in the study of learning.

Chapter 2
A Behavior Taxonomy

The word behavior, *like* habit, inhibit *and* ability, *is related to the Latin* habere, *to hold or have. The prefix* be *became attached in such words as the Old English* behabban. *As a word for how one held oneself, it was closer to the sense of comportment or demeanor than to the more contemporary sense of activity, just as* habit *was once more commonly what was worn than what was habitually done.*

A taxonomy *is a system of classification. The word is derived from the Greek* tassein, *to arrange, plus the Greek* nomia, *method. It shares its first root with the grammatical term* syntax, *with* tactic, *and with* taxis, *a kind of movement; it shares its second root with systematic disciplines such as* astronomy *and* economics, *with* metronome *and* autonomy, *and probably even with* number.

Observing Behavior
Presenting Stimuli
Arranging Consequences
Signaling Events and Procedures
  Signaling Stimulus Presentations
  Signaling Consequences
Establishing the Effectiveness of Consequences
Summary

In this chapter I outline a behavioral taxonomy, a vocabulary in terms of which the various procedures and phenomena of behavior can be organized. Please don't expect explanations of behavior in terms of formal laws. Instead, think in terms of a systematic classification of behavior based on its origins: given any example of behavior, where did it come from? The taxonomy will not be exhaustive, because we can't anticipate everything we'll run into in studying behavior. Like any science, behavior analysis and its taxonomy are and will remain works in progress. But we can at least aim for a descriptive system that organizes the phenomena we've begun to know while not excluding those we have yet to study. This field has evolved in many ways since I first encountered it. The several editions of this book provide some snapshots of its progress. I hope all of them, and especially this one, have been reasonably well focused.

We study behavior through experimental *procedures* or *operations*. We call the changes in behavior they produce behavioral *outcomes* or *processes*. We study the relation between environmental events and the organism's behavior by changing the environment and observing how this affects what the organism does. In the analysis of behavior, procedures or operations are *what the experimenter or the environment does or arranges*, and outcomes or processes are *the changes in behavior that result*. A convenient analogy comes from medicine, where the surgical operation, a procedure, is what the physician does to the patient, and the processes that follow, such as changes in circulation, respiration and so on, are the outcomes of the operation. In this text I'll usually favor *procedure* over *operation* and *outcome* over *process*, but within each pair the terms are somewhat interchangeable.

KEY TERMS: Taxonomy; Operation or Procedure; Process or Outcome; Reflex, Releaser, Fixed action pattern; Stimulus control, Discriminative stimulus, Signal, Occasion; Elicitation, Emission, Evocation; Establishing operation, Motivation, Deprivation, Satiation.

If you are familiar with the standard sections of experimental papers in psychology you will recognize this distinction as similar to that between the Method section, which describes the procedures and other details of an experiment, and the Results section, which describes the data obtained. We can also interpret changes in behavior once we have observed them, and interpretation corresponds most closely to what happens in the Discussion section of an experimental paper. In distinguishing between procedures, what gets done, and outcomes, what happens as a result, I'll emphasize experimental procedures rather than properties of the environment, mainly because this text is about research findings drawn from the laboratory.

But behavioral events occur in natural circumstances as well as in artificial circumstances arranged by an experimenter. We startle in response to an unexpected loud noise whether it was produced by accident on the street or deliberately in a laboratory. For that reason, these terms are often applied not only to artificial procedures and their outcomes but also to natural sequences of events. For example, Chesapeake, my cat, has learned that he often gets fed after going to his feeding dish in the kitchen. It hardly matters whether I deliberately trained his behavior or it came about as a matter of course during routine feedings. Either way, this was a procedure or operation in which being fed had become a frequent consequence of going to that place in the kitchen. It is convenient to call such circumstances procedures or operations whether or not they were explicitly arranged.

The simplest procedure, of course, is merely (1) to *observe behavior*: just watch. We then discover what an organism can do. But we have no control over events when we simply observe, so we may not be able to draw conclusions about the causes of behavior. To learn more we must intervene, and the simplest intervention is (2) to *present stimuli*. A more complicated intervention is (3) to arrange the environment so that it is changed by the organism's behavior or, in other words, so that the organism's behavior has *consequences*. Once behavior has consequences it may occur more or less often, and therefore consequential arrangements lead to the outcomes sometimes called *reinforcement* and *punishment*.

We haven't exhausted the possibilities. We can arrange things so that (4) stimuli *signal* the presentation of other stimuli or so that (5) stimuli *signal* the opportunity to produce consequences. We then speak of *signaling* or *stimulus-control* procedures; these can only occur in combination with one of the simpler procedures, either presenting stimuli or arranging consequences. Behavior may then depend on whether the signaling stimulus is present or absent. We must also consider procedures that (6) change the effects of the consequences of behavior, as when food becomes a more potent reinforcer after a period of food deprivation. Such procedures are called *establishing* or *motivating* operations, in that they establish the conditions under which consequences may become effective as either reinforcers or punishers.

Thus, the basic procedures are (1) observing behavior, (2) presenting stimuli, (3) arranging consequences for responses, (4) signaling stimuli, (5) signaling consequences and (6) establishing the effectiveness of consequences. These categories were not recognized all at once. They were gradually added to our taxonomy over the history of this science and have been refined in various ways as their effects on behavior have been more fully understood. For that reason, what follows surveys several classic experiments, not only illustrating these procedures in more detail but also introducing some researchers who played significant historical roles.

OBSERVING BEHAVIOR

What must we do to observe behavior? Interesting behavior depends on interesting environments. What would happen if we tried to move in the other direction, avoiding the contamination of behavior by the environment? For example, imagine fitting a rat with goggles to exclude visual stimuli and with ear plugs to exclude sounds, and setting up a ventilating system to remove odors. The rat can still touch things, including its own body, so we arrange a suit of hollow tubes that holds the rat's legs, reducing tactile contact at least for its paws. We may still not be satisfied, because the rat's weight produces pressure where the suit meets part of its body and so allows it to orient spatially. Is our next step to send the suited rat up to an orbiting space station, where gravity is eliminated? Yet even if we accomplished this much, what could we say about the rat's behavior? What could we see it doing?

The rat example is hypothetical, but experiments on sensory deprivation place humans in environments approximating the minimal stimulation we imagined for the rat. The problem is that in such environments, for human as well as rodent, there isn't much to do: no place to go and no one to talk to. Humans in such environments report a range of activities during their waking time, from thinking to hallucinating, but—no surprise—they sleep most of the time. So, to observe interesting behavior we have to observe organisms in interesting environments. Let's consider some examples. In early accounts of learning, speculations were typically based upon anecdotal evidence derived from simple observation, as in C. Lloyd Morgan's description of how his dog learned to get through a garden gate:

> The iron gate outside my house is held to by a latch, but swings open by its own weight if the latch be lifted. Whenever he wanted to go out the fox terrier raised the latch with the back of his head, and thus released the gate, which swung open.... How did he learn the trick? In this particular case the question can be answered, because he was carefully watched. When he was put outside the door, he naturally wanted to get out into the road, where there was much to tempt him—the chance of a run, other dogs to sniff at, possibly cats to be worried. He gazed eagerly out through the railings... and in due time chanced to gaze out under the latch, lifting it with his head. He withdrew his head and looked out elsewhere; but the gate had swung open.... After some ten or twelve experiences, in each of which the exit was more rapidly effected with less gazing out at wrong places, the fox terrier had learnt to go straight and without hesitation to the right spot. In this case the lifting of the latch was unquestionably hit upon by accident, and the trick was only rendered habitual by repeated association in the same situation of the chance act and the happy escape. Once firmly established, however, the behaviour remained constant throughout the remainder of the dog's life, some five or six years. (1920, p. 144)

Observing this behavior was perhaps a lucky accident, like the dog's lifting of the latch. More can be learned about learning by arranging the environment within which behavior will be observed. One researcher who did so was Wolfgang Köhler, one of the founders of gestalt psychology. Köhler studied the behavior of chimpanzees maintained from 1913 to 1917 at the Anthropoid Station on Tenerife, an island northwest of Africa (Köhler, 1927). Chimpanzees aren't native to Tenerife and the station was probably a front for German espionage, the observation of Allied shipping, during World War I (Ley, 1990). But Köhler's research does not appear suspect. In some of Köhler's experiments, bananas or oranges were placed in visible but inaccessible locations and the chimpanzees used materials within the area as tools to obtain the fruit. The following describes the behavior of the male chimpanzee, Sultan:

> The six young animals of the station colony were enclosed in a room with perfectly smooth walls, whose roof—about two metres in height—they could not reach. A wooden box..., open on one side, was standing about in the middle of the room, the one open side vertical, and in plain sight. The objective was nailed to the roof in a corner, about two and a half metres distant from the box. All six apes vainly endeavored to reach the fruit by leaping up from the ground. Sultan soon relinquished this attempt, paced restlessly up and down, suddenly stood still in front of the box, seized it, tipped it hastily straight towards the objective, but began to climb upon it at a (horizontal) distance of half a metre, and springing upwards with all his force, tore down the banana. About five minutes had elapsed since the fastening of the fruit; from the momentary pause before the box to the first bite into the banana, only a few seconds had elapsed, a perfectly continuous action after the first hesitation. (Köhler, 1927, pp. 39–40)

In many instances, of course, chimpanzees made fruitless attempts to solve such problems.

Köhler discussed these and other observations in terms of the chimpanzees' intelligence and insight. More important, chimpanzees were much less familiar in those days than now, and Köhler's readers were impressed by his descriptions what chimpanzees could do. But it was impossible to say from observation alone where the behavior came from. Did Sultan solve the problem because of some inherited cognitive disposition? Because the problem had features in common with some situation he had already encountered? Because he'd seen other chimpanzees whose behavior he could imitate? Because of some combination of these and other factors?

The term *insight* seemed appropriate to Köhler because of the suddenness with which a solution to a problem typically emerged. Problem solving that seemed insightful led to further questions: whether learning took place abruptly or gradually, and whether problem solving could be explicitly taught. Debates about how much nonhuman primates can learn continue to the present. But observation alone rarely identifies the sources of behavior and therefore rarely resolves such issues.

Strictly, of course, Köhler did much more than just observe behavior. He arranged environments within which to make his observations. Observation without intervention is difficult to achieve. To observe organisms successfully in the wild, one must know the possible effects of a human presence on their behavior. Even bringing an organism into captivity is itself an intervention. Observation alone is not good enough.

PRESENTING STIMULI

Köhler did in fact present stimuli, by arranging environments for his chimpanzees. So let's turn to procedures that examine the role of stimuli more directly. The following, by the ethologist Niko Tinbergen, describes the first feeding of newly hatched herring gull chicks:

> Sometimes the parent stands up and looks down into the nest, and then we may see the first begging behavior of the young. They do not lose time in contemplating or studying the parent, whose head they see for the first time, but begin to peck at its bill-tip right away, with repeated, quick, and relatively well-aimed darts of their tiny bills. They usually spread their wings and utter a faint squeaking sound. The old bird cannot resist this, and if only the chicks persist it will feed them. First the parent stretches its neck, and soon a swelling appears at its base. It travels upward, causing the most appalling deformations and the most peculiar turnings and twistings of the neck. All at once the parent bends its head down and regurgitates an enormous lump of half-digested food. This is dropped, and a small piece is now picked up again and presented to the chicks. These redouble their efforts, and soon get hold of the food, whereupon the parent presents them with a new morsel. Now and then the chicks peck at the food on the ground, but more often they aim at the parent's bill, and although this aiming is not always correct,

> it rarely takes them more than three or four attempts until they score a hit. (Tinbergen, 1960, p. 178)

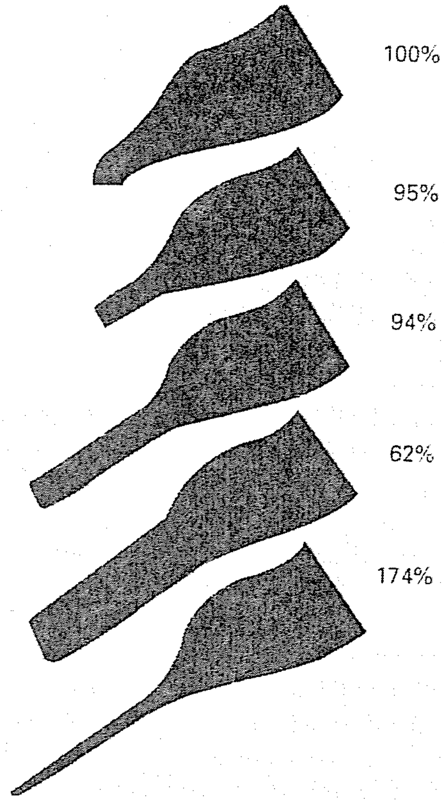

Figure 2–1 A series of models used to analyze which properties of the parent gull's beak produced begging pecks in the hatchling herring gull. Pecks to each model are expressed as a percentage of the reference level (100%) given by pecks to the top model. (Adapted from Tinbergen & Perdeck, 1950, Figure 15)

So far all we have are only some observations of chick behavior. But they involve the effects of stimuli and therefore prompt questions. What exactly are the critical features of these special stimuli presented by the parent gull? Are some more important than others? Are they the most effective ones possible? Tinbergen answered these questions by preparing stimuli that resembled the parent gull in various ways. He then measured the pecking generated when these stimuli were presented to recently hatched herring gull chicks. The herring gull parent has a white head and a yellow beak with a red patch near its tip. A beak with a black or blue or white patch produced less pecking than one with a red patch, but a beak with a patch of any color produced more pecking than one with no patch at all. Compared to the red patch, the color of the beak or the head was relatively unimportant in generating pecking. In fact, as long as the model had a beak with a red patch on it even the presence or absence of a head made little difference.

Tinbergen also varied the shape of the beak, as illustrated in Figure 2-1. Next to each stimulus, pecking is shown as a percentage of pecks generated by the normal beak shape at the top. With the red patch and other color differences eliminated, changes in pecking depended solely on changes in shape. Most models produced less pecking than the one with the normal shape. The model with an elongated beak (Figure 2-1, bottom) produced considerably more pecking than any of the others, including the one with the normal shape. Because of its effectiveness relative to the normal shape, Tinbergen called this one a *supernormal* stimulus, but he speculated that its shape might be more like what the chick first sees of its parent's beak, as it looks up from beneath the parent's head, than the profile view used for most of the other models.

Tinbergen varied other properties of the beak, including movement, slant and height above the ground. His analysis enabled him to construct a truly supernormal stimulus, a red pencil-shaped rod with three narrow white bands that generated more pecking than an accurate model of a herring gull head. More important, he could specify which features of the parent's head were important in generating pecks and which were unimportant. In other words, by presenting stimuli and observing their effects Tinbergen was able to identify the critical *structure* of the stimuli that generated pecking in the hatchling herring gulls.

Stimulus presentations are a common feature of research conducted by ethologists, whose concern is the evolution of species-specific behavior patterns in natural habitats. One effect of stimulus presentations, as we have just seen, is to produce responses. This outcome of presenting stimuli is called *elicitation*; the stimulus is said to elicit a response. In the language of ethology, the critical stimuli or stimulus features are called *releasers*, and the behavior they produce is called a *fixed action pattern*. But variations in vocabulary should not obscure the simplicity of the basic procedure of presenting stimuli.

The eliciting or releasing effects of stimuli can change over time. Data from the laughing gull chick provide an example. Feeding in the laughing gull chick differs in some details from that of the herring gull but includes the begging peck at the parent's beak followed by the parent's regurgitation of partly digested food that the chick then eats. The accuracy of the begging peck was tested by presenting beak models to chicks of various ages. Only about one-third of the pecks of newly hatched chicks struck the model, as opposed to more than three-quarters of those of two-day-old chicks (Hailman, 1969). Did the improved accuracy depend on changes in coordination or visual experience or other factors? Some behavior might be built in or *prewired* whereas other behavior might have to be learned. How do we tell which is which?

The consequences of accurately aimed pecks differ from those of poorly aimed pecks. In the laughing gull's natural habitat, a more accurately aimed peck is more likely to hit the parent's beak and therefore to be followed by the parent's regurgitation of food than a poorly aimed one. Accurately aimed pecks might increase relative to poorly aimed pecks because of their different consequences. Hailman's observations are consistent with that idea:

> If an inexperienced chick is too close to the target at first, its pecking thrust against the bill or model is so strong that the chick is thrown backward as much as an inch. If the chick starts out too far from the target, the pecking thrust misses and the chick falls forward as much as two inches. Older chicks rarely make such gross errors, suggesting that the experience of overshots and undershots has helped the chick learn to adjust its distance. (Hailman, 1969, p. 100)

To study such cases, it isn't enough simply to present stimuli. A more complex procedure must be arranged: stimuli must be arranged as consequences of the organism's behavior.

ARRANGING CONSEQUENCES

Our example this time is from research on animal intelligence by the American psychologist, Edward L. Thorndike. The crucial difference between Thorndike's research and Köhler's was that Thorndike systematically observed changes in behavior over many repetitions of behavior in a given setting rather than looking only at single instances of a problem solution. Thorndike noted gradual changes in behavior over repetitions rather than the abrupt changes typically reported by Köhler, perhaps in part because the problems he studied didn't lend themselves to sudden or insightful solutions. More important, Thorndike's experiments showed how responding often depends on its past consequences. Thorndike described his procedures this way:

> I chose for my general method one which, simple as it is, possesses several other marked advantages besides those which accompany experiment of any sort. It was merely to put animals when hungry in enclosures from which they could escape by some simple act, such as pulling at a loop of cord, pressing a lever, or stepping on a platform.... The animal was put in the enclosure, food was left outside in sight, and his actions observed. Besides recording his general behavior, special notice was taken of how he succeeded in doing the necessary act (in case he did succeed), and a record was kept of the time that he was in the box before performing the suc-

> cessful pull, or clawing, or bite.... If, on the other hand, after a certain time the animal did not succeed, he was taken out, but *not fed.* (Thorndike, 1898, pp. 5–6)

One of Thorndike's problem boxes is illustrated in Figure 2-2. In such devices, Thorndike studied cats, dogs and chicks. He gave the following description as typical of the behavior of most cats:

> When put into the box the cat would show evident signs of discomfort and of an impulse to escape from confinement. It tries to squeeze through any opening; it claws and bites at the bars or wire; it thrusts its paws out through any opening and claws at everything it reaches; it continues its efforts when it strikes anything loose and shaky; it may claw at things within the box.... The cat that is clawing all over the box in her impulsive struggle will probably claw the string or loop or button so as to open the door. And gradually all the other non-successful impulses will be stamped out and the particular impulse leading to the successful act will be stamped in by the resulting pleasure, until, after many trials, the cat will, when put in the box, immediately claw the button or loop in a definite way. (Thorndike, 1898, p. 13)

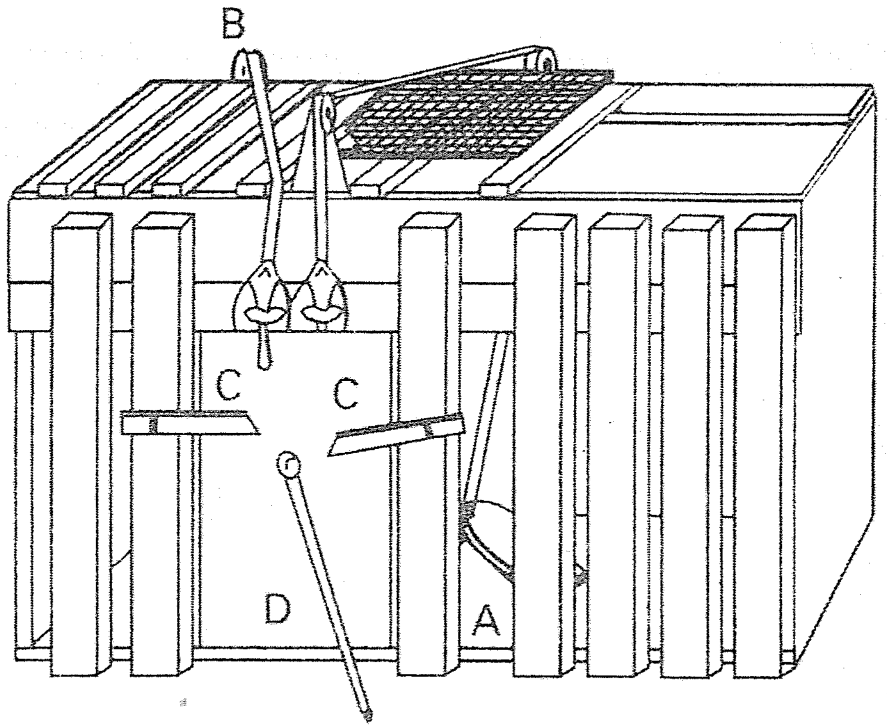

Figure 2-2 In most boxes that Thorndike (1898) used, the animal had only a single way to free the door. In the one shown, three different methods are illustrated: a treadle inside the box (A); a wire or string that can be reached from inside (B); and two outside latches that can be reached from inside (C). The door (D) was usually counterweighted so that it opened by itself once the animal made the appropriate response.

As a consequence of its behavior, the cat escaped from confinement and gained access to food. We can assume that both escaping and eating helped make the successful response gradually dominate over other, unsuccessful ones. In either case, the procedure cannot be reduced simply to the presentation of stimuli. A new part of the environment was not just presented to the cat; it became available as a consequence of what the cat did. As a result, the cat's behavior changed. Contemporary experiments often examine simpler responses in simpler situations but are similarly concerned with relations between responses and their consequences.

We arrange consequences for responses by constructing environments. If we place food in the goalbox of a maze, for example, we create an environment in which a consequence of a rat's movement from the startbox to the goalbox is finding food. After the rat has reached the food once, we can find out how this consequence affects its behavior by seeing what it does next time we put it in the startbox.

The consequences we arrange can vary from events of obvious biological significance such as presenting food or water to relatively minor changes in things seen or heard or touched. But not all consequences involve producing stimuli: responses can alter stimuli, as when turning a dimmer switch changes the brightness of a lamp; they can remove stimuli, as when operating a switch turns off a light; they can prevent stimuli, as when unplugging a lamp before repairing it eliminates the possibility of an electric shock; they can change the consequences of other responses, as when replacing a burned-out lightbulb makes the previously ineffective response of operating the light switch effective again.

Two classes of consequences are often distinguished on the basis of their effects on behavior. *Reinforcing* consequences are those that increase or maintain responding; *punishing* ones are those that decrease or suppress it. (It is also useful to have a term that doesn't prejudge whether the consequences will reinforce or punish. *Consequation* has been introduced for that purpose, though its usage is fairly uncommon. For example, if we don't know whether gold stars will reinforce the classroom behavior of a kindergarten child, it would still be appropriate to speak of *consequating* that behavior with gold stars (e.g., Powers & Osborne, 1976).

SIGNALING EVENTS OR PROCEDURES

When stimuli become effective as signals, we speak of them as *discriminative* stimuli. Procedures that involve signals are called *signaling* or *stimulus-control* procedures or operations. The presentation of stimuli and the arrangement of consequences seldom occur in isolation; they are often signaled by other events. A flash of lightning is often followed by a clap of thunder. A traffic light typically alerts drivers to possible consequences of proceeding or stopping at an intersection. These two examples illustrate that the signaling or discriminative effects of stimuli may be combined either with stimulus presentations or with consequential arrangements. Both demonstrate the signaling functions of stimuli, though the two types of signaling can have very different properties.

Signaling Stimulus Presentations

Stimuli that signaled the presentation of other stimuli were the basis for experiments on conditional or conditioned reflexes by the Russian physiologist, Ivan P. Pavlov. Pavlov (1927) studied how stimuli acquired signaling properties by showing that responses to stimuli such as food were sometimes produced by other stimuli that had reliably preceded the food. Pavlov spoke of the effects of food in a dog's mouth in terms of the alimentary reflex (for Pavlov, its components included both the glandular response of salivating and motor responses such as chewing and swallowing). He concentrated on salivation because the technology available to him made salivating easier to measure than motor responses. Through surgery, he brought the duct of one of the dog's salivary glands to the outside of the dog's cheek, where he connected it to a fluid system that allowed drops of saliva to be counted.

For one dog, the sound of a metronome consistently preceded food presentations. Pavlov gave the following account of the conditions necessary to make a stimulus function as a signal:

> On several occasions this animal had been stimulated by the sound of the metronome and immediately presented with food—i.e., a stimulus which was neutral of itself had been superimposed upon the action of the inborn alimentary reflex. We observed that, after several repetitions of the combined stimulation, the sounds from the metronome had acquired the property of stimulating salivary secretion and of evoking the motor reactions characteristic of the alimentary reflex.... Hence a first and most essential requisite for the formation of a new conditioned reflex lies in a coincidence in time of the action of any previously neutral stimulus with some definite unconditioned stimulus. Further, it is not enough that there should be overlapping between the two stimuli; it is also and equally necessary that the conditioned stimulus should begin to operate before the unconditioned stimulus comes into action. If this order is reversed, the unconditioned stimulus being applied first and the neutral stimulus second, the conditioned reflex cannot be established at all. (Pavlov, 1927, pp. 26–27)

Pavlov's conditioning experiments demonstrated how a signaling procedure can be superimposed on the simpler procedure of presenting stimuli. We'll discuss his procedures in more detail in another chapter.

Signaling Consequences

Instead of signaling the presentation of stimuli, a stimulus may signal when responses will have consequences. The signaling of consequences played an important role in the history of the study of learning long before it began to be studied experimentally. For example, it was involved in the analysis of the case of Clever Hans, a horse that seemed to have been taught to solve arithmetic problems (Pfungst, 1911). The horse apparently took not only addition and multiplication but also square roots in its stride.

> The visitor might walk about freely and if he wished, might closely approach the horse and its master, a man between sixty and seventy years of age. His white head was covered with a black slouch hat. To his left the stately animal, a Russian trotting horse, stood like a docile pupil, managed not by means of the whip, but by gentle encouragement and frequent reward of bread or carrots.... Our intelligent horse was unable to speak, to be sure. His chief mode of expression was tapping with his right forefoot. (Pfungst, 1911, pp. 18–19)

Clever Hans gave his answers by the number of times he tapped. His performance was investigated by Oskar Pfungst, who discovered that the horse performed accurately only in the presence of his master. Furthermore, Clever Hans knew the answers only if his master knew them too. Pfungst therefore turned his attention from the horse to the master, and determined that the horse was responding to subtle cues provided by the master's behavior.

> we sought to discover by what movements the horse could be made to cease tapping. We discovered that upward movements served as signals for stopping. The raising of the head was most effective, though the raising of the eyebrows, or the dilation of the nostrils—as in a sneer—seemed also to be efficacious.... On the other hand, head movements to the right and to the left or forward and back... remained ineffective. We also found that all hand movements, including the "wonderfully effective thrust of the hand into the pocket filled with carrots," brought no response. (Pfungst, 1911, p. 63)

The master had cooperated in the investigation and there was no evidence that he had been aware of the signals he had provided. Pfungst noted that

> Hans's accomplishments are founded... upon a one-sided development of the power of perceiving the slightest movements of the questioner....We are justified in concluding from the behavior of the horse, that the desire for food is the only effective spring to action.... The gradual formation of the associations mentioned above, between the perception of movement and the movements of the horse himself, is in all probability not to be regarded as a result of a training-process, but as an unintentional by-product of an unsuccessful attempt at real education. (Pfungst, 1911, pp. 240–241)

In this case, the master's movements provided the stimuli in the presence of which the horse's taps were followed by food. The case of Clever Hans demonstrates that even very subtle properties of stimuli can signal the consequences of responding, and it is often cited as *The Clever Hans Effect* when critics wonder whether sophisticated behavior depends on cues unwittingly provided by participants.

Facilitated communication, for example, was an attempt to provide therapy for nonverbal children on the autism spectrum. The children's hands were guided on a typewriter keyboard by facilitators, individuals who were supposed to provide emotional support and motor help. Even though the children did not speak, some soon seemed to be producing complex typed messages. When it was discovered that the children could answer questions correctly only when the facilitators knew the questions, however, it became clear that facilitated communication was a modern Clever Hans phenomenon (Montee, Miltenberger, & Wittrock, 1995). The facilitators had been actively guiding the typing, though in many cases like Clever Hans' master they were unaware that they were doing so.

Signaling effects such as those displayed by Clever Hans were eventually examined more systematically. They came to be called the discriminative functions of stimuli, and differed in many ways from the kinds of signaling functions that had been studied by Pavlov. The research that most decisively established the distinction was conducted by the American psychologist, B. F. Skinner, who arranged an environment in which a rat's lever presses produced food when a light was on but not when it was off, as illustrated in the following passage:

> [The apparatus] consists of a dark, well-ventilated, sound-proofed box...containing... a horizontal bar, made of heavy wire, which may be pressed downward approximately 1.5 cm. against a tension of 10 grams. As the lever moves downward, a mercury switch directly behind the wall is closed. We are concerned with the response of the rat in pressing this lever, which we may define as any movement by the rat which results in the closing of the switch. The switch operates a food-magazine, which discharges a pellet of food of standard size into the tray, where it is accessible to the rat. The connection between the lever and the magazine may be broken at will by the experimenter.... The only additional requirement for the investigation of a discrimination is an extra source of stimulating energy... a small (3 c.p.) electric bulb.... The experimenter controls the current to the light and the connection between the lever and the magazine in such a way that the response to the lever-plus-light is always followed by the discharge of a pellet of food into the tray, while the response to the lever alone is never so reinforced. The animal eventually learns to respond to the lever when the light is on but not to respond when the light is off. (Skinner, 1933, pp. 304–305)

In this example, the light signals the consequences of pressing the lever, in that the lever press is reinforced in the presence but not the absence of the light. The light is a *discriminative* *stimulus*, and the rat comes to press the lever more often when it is on than when it is off. As the rat begins to respond differently in the presence than the absence of the light, its behavior is said to come *under the control* of the light as a discriminative stimulus; the light is also said to *occasion* the behavior. The development of this differential responding has been called *discrimination learning* (this is a technical usage, but later we'll discuss how it is related to discrimination in social behavior, as in racial profiling and in other differential treatments of people along dimensions such as gender and ethnicity).

The relations between a discriminative stimulus and the consequences of responding are elaborated in the following passage by Skinner (the term *operant* refers to a class of responses having certain consequences, and the term *reinforcement* refers to these consequences):

> the operant must *operate* upon nature to produce its reinforcement. Although the response is free to come out in a very large number of stimulating situations, it will be effective in producing a reinforcement only in a small part of them. The favorable situation is usually marked in some way, and the organism... comes to respond whenever a stimulus is present which has been present upon the occasion of a previous reinforcement and not to respond otherwise. The prior stimulus does not elicit the response; it merely sets the *occasion* upon which the response will be reinforced.... Three terms must therefore be considered: a prior discriminative stimulus (S^D), the response (R^O), and the reinforcing stimulus (S^1). Their relation may be stated as follows: only in the presence of S^D is R^O followed by S^1. (Skinner, 1938, p. 178)

Skinner explored this three-term relation experimentally with lights as discriminative stimuli, rats' lever presses as responses, and food pellets as reinforcing consequences, but his next example, reaching for and touching objects in the seen environment, illustrates the broad range of situations to which the relation can be applied.

> A convenient example is the elementary behavior of making contact with specific parts of the stimulating environment. A certain movement of my arm (R^O) is reinforced by tactual stimulation from a pencil lying on my desk (S^1). The movement is not always reinforced because the pencil is not always there. By virtue of the visual stimulation from the pencil (S^D) I make the required movement only when it will be reinforced. The part played by the visual stimulus is shown by considering the same case in a dark room. At one time I reach and touch a pencil, at another time I reach and do not.... In neither the light nor the dark does the pencil *elicit* my response (as a shock elicits flexion), but in the light it sets the occasion upon which a response will be reinforced. (Skinner, 1938, p. 178)

The three-term relation, discriminative-stimulus—response—consequence, will be a recurrent theme. Each term is critical. Their combination distinguishes them from other, simpler behavioral relations. In the Pavlovian situation, for example, in which a stimulus is signaled, the organism's behavior has no effect on the sequence of events; no consequences are arranged for responses.

Consider the earlier examples. My blinking or startling at the lightning flash won't prevent the subsequent clap of thunder. But if a traffic light is red as I approach an intersection, my stepping on the brakes is occasioned by this stimulus only because I have learned the potential consequences of doing or not doing so. The second of these two examples is the only one that involves all of the terms of Skinner's three-term contingency. An important difference in vocabulary accompanies these distinctions: (i) **when a stimulus is the primary cause of a response, we say that the stimulus elicits the response or that the response is elicited**; but (ii) **when a response occurs in the presence of a stimulus because the stimulus signals some consequence of responding, we say that the stimulus occasions the response and that the response is emitted**.

Early animal experiments were often concerned not so much with the nature of discrimination learning as with the sensory capacities of organisms. For example, rodent vision was studied by arranging two paths only one of which led to food (Yerkes & Watson, 1911). Where a rat had to choose between left and right, two stimuli were presented (e.g., a black card and a white card). The path to food varied from left to right but was always indicated by the same card (e.g., black). Once the rat learned to take the path indicated by the card correlated with food, the limits of its vision could be studied by substituting other cards for the original pair (e.g., light and dark grays). Such experiments were laborious; demonstrating discrimination learning might take hundreds of trials if the rat learned at all. Several problems existed in this type of study, not least of which was ensuring that the rat looked at the cards when it reached the choice point.

Apparatus improved over time. Figure 2-3, for example, shows the jumping stand developed by Karl S. Lashley (1930). Lashley described its advantages:

> it requires the animal to jump against the stimulus patterns from a distance, instead of to run past them.... I have usually trained the animals by placing the stand against the screen and allowing the animals to step through the open holes to the platform, then gradually withdrawing the stand until, in ten or fifteen trials, the distance of 25 cm. is reached. Cards are then placed in position and training in discrimination begun. (Lashley, 1930, pp. 454–457)

In Lashley's apparatus, rats typically learned to discriminate black from white with perfect accuracy within four or five trials, and even more difficult discriminations, such as vertical versus horizontal, could usually be mastered within less than 50 trials.

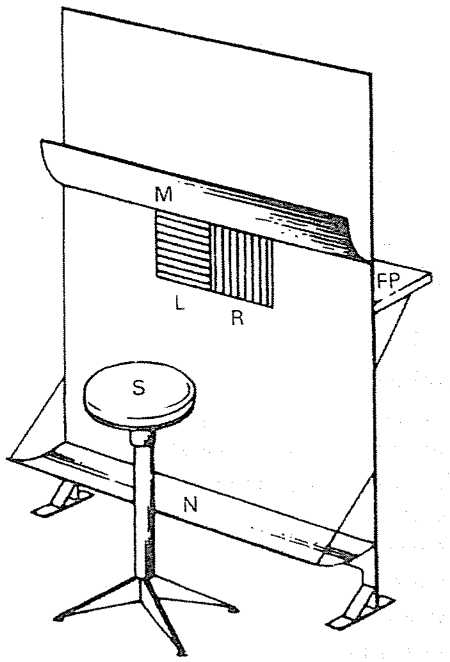

Figure 2-3 The Lashley jumping stand (Lashley, 1930, Figure 1). A rat was trained to jump from the stand (S) to one of two doors (L and R). If it jumped to the correct door, the door gave way and the rat reached the food platform (FP). If it jumped to the incorrect door, the door remained fixed and the rat fell into the net below (N). The projecting metal sheet (M) prevented the rat from jumping too high. In the illustration, the right door (R) would be correct for a rat being trained to jump toward vertical lines.

These cases in which discriminative stimuli signal consequences are more complex than the example in which a rat's lever presses produced food in the presence but not the absence of light. There our concern was only with how often presses occurred when the light was on and when it was off. The jumping stand, however, seems to involve just two responses and their consequences: jumping toward vertical and finding food and jumping toward horizontal and landing in the net. But left and right are not irrelevant to the rat. The situation involves at least four responses, each with its own consequence: jumping to vertical on the left, to vertical on the right, to horizontal on the left and to horizontal on the right. Rats might respond on the basis of position, left or right, rather than on the basis of stimulus cards. For example, if the first three trials of vertical-horizontal training were set up with vertical on the right, as in Figure 2-3, it would be no surprise if on trial 4, with vertical on the left for the first time, the rat jumped right, toward horizontal. Until trial 4, jumping to the right led to food just as reliably as jumping toward vertical.

A discrimination in which stimulus conditions alternate, as in the lever-pressing example, is called a *successive* or *go–no go* discrimination. One in which two or more stimuli are present at the same time and in which each is correlated with a different response, as in the jumping-stand example, is called a *simultaneous* discrimination. Both illustrate signaling procedures superimposed on consequences of responding. Such procedures come in varying degrees of complexity.

ESTABLISHING THE EFFECTIVENESS OF CONSEQUENCES

Some consequences of behavior are more important than others and their effectiveness can vary over time. For example, water is likely to be an effective reinforcer if you've been deprived of water for a long time, but it is less likely to be so if you just drank a lot of water. The things that can be done to change its effectiveness are called *establishing* or *motivating operations*. *Deprivation* and *satiation* are two examples but are not the only possibilities. For example, heavy exertion in a hot and dry climate or a mouthful of very salty food can have the same effect as a period of water deprivation. Establishing operations change the effectiveness of consequences by changing the likelihood of behavior, as illustrated in the following passage from B. F. Skinner:

> the probability of drinking becomes very high under severe water deprivation and very low under excessive satiation.... The biological significance of the change in probability is obvious. Water is constantly being lost through excretion and evaporation, and an equal amount must be taken in to compensate for this loss. Under ordinary circumstances an organism drinks intermittently and maintains a fairly steady and presumably optimal state. When this interchange is disturbed—when the organism is deprived of the opportunity to drink—it is obviously important that drinking should be more likely to occur at the first opportunity. (Skinner, 1953, p. 141–142)

The vocabularies of motivation and establishing operations have become somewhat interchangeable. Skinner also spoke of these phenomena in terms of drives:

> The term is simply a convenient way of referring to the effects of deprivation and satiation and of other operations which alter the probability of behavior in more or less the same way. It is convenient because it enables us to deal with many cases at once. There are many ways of changing the probability that an organism will eat; at the same time, a single kind of deprivation strengthens many kinds of behavior. (Skinner, 1953, p. 144)

Behavior that occurs because of establishing operations is sometimes said to be *evoked*. Skinner pointed out, however, that the effects of establishing operations must not be equated with those of stimuli:

> A common belief is that deprivation affects the organism by creating a stimulus. The classic example is hunger pangs. When an organism has been without food for a sufficient time, the contractions of the stomach stimulate it in a characteristic way. This stimulation is often identified with the hunger drive. But such stimulation is not closely correlated with the probability of eating. Hunger pangs are characteristic of only a small part of the range through which that probability varies continuously. We usually eat our meals without reaching the condition of deprivation in which pangs are felt, and we continue to eat long after the first few mouthfuls have stopped any pangs which may have occurred. (Skinner, 1953, pp. 144–145)

As Skinner's example indicates, effects of discriminative stimuli must be distinguished from those of establishing operations. Another relevant term is *incentive*. It is often defined as a motivator or incitement to action, but it usually implies some stimulus correlated with its availability. Food pellets that a rat has not yet seen or eaten may serve as reinforcers but they are not yet incentives. Once the rat has produced and eaten some, however, their sight or smell or predictable availability will function as incentives.

The most straightforward way to distinguish between consequential procedures and establishing or motivational ones is to consider whether the consequences of a response change or stay the same. Consider a flashlight (cf. Michael, 1989). It lights when I press the button that turns it on. It does so whether I press the button in light or in dark, but turning it on matters to me only when it is dark. Thus, changes from outdoor daylight to darkness or from indoor lamplight to the darkness of a power outage are examples of establishing operations with regard to whether I'm likely to turn on the flashlight. In each case, something happened that made it important to turn on the flashlight, but I could have turned it on even if those events hadn't occurred.

If my flashlight battery goes dead, however, pressing the button that usually turns it on no longer does anything. The consequences of pressing the button have changed. It used to work. Now it doesn't. Thus, the dying of the battery is not an establishing operation. It is a consequential operation: it changes whether my button press will be reinforced by the onset of light.

But the dead battery may be establishing or motivational in a different way. It might not have mattered to me before, but now finding a fresh battery has become important. Once I find one and replace the dead battery, my flashlight works again. In other words, the battery going dead had two effects at the same time: it had a consequential effect, because it changed what happened when I tried to turn on the flashlight, but it also had an establishing or motivational effect, because it made finding a fresh battery important. And if I cannot find a fresh battery, I might start looking for candles and matches; not everything that becomes established as a potential reinforcer will necessarily be available when it is needed.

Establishing operations and consequential operations work together. Usually we can't have one without the other, but it is important to be clear about which behavior is related to each. In these examples, turning on the flashlight was a response with consequences, but the light versus dark conditions established whether it was important for me to turn it on; similarly, when the battery went dead, replacing the battery was a response with consequences, but the failure of the flashlight to work established whether it was important for me to change the battery. Another point of this example is that these procedures encompass far more than the physiological effects such as hunger and thirst implied by the language of drive or motivation.

In contrast with the language of stimulus control, in which changes in discriminative stimuli are said to occasion responses, the responding said to be evoked by establishing operations may occur in a relatively constant environment. For example, if you are more likely to check the refrigerator a long time after than right after a meal, it is your behavior and not the refrigerator that changes from one time to another. In this case, checking the refrigerator is said to be evoked by the establishing event, food deprivation. And when we want to talk about such behavior without mentioning either the establishing operations or the discriminative stimuli, it remains okay to say that the behavior was emitted. In other words, establishing or motivational procedures *evoke*, discriminative stimuli *occasion*, and the responses they evoke and/or occasion *are emitted*.

SUMMARY

The study of behavior is concerned with relations between environmental events, *stimuli*, and the organism's actions, *responses*. We can examine these relations by analyzing how changes in the environment produce changes in responding. A critical first step is to *observe* behavior, but just watching is typically not enough. To understand behavior we must intervene by changing the environment. We describe environmental changes in terms of classes of experimental procedures: *presenting stimuli, arranging consequences, signaling stimuli or consequences*, and *arranging establishing operations*. We'll return to these procedures often throughout this book, and they will especially help us to organize the topic of learning without words in Part III. They are summarized in Table 2-1.

Behavior is complicated. Different stimuli can have different effects on different responses, and different responses can have different consequences. Nevertheless, as we will see, a wide range of learning situations can be treated in terms of combinations of these basic types of experimental interventions. As our behavioral taxonomy these categories will take us a long way. But later, especially when we consider the transition from learning without words to learning with words, we will find that they do not exhaust the possibilities.

Table 2-1 Behavioral Procedures

Procedure	*Description*	*Examples*	*Usage*
1. Observation	No intervention.	We watch an animal behave.	---
2. Stimulus-presentation procedure	Stimulus *A* is presented.	*Loud noise (A)* startles child. Physician shines *light (A)* in patient's eye.	Stimulus *elicits* response; response *is elicited by* stimulus.
3. Consequential procedure	Response *B* has consequence *C* (e.g., a stimulus is produced or terminated).	*Putting money in vending machine (B)* produces *soft drink (C)*. *Touching hot stove (B)* produces *burn (C)*. *Light goes out (C)* when *switch is thrown (B)*.	Response is *emitted*.
4. Signaling or stimulus-control procedure: Superimposed on stimulus presentation	Stimulus *D* signals presentation of stimulus *E*.	*Lightning (D)* precedes *thunder (E)*.	Stimulus *elicits* response; response *is elicited by* stimulus.
5. Signaling or stimulus-control procedure: Superimposed on consequences	Stimulus *F* signals that response *G* will have consequence *H*.	*Red traffic light (F)* signals that *driving through intersection (G)* may lead to *traffic ticket (H)*. *Ringing telephone (F)* signals that *answering (G)* may provide *opportunity for conversation (H)*.	Stimulus *occasions* response; response *is emitted* in presence of stimulus.
6. Establishing operation	Effectiveness of consequence *I* as a reinforcer or punisher is established.	*Food (I)* becomes an effective reinforcer after food deprivation. The presentation of shock makes *shock removal (I)* a reinforcer. When it is important to unlock a door, *the key to the door (I)* becomes a reinforcer.	An event is *established* as a reinforcer or punisher. Behavior is *evoked* by the establishing operation.

PART II
BEHAVIOR WITHOUT LEARNING

Chapter 3
Evolution and Development

Evolution *and* revolution *are descendants of the Latin* volvere, *to roll; they differ in that evolution implies an unrolling or rolling out whereas revolution implies a rolling over or turning around. They appear unrelated to* develop, *perhaps of Celtic origin via* des- *plus* voloper, *an unwrapping or unfolding.* Selection *can be traced to the Latin* legere, *originally to gather or to choose; via* logos *it is a relative of* logic *and* lexical. *The prefix,* se, *adds the implication of a weeding out from a large number, as contrasted with the bringing together implied by* con, *the root prefix for* collection.

Phylogeny, *evolutionary history, and* ontogeny, *the life history of the individual organism, share the Greek root* gen, *in the sense of kind or sort, through which both are related to* generalization. Phylo- *has a Greek root implying a tribe or clan or racial stock and* onto *has one implying being or reality. In their combination with* gen, *each implies origin: the origin of a biological phylum or population or the origin of a living entity.* Memetic *is derived from a 1976 coinage by Dawkins,* meme, *a term that echoes both* gene *as in genetics and* mime *as in mimicry.*

It has not been many human lifetimes since Copernicus and Galileo and Kepler pressed their cases that the earth was part of a system of planets orbiting our sun rather than being at the center of things. However humbling that displacement may have seemed at the time, it pales in comparison with the twentieth-century discoveries that our sun is just one among hundreds of billions of stars in our galaxy and that our galaxy is just one among hundreds of billions of galaxies in our universe.

The cumulative evidence has converged on an estimate for the age of our universe of roughly fourteen billion years. Our galaxy formed within the first billion years or so and our sun ten billion years or so later (e.g., Marochnik & Suchkov,

KEY TERMS: Evolution; Variation and Selection; Recipe versus Blueprint; Phylogeny and Ontogeny; Development and Evo Devo; Modules, Compartments and Boundaries; Natural Selection, Sexual Selection, Operant Selection, Cultural or Memetic Selection.

1995). Our planet formed not long after. The earth is about four and a half billion years old. That's a very long time. If I tried to count to a billion, I'd take more than thirty years to finish even if I kept up an uninterrupted count of one number per second. If I took time out to sleep, of course, I'd take much longer. The one-per-second estimate, by the way, is very generous: I can easily manage with small numbers, but surely I'd slow down when I got to bigger ones like 9,275,078, especially if I was worrying about losing count.

Life existed on earth for most of those years (Fortey, 1998, provides a detailed account). Chemical and fossil evidence indicates that it began within the first billion years or so. Over most of the next three billion years it consisted of single-celled organisms. Some of those single-celled organisms congregated, but multicellular organisms in which different cells had different functions only appeared roughly 550 million years ago, in the geological period called the Cambrian. An explosion in the diversity of multicellular life during that period provided the major groupings from which contemporary species evolved. One of these was the vertebrates. The evolution from fish to amphibians to reptiles included many significant events, such as the colonization of land. The dinosaurs were a spectacular part of the story, but by 65 million years ago they were gone, surviving only in those forms that have since evolved into birds. The passing of the dinosaurs made room for the evolution of mammals, and by perhaps four million years ago primates that walked upright had evolved. We humans eventually emerged from that hominid line only a little more than 100,000 years ago. We are all cousins, all the descendants of a very long line of survivors.

CHAOS THEORY AND DARWIN'S BUTTERFLY

In the face of these vastnesses of time and space it is all too easy to conclude that we are insignificant. But if nothing else, our planet is a place that has produced organisms who have begun to see how the world in which they live works. We are those organisms. Writing on the issue of God having created humans from mud, Kurt Vonnegut put it this way: "And I was some of the mud that got to sit up and look around. Lucky me, lucky mud" (Vonnegut, 1963, ch. 99). His point is not novel: "Our word 'human' comes from the proto-Indo-European root *dhghem*, meaning simply 'earth'" (Thomas, 1992, p. 19). So let's look around. There is plenty to marvel at.

In those vastnesses the odds *were* against us, individually as well as collectively. Charles Darwin can give us a sense of how much so. Butterfly collecting was one of Darwin's many interests, and it is safe to assume that in his youth in the summer of 1828 in North Wales he caught a butterfly. Actually, he almost certainly caught more than one, but the capture of any butterfly by Darwin or by any other butterfly collector during the first half or so of the nineteenth century would suit our purposes. Whichever capture we consider, all our lives hung upon it.

This conclusion follows from what in Chaos Theory is called the Butterfly Effect (Gleick, 1987). Chaos Theory is derived from the mathematics of nonlinear systems, which involves recursive computations, computations in which the output of an equation serves as the input for its next iteration. For example, if x_n begins as .50 in the equation $x_{n+1} = 3x_n(1-x_n)$, it first becomes .75, then substituting this as the new x it becomes .56, then .74, and .58, and .73, and so on. Many natural phenomena, including the weather, are best described in terms of such systems. One significant property of some of them is that they are drastically affected even by very tiny changes in initial value. Mathematical models for predicting the weather made significant contributions to Chaos Theory (Lorenz, 1963). The Butterfly Effect refers to the finding that when predicting weather patterns using models that incorporate nonlinear equations, the entry of initial values differing by as little as the energy produced by the flap of a butterfly's wings can alter the prediction of the direction in which a storm system will move some weeks from now.

Now consider the implications of Darwin's capture of that butterfly. If it had remained free to flap, weather patterns throughout the world, little by little, would have begun to deviate from those in our own history. We could imagine their impact on major historical events. For example, the battle of Gettysburg would probably have gone differently had the weather been different. And even if Darwin had set sail on the *H.M.S. Beagle* in that world as in ours, his voyage might have come to a different conclusion.

But for us there would have been an even more profound effect, because each of us is the product of a particular union of sperm and egg. Would that particular union have come to pass—would any of us have ever come into existence—if our parents' act of procreation had occurred at a somewhat different time or place? Almost certainly not. A lovely sunset where in our world it was overcast; a meeting at an agreed time where in our world a delay was caused by rain; a wedding forced indoors by storm where in our world it occurred in open air; illness following from rain and damp where in our world a parent remained healthy. Spun out over days and weeks and months, these alternative circumstances would very soon extend to every person on our planet.

Long before the twentieth century, things would have differed so much that no individual now alive would have been conceived in that other world. No doubt some individuals would have been given the same names as their existing counterparts, but all would have been different, both in their genetic endowments and in their upbringings. Would an Igor Stravinsky in that world have composed *The Rite of Spring*? Would a Pablo Picasso have painted *Guernica*? Would an Albert Einstein have formulated the Theory of Relativity? Of those who determined the fates of so many in our world, no Roosevelt or Hitler or Churchill or Stalin or Gandhi or Mao. Others in their places and in ours, but neither you nor me. Had that butterfly evaded Darwin's net, we would not be here. Lucky we. And if our being here now sounds improbable enough, think of all the flaps of all the butterflies across all the lepidopteran millenia of our world.

Now some might argue that a single butterfly couldn't have made all that difference, and that whether it flapped or not would have been washed out by other events. But if effects on the weather seem too counter-intuitive to be persuasive and the butterfly effect seems too metaphorical, consider that in taking flight in one direction or another that butterfly might have led Darwin to trip when he did not trip in our world, perhaps resulting in a sprained ankle and so on through a cascade of events ending with someone other than Darwin eventually taking his place on the deck of the *H.M.S. Beagle*. A world without his revelatory trip and his momentous book would have been different enough for the purposes of our example.

THE NATURE OF EVOLUTION

My interest in evolution began early. I was a preschooler when my parents took me to see Walt Disney's animated film, *Fantasia*, not long after its 1940 premiere. Igor Stravinsky had written *The Rite of Spring* as a ballet evoking primitive human rites, but Disney's animations for it instead began with the volcanic spasms of an ancient earth, followed by the emergence of life in the sea and its movement onto land, and culminating in prehistoric scenes dominated by the dinosaurs. The dramatic entrance of Tyrannosaurus Rex scared me at the time, but it was not too long afterward that it was a treat to be taken to see actual dinosaur skeletons in the exhibit halls of the Museum of Natural History in New York City. All of the segments of *Fantasia* were cartoons but on that screen the dinosaurs had a reality not shared by Mickey Mouse as the Sorcerer's Apprentice or the pirouetting hippopotami of *The Dance of the Hours*. I was hooked. But it is one thing to have one's interest aroused and another to grapple with real events rather than human creations such as movie cartoons. What kept me hooked was the coherent and persuasive science I later encountered.

We live in a time when some still think evolution is a topic from which schoolchildren must be protected. Evolution is of course hardly the only example of reluctance to face facts. For example, the evidence for the Nazi Holocaust of World War II remains overwhelming. It happened during my lifetime and I vividly recall the newsreels of the time. But when they came home from the war, many of the US troops who had participated in the liberation of the concentration camps didn't talk much about what they had seen (Hirsh, 2010). Despite the overwhelming evidence, some of those born too late to remember it came to deny that the Holocaust had ever happened. Talk is too often unconstrained by reality. Like everything else we do, human talk is behavior, and as we shall see when we later consider verbal behavior, it is too often shaped more by social consequences than by actual events. Anyway, despite the naysayers, evolution happens.

Evolution has operated throughout the history of life on our world. It is rapid enough that it has been observed within individual human lifetimes (Weiner, 1994). It occurs in natural habitats, such as the Galapagos Islands, where different species of finches continue to evolve with changes in the local habitats on each island. It also occurs as a result of human interventions, for example when disease organisms become resistant to antibiotics or insect pests become resistant to insecticides. Evolution is the foundation of the entire discipline of contemporary biology: "Biology without evolution is like physics without gravity" (Carroll, 2005, p. 294). We are only beginning to reap the fruits of this expanding science, ranging from the legal ramifications of DNA testing to a broad spectrum of medical applications.

The fossil record is inevitably incomplete, but many lines of evidence from biology and geology and other disciplines have decisively shown that evolution has happened and continues to happen. In other words, **evolution is not a theory; it is a name for certain kinds of changes that happen to the biological populations we call *species*.** Theories of evolution are not about whether contemporary species are descended from the very different ancestors we find in the geologic record. All theories of evolution take that for granted. They differ in what they say about how evolution came about. The theory that has been most successful in accommodating the facts of evolution is Charles Darwin's account in terms of natural selection. *Natural selection* refers to Darwin's account of evolution in terms of the differential survival and reproduction of the members of a population; the environment selects the individuals who pass on something of themselves from one generation to the next and it thereby shapes the characteristics of those in later populations (Carroll, 2005, 2006; Dawkins, 1976, 1986). Evolution by natural selection requires variations within populations; these variations are the stuff upon which selection operates.

Selection was well-known even before Darwin but was the sort used by humans in horticulture and animal husbandry. People knew how to breed plants or livestock selectively for hardiness or yield or other characteristics. This selective breeding was called *artificial selection*, and it created new varieties of vegetables and flowers and so on. Workhorses were selected for strength and racehorses were selected for speed. One part of Darwin's insight was that a similar kind of selection occurred in nature, without human intervention; that was *natural selection*. Darwin's main arguments were first published in his book, *On the Origin of Species* (Darwin, 1859). They were warmly received in some quarters but in others were strongly resisted. The resistance grew and by the end of the nineteenth century the belief was widespread that Darwinism was dead. It didn't recover until well into the twentieth century. The half century or so that preceded that recovery has been called the eclipse of Darwinism (Bowler, 1983; cf. Catania, 1987).

The reason for the eclipse wasn't simply that evolution itself had temporarily been discredited but rather that other theories than Darwin's became dominant. The main alternatives to Darwin's *natural selection* were *Lamarckism*, *orthogenesis*, and *Mendelian genetics*. Lamarckism was based on the work of an eighteenth-century French scientist who in his time had done much to make a case for the fact of evolution (Jordanova, 1984). Lamarck's guess was that characteristics acquired during an organism's lifetime could be passed on to its offspring, through changes in its own genetic material or germ plasm. One problem with his theory was that it couldn't show why advantageous acquired characteristics should be any more likely to be passed on than disadvantageous ones such as injuries.

According to the theory of orthogenesis, evolution was a developmental unfolding directed by forces within organisms, without reference to demands of the environment. One of its manifestations was supposed to be the recapitulation of phylogeny by ontogeny. *Ontogeny* is the development of the individual organism and *phylogeny* is its evolutionary history. During ontogeny the embryo was thought to pass through stages corresponding to its phylogeny. But the idea of recapitulation, based on only superficial properties of embryos, is no longer central to evolutionary theory (Gould, 1977).

The work of Gregor Mendel (Henig, 2000) provided a crucial foundation for genetics, but its problem was that it provided no mechanism for variation. In strict Mendelian descent, dominant and recessive genes in one generation determined their proportions in the next. Without variation, natural selection had nothing to work on. To provide for the appearance of new forms, later Mendelian accounts added mutation theory, which held that evolution proceeded through spontaneous and usually large genetic changes. Mutations have since played a pivotal role, but too little was known about them at that time to support a convincing account.

Recipes and Blueprints

In the nineteenth century, genes were theoretical entities. The techniques of cell biology hadn't yet reached the point at which genes had been identified in actual cells; the discovery of the structure of DNA (Watson & Crick, 1953) would follow nearly a century after Darwin's revolutionary book. Nevertheless, all of these evolutionary theories assumed that hereditary material of some sort was passed on from one generation to the next and that evolution was determined by the properties of this material. A recurrent major flaw was the assumption that genetic material constituted a representation or copy of the organism. In the earliest versions of orthogenesis, called *preformationist*, the embryo was literally a homunculus, a tiny individual complete in all its parts; in later variations it was seen as taking on ancestral forms, as ontogeny was said to recapitulate phylogeny. As for Lamarckism, the transmission of acquired characteristics required that they be preserved in the germ plasm in some way, so the germ plasm had to contain some kind of plan for those parts of the organism to be altered in subsequent generations. In each case the germ plasm could be regarded as a representation or copy of the organism.

A *recipe* is a sequence of procedures or instructions. It describes how to create a product but doesn't necessarily incorporate a description of the product (a recipe for a cake doesn't look like a cake). A *blueprint*, on the other hand, doesn't ordinarily say how to construct the structure that it shows. A blueprint is a representation or copy but a recipe isn't.

A major achievement in contemporary biology was to reinterpret genetic material not as blueprints for the organism's structure but rather as recipes for its development (Dawkins, 1986, Ch. 11). The modern formulation demanded rethinking of the sense in which genetic material contains information, whether about evolutionary history or about the organism's structure (Dawkins, 1982, Ch. 9). Genetic materials provide some information about the past environments in which they've been selected, especially when those materials can be found in many species, but they don't include the genetic materials of all those other organisms that didn't survive. In combination with developmental environments they determine the eventual structure of an organism, but **they do so as recipes for the production of proteins rather than as blueprints for body parts.** One implication was that Lamarckism and orthogenesis were untenable alternatives to Darwinian selection because their implicit copy theories were inconsistent with what we've learned about how the genetic material works.

It is ironic that Mendelian genetics had also been seen as a serious challenge to Darwinian selection. The integration of Mendelian genetics with Darwinian selection in the 1920s and 1930s, known as the *modern synthesis*, became the core of twentieth-century biology. Mendelian genetics had provided no mechanism for variations. Then genetic experiments with fruit flies by T. H. Morgan and others not only elaborated on genetic mechanisms but also brought mutations into the laboratory. With fruit flies, many generations could be studied within a relatively short time. The evidence yielded natural rates of mutation and the magnitudes of mutation effects, which were relatively small compared to the changes assumed by prior mutation theories. The combination of Mendelian genetics with the facts of mutation provided the variability needed for the workings of natural selection. The case has been further strengthened in recent decades as developmental biology has been incorporated into the story, in an account sometimes called Evo Devo (Carroll, 2005).

The Darwinian view had to face and overcome other hurdles besides competing theories (cf. Mayr, 1982). One was the incompleteness of the fossil record. Our understanding of prehistoric life depended on finding occasional preserved members of earlier species, but the accidents of their preservation and discovery left inevitable gaps. Furthermore, hard parts such as bones or shells were far more likely to be preserved than soft parts. And even when we found all parts intact, our information about how those creatures behaved was limited. It was often necessary to resort to indirect evidence (e.g., analogies with living species; fossil records of behavior such as fossil footprints).

The age of the earth was another problem. In those days the estimate was too short to make evolution through natural selection plausible, but that age has been revised vastly upward. Another misunderstanding was the likelihood of improbable events coming to pass when many opportunities for those events occur over extended time periods. For example, suppose some organic molecule is a crucial prerequisite for life and occurs naturally with odds of one-in-a-million only when lightning creates it by passing through some mixture of gases present in earth's early atmosphere. Its creation might seem a very long shot. But many lightning storms over many millions of years during our planet's early history would have repeated those conditions many millions of times, making it a virtual certainty that the molecule would be created not just once but again and again, even though the exact moments of its creation would have been unpredictable (cf. Dawkins, 1986; Gleick, 1987).

Variation and Selection

Variation and selection will be central to much of our later discussion of learning and behavior, in phenomena called shaping and reinforcement. We'll examine it now in the context of natural selection. Let's start with a population of prey animals (e.g., antelopes) the members of which vary in how rapidly they can outrun predators. The reasons might include differences in anatomy (e.g., bone length, muscle size), sensory differences that allow some to get off to a quicker start than others, metabolic differences that affect endurance, and so on. Everything else being equal, the slowest are the ones most likely to be caught.

The proviso that everything else must be equal is important. Speaking only of speed is an oversimplification. For example, an antelope that is fast at the expense of needing more to eat might forage longer and thus run a higher risk of being seen by predators during foraging. If this antelope is more often chased than others, that higher risk might counterbalance its speed advantage. Or, one antelope might run faster than another but the other is harder to catch because it can change direction more unpredictably. As long as the antelope population varies, we can recast our argument in terms of how such factors affect the chances of being caught. Speaking of speed is convenient, but the effective dimensions of escape from predators are more complex than that.

At any time, our antelope population has some mean or average speed, with some members above that mean and others below it. The ones below are those most likely to be caught and so are less likely to pass their genes on to the next generation. The next generation will then include more descendants of those above the mean than of those below or, in other words, fewer of the previous slow and more of the previous fast runners. The mean speed in this generation will again be higher than in the last one. But the same kind of selection still operates: again, slower ones are more likely than faster to be caught. Over many generations, therefore, the mean speed becomes faster and faster. (Selection will operate similarly on the big cats and other predators, because their effectiveness in catching antelopes will vary across individuals too.)

The evolution of the horse provides striking evidence for such selection (Gould, 1996a; Simpson, 1951). Over the 50 million years or so since *eohippus*, the so-called dawn horse (technically, its proper name is *Hyracotherium*), individuals in the populations from which modern horses are descended gradually became larger. These size changes were accompanied by other changes (e.g., toes becoming hooves), presumably including changes in behavior. Eohippus was the ancestor of modern horses, but it is unlikely that an eohippus population could survive in contemporary habitats. The fact that eohippus is extinct is relevant to our story. Many descendants of eohippus must have been the most evasive of their kind in their time, but they are no longer around. When selection operates on some relative property, such as speed relative to a population mean, the mean for the population changes. After capture by predators has repeatedly selected faster escape in a population, few descendants of the originally slow runners will be left even if that slower running speed provided a selective advantage at a time when it was very fast relative to the mean. In other words, as eohippus demonstrates, we should not expect to find examples of ancestral forms within current populations.

According to these arguments **the environment does the selecting** (the environments of predators include their prey and the environments of prey include their predators). It maintains as well as creates and alters the characteristics of organisms. The ancestors of whales were once land mammals. After they moved back into the sea, the environmental contingencies that made legs advantageous no longer selected well-formed legs. Instead, selection began to favor limbs effective for movement through water. The legs of the ancestors of whales gradually disappeared; in a sense it is appropriate to say that the legs had extinguished or become extinct (Provine, 1984; Skinner, 1988, p. 73). Selection operates on species by acting on particular organs and systems and body parts, and all of these become important because of the ways in which they serve an organism's behavior. Sensory organs and nervous systems and muscles and so on determine what an organism can do.

Consider another example. Environments in which a major food source consists of hard nuts favor beaks that work well at crunching nuts. Although Darwin did not recognize the significance of the observation until some time after visiting the Galapagos Islands during his voyage on the *Beagle*, such beaks were a characteristic of one population of finches there (Weiner, 1994). Such selection hadn't occurred with finches on other islands where softer foods were readily available. The hard nuts had set the occasion for the selection of finches with good nut-crunching beaks.

Selecting environments include members of one's own species. In the discussion of natural selection, the neck of the giraffe has often been offered as an example of selection, on the grounds that environments with food high on tall trees selected for long necks. But the neck of the male is typically longer than that of the female though the same selective contingencies should have acted upon both genders. Despite the plausibility of this evolutionary story, the long neck of the giraffe appears not to be a product of selection by such environments. Instead, female giraffes prefer males with long necks, and this sexual selection by the female is more likely what drove the evolution of long necks (Coe, 1967; Gould, 1996b; Simmons & Scheepers, 1996).

The role of *sexual selection* has long been appreciated in accounting for other evolutionary extravagances, of which the tail of the peacock is a familiar example (Darwin, 1871). The peahen, the female of that species, is more likely to mate with a male with a larger and/or more colorful tail. Despite the metabolic and other costs of their elaborate tails, such males are more likely to be healthy and therefore to provide favorable genes to their offspring. There is no inconsistency between these contingencies of sexual selection and those involved in other aspects of Darwin's account of natural selection. We'll encounter a related example later, in an account of how the female cowbird may shape the dialect of a male cowbird's song through differential attention (see Chapter 9).

We've discussed phylogenic selection involving gradual changes taking place over long periods of time (it has much in common with a kind of selection that occurs within the lifetime of the individual; see Chapter 9 on shaping). Some controversies about evolution have been about whether evolution takes place through *gradual changes*, as in the example of the horse, or in fits and starts (*punctuated evolution* or *saltation*). The fossil record includes evidence of major changes in species over periods of time that are relatively short by evolutionary standards. Examples include the explosion of multicellular life in the Cambrian period, and at the end of the Cretaceous period the extinction of the dinosaurs, perhaps triggered by the impact of a comet or some other planetary catastrophe, and the later proliferation of large mammals. Given the strong evidence for both kinds of evolutionary change, it is reasonable to conclude that evolution can take place either way, with some features selected gradually and continuously relative to a population mean and others selected following punctuated events that produced massive environmental changes.

The Origins of Complexity

Evolution by natural selection involves more than changes along a single dimension. An example is animal mimicry. A stick insect may look so much like a stick that a bird that otherwise would have eaten it will pass it by. But to get to look like a stick requires changes in surface shape and texture and color accumulating over many generations. As that selection got going, how much good did it do for the insect's ancestors to have merely a 5% resemblance to a stick? In response to this question, Dawkins (1986, pp. 83–84) points out that a 5% resemblance might be just enough to make a difference in twilight or in fog or if the bird is far away. Once individuals in the population vary in their resemblance to sticks, natural selection based even on small differences can drive populations to more and more convincing mimicry.

Another example of organized complexity is the intricate structure of the human eye. Is it reasonable to believe that natural selection could have produced such organized complexity? Using an analogy from aeronautical design, Dawkins (1982) posed the problem this way for engineers designing a jet engine:

> Imagine what they would have produced if they had been constrained to "evolve" the first jet engine from an existing propeller engine, changing one component at a time, nut by nut, screw by screw, rivet by rivet. A jet engine so assembled would be a weird contraption indeed. It is hard to imagine that an aeroplane designed in that evolutionary way would ever get off the ground. Yet in order to complete the biological analogy we have to add yet another constraint. Not only must the end product get off the ground; so must every intermediate along the way, and each intermediate must be superior to its predecessor. (Dawkins, 1982, p. 38)

If the eye is a product of natural selection, it couldn't have emerged full blown. But what good is part of an eye? The answer is that even 1% of an eye is a substantial advantage if all of one's contemporaries have even less. Any sensitivity to light is better than none, 2% is better than 1%, 3% is better than 2%, and so on (cf. Dawkins, 1986, p. 81).

But that is only part of the story. Not only have eyes evolved many times in many species. We now know that the same gene, Pax-6, provided the source of the evolution of all eyes across the entire animal kingdom, from worms and flies and squid and crabs and fish to mammals like us (Carroll, 2006, pp. 194–196). "The eye, far from being one of the most difficult structures to account for by evolution, has become one of the leading sources of insights into how evolution works with common genetic tools to build complex organs… Common genetic tools are used to build the very different hearts, digestive tracts, muscles, nervous systems, and limbs of all sorts of animals" (Carroll, 2006, p. 202) .

The earliest light sensitive tissues emerged half a billion years or more ago, during Cambrian times. "Natural selection has not forged many eyes completely from scratch; there is a common genetic ingredient to making each eye type…. These common genetic ingredients must date back deep in time, before there were vertebrates or arthropods, to animals that may have first used these genes to build structures with which to see, sense, eat, or move" (Carroll, 2005, p. 72). With sensitivity to light, predators and prey could begin to orient and move with respect to each other not just in response to accidental contact but also when at some distance from each other. From even those early times, the kind of selection relative to the population mean that we considered for antelopes and horses would have produced rapid evolutionary change, as predators became progressively more efficient in capturing prey and prey in turn became progressively more efficient at evasion (cf. Parker, 2003).

Evolution and Development

Science too evolves. Much of what we can now say about natural selection is based upon research on human and other genomes first reported only within the first decade of this century. The field continues to grow and change. I cannot do the topic justice here, but some parts of it have enough relevance to other parts of our behavioral story that it is appropriate to sketch out some of its features (cf. Carroll, 2005, 2006; Kirschner & Gerhart, 2005; West-Eberhart, 2003). The significant developments include demonstrating the modular functions of the compartments of the developing embryo, recognizing the independence of these modules in development and evolution, and working out the roles of the conserved evolutionary processes that built structural support, circulation, neural innervation and other basic biological functions into multicellular organisms.

A key was the integration of embryology with genetics and evolution: "every animal form is the product of two processes—development from an egg and evolution from its ancestors" (Carroll, 2005, p. 4), leading to the conclusion that "the development of various body parts such as eyes, limbs, and hearts, vastly different in structure among animals and long thought to have evolved in entirely different ways, was… governed by the same genes in different animals. The comparison of developmental genes between species became a new discipline at the interface of embryology and evolutionary biology, or 'Evo Devo' for short…. Evo Devo can trace the modifications of structures through vast periods of evolutionary time—to *see* how fish fins were modified into limbs in terrestrial vertebrates, how successive rounds of innovation and modification created mouthparts, poison claws, swimming and feeding appendages, gills, and wings from a simple tubelike walking leg, and how many kinds of eyes have been constructed beginning with a collection of photosensitive cells" (Carroll, 2005, p. 9–10).

The earliest genetic contributions to life on our planet were those involving the genetic system itself, including the replication of DNA and an array of basic metabolic and structural functions called *conserved processes* (e.g., Kirschner & Gerhart, 2005). With the emergence of multicellular forms these processes expanded to determine the structure of the growing embryo, and in particular to geographic arrangements of compartments or modules, each containing genetic material destined to control various body parts and functions (Goodsell, 2009). Compartments remain important throughout development and manifest themselves in various parts of the mature organism, including the nervous system; furthermore, the boundaries between compartments can be affected by their environments as well as by their genetic context (Bolker, 2000; Irvine & Rauskolb, 2001; Kiecker & Lumsden, 2005). Later we'll see that modules and the formation of boundaries are relevant to behavioral units, such as classes of responses with common functions.

Genes are, in effect, chemical switches, components that can be turned on or off at various stages of development. "A critical property of these switches is that changes in one switch will not affect the function of the other switches. And therein lies a huge insight into how form evolves" (Carroll, 2006, p. 206). "Modularity allows for the modification and specialization of individual body parts…, independent of other body parts… [allowing] evolutionary change to occur in one part of structure, independent of other parts" (Carroll, 2005, p. 195). Modules that isolate various developmental functions are crucial to the evolution of organized complexity. Without them a mutation that influenced, say, the form or size of a limb, would too often have lethal effects on other systems. Within compartmentalized systems, genes can have very specific effects without disrupting other functions.

As such systems play out, the conserved processes that determine skeletal structure and blood circulation and muscular attachment and neural connection and so on each do their job to accommodate to the developing organism within which they find themselves. If a mammalian forelimb changes its form or its size, blood supply will serve it in its entirety based upon chemical gradients produced by metabolizing cells, muscles and tendons will attach to bones and joints at appropriate places, neurons will find their way to sensory cells or to the sites of neuromuscular junctions, and so on, all independently of the details of the developing hindlimbs or other growing parts of this particular organism. The restriction of genetic changes to single sites opens up a vast evolutionary potential; effective mutations do not require synchronized changes at other sites.

Thus, the evolution of limbs is independent of the evolution of the nervous system that controls them: "the connectivity of the nervous system can change as the anatomy of the organism changes, without the need to alter the processes that generate that connectivity" (Kirschner & Gerhart, 2005, p. 160; see also p. 203). Earlier conceptions of self-organizing processes in developing organisms, and especially in developing nervous systems, properly appealed to common function, as in the assumptions that cell assemblies would come about because "cells that fire together wire together" (e.g., Hebb, 1949), but those systems were insufficiently concerned with the role of boundaries in the development of functional units.

Another property of these systems is that they too involve varieties of selection. For example, cell death is a crucial component of the developing nervous system. As cells grow from the central nervous system to sensory receptors at the periphery, only some find their way; those others that cannot find a binding site die (e.g., Edelman, 1987). The genes of the developing organism do not give instructions about where structural parts will go; instead, the parts go their variable ways and are selected on the basis of their functional engagement with other parts. This is the sense in which brains and other organs adapt themselves to the bodies within which they find themselves. Though the variable workings of these systems preclude predictions of particular outcomes, once they are in place there is a certain inevitability to the proliferation and elaboration of organized complexity. "This is a general theme in evolution, that one innovation creates the opportunity to evolve additional innovations" (Carroll, 2006, p. 113). Where there is life there is evolution.

KINDS OF SELECTION

The selection we've emphasized so far, the selection of populations of organisms over evolutionary time, is called *phylogenic selection*. But it is not the only kind. Of the various kinds of learning we'll explore later, one is the case in which responses are affected by their consequences. For example, if an organism is food-deprived and some response produces food, that response is likely to occur more often. This is an instance of reinforcement. It can also occur in particular environments or settings, when we say that a situation sets the occasion on which responses are reinforced. This is selection operating within the lifetime of the individual rather than over successive generations: the reinforcer selects those responses that will continue to occur.

Such selection has been called *ontogenic selection*; it involves *selection by consequences* (cf. Skinner, 1981). For the food-deprived organism, for example, responses that produce food continue to occur; other responses don't. Food is the consequence that selects some responses and not others. This is a way of saying that the responding is selected by its environment (notice how very different this is from saying that the organism itself selected some way to respond). We could say that responses that produce food survive and others that don't extinguish. Parallels between these two varieties of selection, phylogenic selection or Darwinian natural selection and ontogenic selection or the selection of behavior by its consequences, have been explored in considerable detail (e.g., Catania, 1978, 1995b; Skinner, 1981; Smith, 1986); we'll note some parallels as we explore learning.

Behavior acquired through learning during an individual's lifetime will disappear unless it is somehow passed on. A third variety of selection occurs when behavior can be passed on from one organism to another, as in imitation or, more important, in language. For example, what someone has said or written can survive the person's death if it is passed on to and repeated by others. The verbal behavior that survives within and is shared among the members of a group is part of the culture of that group. We'll give special attention to this third kind of selection in the chapters on social learning and on verbal behavior (Chapters 19 through 24). It has been called *cultural selection* or *memetic selection*. The term *memetic* is derived from *meme*, a word coined by Dawkins (1976) for an event passed along among individuals, (e.g., a vocabulary item, a catchy melody).

These varieties of selection may complement one another, but selection at one level also may oppose selection at another. For example, drug abuse is behavior that has been reinforced in ontogeny, but it will have unfortunate phylogenic consequences if it damages the fetus. At the level of cultural selection, the celibacy of some of its members may benefit a group, but it disadvantages those individuals at the phylogenic level. Conversely, at the phylogenic level opposition to birth control may benefit individuals who produce more offspring than those who practice it, but if population growth exceeds economic growth then per capita income decreases; the cultural consequences may be spreading poverty and environmental degradation. Phylogeny may have produced the capacity for behavior to be selected by its consequences in ontogeny, but this does not imply that individuals or their progeny will necessarily profit from that capacity. Similarly, beneficial social behavior transmitted from one individual to another only with difficulty may be displaced by competing behavior less likely to be reinforced but more easily transmitted, as when extreme political rhetoric overrides reasoned discussion (Catania, 1994, 2001b). "The dilemma of good and evil was created by multilevel selection, in which individual and group selection act together on the same individual, but largely in opposition to each other" (Wilson, 2012, p. 243).

We've considered three kinds of selection: (1) phylogenic selection, the evolution over biological time of populations of organisms and their characteristic features, such as behavior; (2) ontogenic selection, the shaping of behavior by its consequences during the lifetime of an individual organism; and (3) cultural or memetic selection, the survival of patterns of behavior as they're passed on from some individuals to others. These kinds of selection depend on behavior that changes during ontogeny as well as during phylogeny. We must understand each kind of selection to see where behavior comes from.

ADDENDUM 3A: PHYLOGENY, ONTOGENY AND BEHAVIOR

As research in Evo Devo has shown, development and evolution are intricately interrelated. This makes sense because both involve interactions with the environment. Consider the butterfly, *Bicyclus anynana*, and its adaptation to the seasonal swings of its habitat in Malawi (Carroll, 2005, pp. 214–215). In the wet season, broods of these butterflies emerge with big eyespots on their wings, which amidst lush foliage contribute to their evasion of predatory birds and lizards; in the dry season, however, amidst withered foliage and brown leaf litter, those spots make a distinctive target, and broods of these butterflies emerging at the start of the dry season have only flecks of color in place of the eyespots. The system works because these caterpillars have evolved to be sensitive to the ambient temperature in such a way that when they metamorphose into a butterfly, the gene that produces their spots is activated if it develops in warm wet season temperatures but not if it develops in cooler dry season temperatures. Thus, eyespots may vary from one generation to the next, depending on the environment, though the generations share a single gene pool. They have common *genotypes*, or genetic makeups, but different *phenotypes*, in the individual variations occasioned by the different developmental courses in their different environments.

The interaction of development and evolution becomes even more complex when development can be influenced by heritable units other than genes, such as chemical complexes that accompany the transcription of genes (e.g., Ferrell, 2012). Such phenomena are the concern of epigenetics, which examines how both genes and intracellular environments can have major effects on the path of development, in what Waddington (e.g., 1957, 1959) called the epigenetic landscape, the hills and valleys of potential that both limit and open up paths of development.

Another instance of the interacting roles of environments and genotypes is the enhancement of the likelihood of functional mutation in a phenomenon called *the Baldwin effect* (Baldwin, 1896). When environments change, organisms adapt, as when a fish population adjusts to an increase in the water temperature of its habitat. Selection will soon work on the fish population in favor of those best fitted to the higher temperature. If a mutation arises that makes the fish even better adapted to the new temperature, then those carrying that mutation will have a selective advantage even though the mutation followed the adaptations produced by the temperature change. First came ontogenic changes produced by the new environment, with phylogenic changes following later (cf, Kirschner & Gerhart, 2005, pp 76, 222). Still another possibility is that the new temperature will select a genetic feature that was long latent in some members of the fish population, becoming effective only after environment change. In this instance, the phylogenic changes came first and the ontogenic changes later (Kirschner & Gerhart, 2005, pp. 251–252).

Together these examples show how new mutations can become functional even if their origins were unsynchronized with relevant environmental changes. We must also allow the possibility that variability itself can be selected, in what has been called *facilitated variation*: "the capacity for generating a broad range of somatic adaptations is as heritable as anything else" (Kirschner & Gerhart, 2005, p. 83). We will explore a parallel issue when we examine the reinforcement of variable responding in Chapter 13.

So much for our whirlwind tour of evolution and development. Let us now focus again on behavior, which is a joint function of *phylogenic contingencies*, those operating in ancestral environments during the evolution of species, and *ontogenic contingencies*, those operating during interactions between organisms and their environments within their own lifetimes (cf. Skinner, 1966). Ontogeny does not recapitulate phylogeny, so we cannot trace the evolution of behavior by following the development of behavior in an individual (or vice versa). Whether phylogeny or ontogeny contributes more to behavior has been a common question. How much does behavior depend on evolutionary history and how much on learning? When such questions are addressed to socially significant issues such as the inheritance of intelligence, especially when the alternatives are presented as oppositions (e.g., nature versus nurture, heredity versus environment), heated controversy often follows.

Douglas Spalding, a nineteenth-century British naturalist, provides an elegant example of the role of phylogeny in behavior:

> we have only to look at the young of the lower animals to see that as a matter of fact they do not require to go through the process of learning;... chickens, for example, run about, pick up crumbs, and follow the call of their mother *immediately* on leaving the shell.... I have observed and experimented on more than fifty chickens, taking them from under the hen while yet in the eggs. But of these, not one emerging from the shell was in a condition to manifest an acquaintance with the qualities of the outer world.... (Spalding, 1873/1954, pp. 2–3)

Spalding noted that the chicks advanced rapidly. Within four or five hours of hatching they were pecking at objects and preening their wings. But he also recognized that a lot could be learned in four or five hours.

> To obviate this objection with respect to the eye, I had recourse to the following expedient. Taking eggs just when the little prisoners had begun to break their way out, I removed a piece of the shell, and before they had opened their eyes drew over their heads little hoods, which, being furnished with an elastic thread at the lower end, fitted close round their necks. (Spalding, 1873/1954, p. 3)

Spalding kept the chicks blind for one to three days and then removed their hoods.

> Almost invariably, they seemed a little stunned by the light, remained motionless for several minutes, and continued for some time less active than before they were unhooded. Their behaviour, however, was in every case conclusive against the theory that the perceptions of distance and direction by the eye are the result of experience, of associations formed in the history of each individual life. Often at the end of two minutes they followed with their eyes the movements of crawling insects, turning their heads with all the precision of an old fowl. (Spalding, 1873/1954, p. 3)

Our primary interest in this text is in behavior that is learned, but we must always entertain the possibility that the behavior we study has phylogenic sources. We can try to create arbitrary environments to minimize the role of phylogeny. For example, a standard pigeon chamber is an arbitrary environment, because natural environments don't include keys on which a pigeon's pecks produce food only when the key is lit. But arbitrary environments aren't always arbitrary enough; they won't necessarily make the behavior that occurs in them arbitrary. Consider a pigeon's key pecks:

> Such responses are not wholly arbitrary. They are chosen because they can be easily executed, and because they can be repeated quickly and over long periods of time without fatigue. In such a bird as the pigeon, pecking has a certain genetic unity; it is a characteristic bit of behavior which appears with a well-defined topography. (Ferster & Skinner, 1957, p. 7)

Behavior may start very early in an organism's life, but that in itself is not evidence that its sources are phylogenic rather than ontogenic. Recall that Spalding had misgivings about how much a chick could learn within just a few hours after hatching. Creatures may be prepared by phylogeny to do the same sorts of things that their ancestors have done, but they also may be prepared to start learning right away. Behavior begins in the embryo (e.g., Hall & Oppenheim, 1987). Both prenatally and postnatally some of that behavior is independent of sensory input and of consequences. But behavior changes even with the organism's earliest interactions with its environment (e.g., Johanson & Hall, 1979; Rudy, Vogt, & Hyson, 1984). By this point it should be evident that the answer to the question of whether behavior is a product of phylogeny or ontogeny is that it is a product of both. To questions about the relative magnitudes of their contributions, of course, the safest answer is typically, "It depends."

Chapter 4
Elicited and Emitted Behavior

Stimulus *can be traced to an Indo-European root* steig- , *to stick. The same root also generated* distinguish, instinct, *and, via* stylus, *a writing instrument,* style *(but it is probably not closely related to* extinction*). Through the Old French* estiquet, *to impale and later to label, it also produced a modern French word now related to behavior,* etiquette.

The Indo-European root spend-, *to pour a libation or to make a treaty, led to the Greek* sponde, *a drink offering, and the Latin* spondere, *to promise. Through these words,* response, *originally an undertaking in return, is linked to* spouse, sponsor, *and perhaps even* spontaneous. *The latter relation is interesting because* response *now refers to a unit of behavior that needn't be produced by a stimulus; responses may be elicited by stimuli but they may also occur spontaneously, when they are said to be emitted. Maybe they should be called* sponses.

The concept of the reflex was central to early studies of learning. We'll start with what happens when we present stimuli. Some relatively simple examples of stimuli producing or eliciting responses will demonstrate that we can't judge the effects of a stimulus unless we also know about the responding that occurs in its absence. The concepts of probability and conditional probability will allow us to cope with this problem; these terms will also be useful in the treatment of other phenomena later in this book.

After we examine simple reflex relations in terms of conditional probabilities, we'll deal with how elicited behavior may change over successive stimulus presentations. That will set the stage for exploring how behavior can emerge when it has not been produced by an eliciting stimulus; such behavior is said to be emitted. The chapter closes by treating how the behavioral significance of a stimulus can change over time, with imprinting as an example that relates such effects to the concept of drive or motivation.

KEY TERMS: Reflex; Elicited and Emitted Responses (Elicitation and Emission); Probabilities and Conditional Probabilities; Excitation and Inhibition; Habituation and Potentiation; Exercise; Imprinting.

THE REFLEX: ELICITATION

Presenting a stimulus is a simple way to change behavior. If I'm standing and talking, for example, a sudden loud noise will probably stop my talking and produce the change in my posture called a *startle reaction*. This reliable relation between a stimulus and the behavior change it produces has been called a *reflex*. The application of the vocabulary of the reflex to behavior began with René Descartes, a seventeenth-century French philosopher (Fearing, 1930). Descartes was familiar with hydraulic devices constructed to amuse visitors in the royal gardens of France. Stepping on a concealed trigger released a flow of water that made statues move. Descartes saw a similarity between such devices and behavior. According to Descartes, stimuli were comparable to the garden visitors who,

> entering into one of the grottoes containing many fountains, themselves cause, without knowing it, the movements which they witness. For in entering they necessarily tread on certain tiles or plates, which are so disposed that if they approach a bathing Diana, they cause her to hide in the rosebushes, and if they try to follow her, they cause a Neptune to come forward to meet them threatening them with his trident. (Descartes, translated in Fearing, 1930, pp. 20–21)

Just as a step on the concealed plate triggered the movement of a statue, a stimulus triggers a response. For Descartes, the role of pipes and water in those statuary systems was played by nerves and animal spirits in living organisms.

For our purposes, the most important part of Descartes' concept of the reflex was that it captured the fact that behavior is sometimes caused by environmental events, as when I quickly withdraw my hand upon touching an open flame. Eventually physiologists turned their attention to the mechanism of such behavioral relations and explored the components of the reflex arc, the path from the original sensory impact of the stimulus through the central nervous system and then back to the muscular or glandular system within which the response occurred. Analyses of the reflex became more and more sophisticated (Sherrington, 1906), and the conditioned-reflex concepts of Pavlov (1927) and the related behaviorism of Watson (1919) treated the reflex as the basic unit of behavior.

To the startle reaction and the withdrawal of a hand from a flame we can add many other familiar examples: the knee jerk produced by a tap on the patellar tendon; salivation caused by food in the mouth; the postural adjustments triggered by an abrupt loss of support. These have the common feature that some stimulus reliably produces some response. This is what defines a reflex. In such circumstances, we say that the stimulus *elicits* the response, or that the response is *elicited* by the stimulus; the stimulus is an *eliciting* stimulus, and the response is an *elicited* response (we never use *elicited* the other way around, to refer to the effect of a response on the occurrence of a stimulus).

The reflex is neither stimulus nor response; it is the relation between them (Skinner, 1931). For example, we wouldn't speak of a reflex if we set off a firecracker but observed no startle response; the noise alone is insufficient to justify the term. Neither would we speak of a reflex if we observed a response without an identifiable stimulus; by itself, the startle reaction isn't a reflex. And we should not confuse startle reactions produced by loud noises with those produced in other ways. Many reflexes have been named after their characteristic responses. The startle and patellar reflexes are examples. But **these are *not* names for responses**; applied to knee jerks in the absence of any eliciting stimuli, the reflex terminology is inappropriate.

Consider another example. Pupillary constriction produced by bright light to the eye has usually been called the pupillary reflex, but this is technically incorrect. This relation involves a response that adjusts continuously to stimulus levels. As brightness increases, the pupil constricts; as it decreases, the pupil dilates. When my eye is relaxed in total darkness my pupil will be about twenty times larger in area than when my eye is exposed to very bright light. By focusing on just one particular change in pupil size taken from a continuous and broad range of possibilities, the language of the reflex here diverts attention from the inverse functional relation between pupil size and brightness.

The language of the reflex may also be misleading in other cases that superficially seem to involve reflex relations. In everyday talk, for example, we sometimes speak of people having quick or slow reflexes, often with regard to skilled athletic behavior in competitive sports. But such usages are typically inaccurate. For example, a fast reaction to the starting gun of a race is operant behavior, behavior that depends on both its antecedents and its consequences, even though it occurs very soon after the shot.

Reflexes are just one subset of many kinds of relations that can exist in behavior. We'll restrict our attention here to reflex relations that involve external stimuli. We won't be concerned with reflex coordinations among responses, as when, in limb flexion or extension, the contraction of one muscle is accompanied by the relaxation of an opposing one, or with other more complex cases, as in rhythmic coordinations during running or other forms of locomotion (Gallistel, 1980).

Properties of Elicited Behavior

Let's assume we've identified a reflex relation between some stimulus and some response (cf. Skinner, 1938, on the Laws of the Reflex). Consider an eliciting stimulus such as an acid- or sour-tasting solution on the tongue (e.g., vinegar) and an elicited response such as salivation. Above some *threshold* value, acid on the tongue reliably elicits salivation, but a low concentration or small quantity may not do so. When the magnitude of a stimulus is too small to elicit a response, the stimulus is said to be *below threshold*.

The threshold isn't a fixed quantity; it is a statistical summary. As we concentrate the solution or increase the amount, we're more likely to produce salivation. At a given concentration, we can find out how many drops on the tongue will reliably elicit salivation; conversely, for a given number of drops, we can find out what concentration is required. Some concentrations and quantities will have only marginal effects, sometimes eliciting salivation and sometimes not. Note that a stimulus doesn't itself have a threshold. Rather, threshold values are determined for some features of stimuli with others held constant. For example, the threshold number of drops for eliciting salivation would typically be smaller given a strong than a weak acid solution; it wouldn't be meaningful to specify the threshold number of drops without also specifying acid concentration, and vice versa.

Some time always elapses between stimulus and response; this time period is called the response *latency*. In addition, the response must exist in some *magnitude* and have some *duration*. Because these properties often vary together, they're sometimes given a common name, *reflex strength*. Thus, responding with long latency, small magnitude and short duration corresponds to weak reflex strength, whereas responding with short latency, large magnitude and long duration corresponds to strong reflex strength. The rate at which the elicited response occurs isn't relevant to reflex strength, however, because it is determined by the rate at which the eliciting stimuli are presented.

Sechenov (1863/1965), a Russian physiologist, noted that the organism's energy expenditure in many responses (e.g., sneezes and coughs) far exceeds the energy provided by the eliciting stimulus. The eliciting effect of a stimulus doesn't depend on a direct transfer of energy from environment to organism; rather, observed Sechenov, the stimulus is a trigger, releasing energy that the organism already has available in muscles or glands or other structures. This view was, of course, consistent with how Descartes conceived of the reflex.

But even though the eliciting stimulus is most accurately regarded as a trigger that releases energy that is already available, its magnitude may affect the elicited response. Response latency typically varies *inversely* with stimulus magnitude; in other words, response latency typically gets shorter as stimuli get bigger. Response magnitude and duration typically vary *directly* with stimulus magnitude; in other words, bigger stimuli typically produce bigger and longer responses. Saying that reflex strength increases with stimulus magnitude just summarizes these relations.

Once talk of reflex relations shifted from specific measures to global abstractions such as strength, it became easier to regard the reflex relation as a fundamental unit of behavior. Soon reflexes were treated as basic building blocks from which more complex behavior could be constructed. The concept of the reflex had a tempting simplicity. As it became more widely accepted as a behavioral unit, it seemed reasonable to conclude that reflex relations could be the basis for understanding behavior in general. Once some stimuli were identified as causes of some responses, the faulty assumption was made that for *every* response a corresponding eliciting stimulus must exist.

Pavlov's conditioned-reflex system and Watson's behaviorism of the 1920s and 1930s grew out of such an assumption. With reflexes serving as units of behavior, complex behavior was treated as nothing more than the combination of such units. When a response occurred with no observed eliciting stimulus, stimuli with appropriate properties were hypothesized. In addition, the responses of one reflex were assumed to have stimulus properties that enabled them to elicit other responses in turn. Thus, behavior extending over long periods of time could be interpreted as a sequence or chain of reflexes, with each response functioning simultaneously as the elicited response of one reflex and the eliciting stimulus of the next. These reflex systems were elaborated in sometimes ingenious ways, but they no longer command major attention in the psychology of learning. The concept of the reflex has its place but its scope is limited. It cannot stand alone.

Eliciting Stimuli and Response Probabilities

Of the possible relations between stimuli and responses, the reflex is only one relation among many. In a reflex, some stimulus reliably produces some response. But the stimulus of that reflex may affect other responses differently, and the response of that reflex may be affected differently by other stimuli. Any stimulus may raise the likelihood of some responses, lower the likelihood of others, and have no effect on still others. Similarly, the likelihood of any response may be raised by some stimuli, lowered by others and unaffected by still others.

Furthermore, simply noting that a response reliably follows a stimulus isn't enough to justify talking about reflexes. We must also know how likely the response is without the stimulus. For example, if a rat in a running wheel spends most of its time running and runs when a noise is on as well as when it is off, we couldn't say that noise had elicited running just because we happened to see the rat run after we presented noise; it might have run anyway, without the noise. To speak of reflex relations, we must know how likely the response is in the absence as well as in the presence of the stimulus.

Probabilities or Relative Frequencies. Such relations among stimuli and responses can best be described quantitatively in terms of probabilities or relative frequencies. We can define the effects of stimuli by comparing response probability when a stimulus is present with response probability when the stimulus is absent. A probability or relative frequency is simply a proportion or ratio: the number of times an event occurs compared with the number of times it could have occurred.

For example, in the Babinski reflex, a splaying or spreading of the toes is elicited by a light stroke across the bottom of an infant's foot. We calculate response probability by counting how many stimulations produce responses. In the newborn, the splaying of the toes may be produced by each of 20 touches to the foot; the proportion of responses is 20 out of 20, a probability of 1.0. The Babinski reflex ordinarily diminishes with age, so if we test again later, only 6 out of 20 touches may produce a response; the probability is now 0.3. Eventually the reflex completely disappears. None of the 20 touches produces a response, and the probability is now 0.0. Probabilities always must fall in the range from 0.0, when the event never occurs, to 1.0, when the event always occurs. In the mathematical notation for probabilities, these examples may be written respectively as:

$$p(R_1) = 1.0; p(R_2) = 0.3; \text{ and, } p(R_3) = 0.0.$$

Here, *p* stands for probability and the parentheses contain abbreviations for the specified events (in this instance, R for responses, with subscripts 1, 2 and 3 corresponding to the three tests).

Conditional Probabilities. The Babinski reflex is uncomplicated for our purposes, because the toe splaying doesn't occur often without its eliciting stimulus. But what about blinking elicited by a puff of air to the eye? If we study this reflex with an infant who already blinks every second or so, how can we tell elicited blinks from those that would have occurred even without a puff (Spence & Ross, 1959)? In fact, we may be unable to say whether any given blink is elicited. We can assess the overall effect of the stimulus, however, by comparing the probability of a blink after a puff with the probability after no puff.

Our procedure is illustrated in Figure 4-1. We watch the eyelid and record blinks within some time period after each stimulus, and we compare responding during those times with responding during equivalent times that don't follow a stimulus. The probabilities in this example are expressed as:

$$p(R/S) = 1.0; p(R/\text{no } S) = 0.6.$$

In this notation, R is a response and S is a stimulus. The first equation says that the probability of a response given a stimulus (blink given puff) equals 1.0. The second says that the probability of the response given no stimulus (blink given no puff) equals 0.6. This kind of probability, in which the probability of one event is given in terms of the presence or absence of another event, is a *conditional probability* (the terminology must not be confused with that of Pavlov's conditioned or conditional reflexes, even though both have the feature that one event is a condition for some other event). In other words, $p(A/B)$ can be read as: probability of A given B, or probability of A on the condition that B is present. Similarly, $p(X/\text{no } Y)$ can be read as: probability of X given no Y, or probability of X on the condition that Y is absent.

Our later analyses will favor the language of probabilities over other ways of describing behavior. Saying that a stimulus caused a response is pretty much the same as saying that the stimulus elicited the response. Either usage can be applied to single instances of a reflex relation (as when this particular stimulus, M, is followed by this particular response, N). Similarly, saying that a stimulus increased response likelihood is pretty much the same as saying that it raised response probability. Either usage can be applied to average effects over many instances (as when stimulus O usually produces response P).

Types of Stimulus-Response Relations

So far we've discussed cases involving one stimulus class and one response class. Behavior is usually more complicated. Environments include various stimuli and organisms produce various responses. For example, we might notice two kinds of chirps produced by a recently hatched duckling. Those produced when the mother duck is present are contentment calls and those when a strange bird appears are distress calls (e.g., Hoffman, 1996). We could lower the likelihood of contentment calls and raise the likelihood of distress calls either by removing the mother duck or by presenting the strange bird. We should be able to describe the effects of either stimulus on either response. We don't deal with stimuli alone or responses alone; we deal with stimulus-response relations.

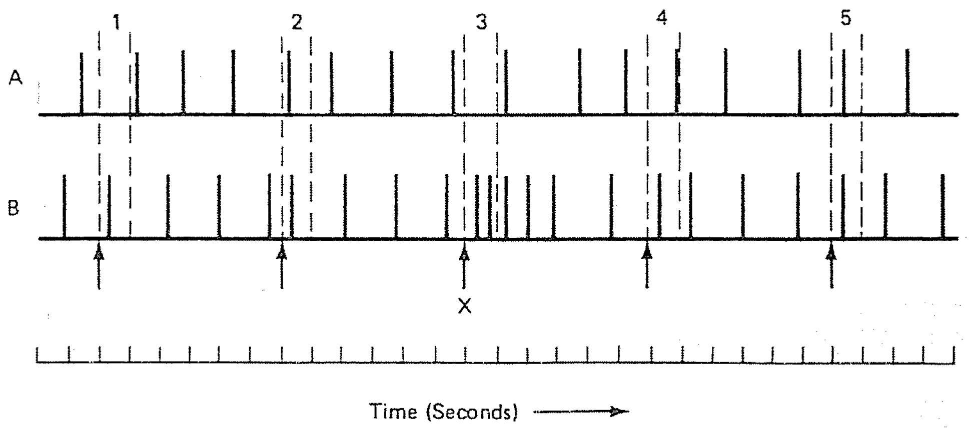

Figure 4-1 Estimating eyeblink probability with and without eliciting puffs of air. Each solid vertical in A and B represents a blink. Dashed lines mark off five 1-s periods during which blinks were recorded. Arrows in B represent puffs to the eye. In A, no puffs were presented; blinks occurred within periods 2, 4 and 5, so blink probability was 3/5 = 0.6. In B, a blink occurred within each 1-s period following a puff, so blink probability given a puff was 5/5 = 1.0. Only the first of the two blinks following stimulus X counted toward this probability; we score time periods as either containing no blinks or containing at least one blink and then calculate probability by dividing time periods with at least one blink by total time periods. The puff raised blink probability from 0.6 to 1.0. (An actual experiment would use a much larger sample to calculate probabilities.)

We can summarize these relations in the coordinate systems of Figure 4-2. The y-axis represents response probability given a stimulus has been presented; the x-axis represents response probability given the stimulus hasn't been presented (cf. *coordinate* in the glossary). In other words, any point on these graphs represents two conditional probabilities: response probability given a stimulus, $p(R/S)$, and response probability given no stimulus, $p(R/\text{no } S)$. For example, point A represents a stimulus-response relation in which response probability is 0.75 when the stimulus is presented and only 0.50 when it isn't; in this instance, the stimulus raises the probability of the response.

The diagonal in Figure 4-2 (left) is of special interest. A response that occurs without being elicited by a stimulus is said to be *emitted.* The diagonal represents stimulus-response relations in which response probability is unaffected by or is independent of the stimulus. Thus, at point B response probability is 0.25 whether or not the stimulus is presented. For example, the duckling's contentment calls and distress calls might be unaffected by the presence or absence of some of its siblings. To say a response has been elicited we must know more than that it followed a stimulus. A response can just happen to follow a stimulus; we can't say it was elicited unless we know that it was caused by the stimulus.

A third class of stimulus-response relations is illustrated by point C, for which a probability of 0.90 without the stimulus is reduced to 0.10 by the stimulus; in this instance, the stimulus reduces response probability. Such reductions of response probability by a stimulus are sometimes called *reflex inhibition*. For example, if the duckling's contentment calls stopped when the strange bird appeared, we might say that the strange bird inhibited the contentment calls.

Point D represents a case in which a stimulus raises response probability from about 0.1 to about 1.0. This is the sort of stimulus-response relation we see in reflexes. Switching to a more traditional example, this might be a case where a dog rarely

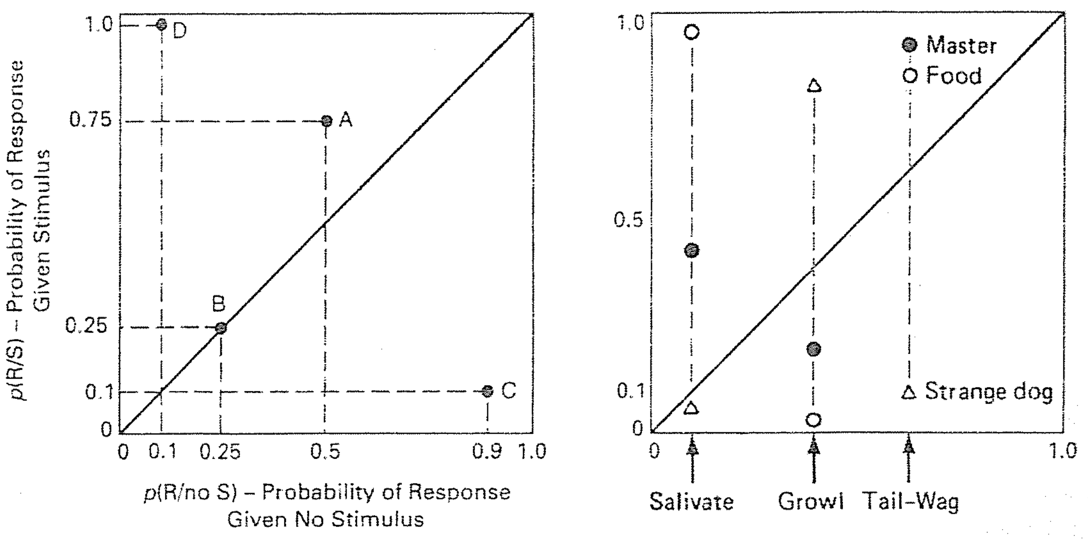

Figure 4-2 Stimulus-response relations represented as conditional response probabilities given the presence of a stimulus, p(R/S), and its absence, p(R/no S). On the left, examples are shown in which a stimulus raises response probability (A), has no effect on response probability (B), or lowers response probability (C); the increase in response probability called a reflex, in which the stimulus reliably produces or causes the response, is illustrated at D. On the right, the effects of three stimuli, a dog's master, food, and a strange dog, are shown in terms of the locations of three responses, salivating, growling and tail-wagging.

salivates in the absence of food but almost always does so given food in the mouth. The graph shows that this is only one among a range of possibilities. Somewhere between the upper edge, where p(R/S) is near 1.0, and the diagonal, where S has no effect on R, we must decide that the eliciting effect is no longer reliable enough to justify calling the relation a reflex. But where should that boundary be? Probably we'd include cases in which p(R/S) is just a little less than 1.0 (e.g., 0.95; maybe even 0.90). But probably we'd also exclude small effects of stimuli. For example, if a stimulus raised response probability from 0.26 to 0.32, would this relation qualify as a reflex? Almost certainly not.

All points above and to the left of the diagonal represent *excitation*, cases in which stimuli raise response probability; those below and to the right of the diagonal represent *inhibition*, cases in which stimuli lower it. Within those areas, any boundary marking off reflex relations would be arbitrary. This conclusion is important. We noted that some earlier behavioral systems, such as those of Watson and Pavlov, were based on reflexes as fundamental units of behavior. Highly reliable reflex relations were surely easier to work with than other less reliable stimulus-response relations. But if Figure 4-2 (left) is the way to describe stimulus-response relations and if the reflex is just one special case among them, then any system of behavior built solely on the reflex as a behavioral unit is bound to be incomplete.

When we defined the reflex, we argued that the reflex is neither stimulus nor response but rather the relation between them. The left graph of Figure 4-2 supports this point. It represents stimulus-response relations; it cannot represent stimuli by themselves or responses by themselves. Neither the strange bird as a stimulus nor the duckling's distress call as a response has a location in the graph. The right graph illustrates it further. The place where we'd locate a dog's salivation, growling or tail-wagging in these coordinates would depend on whether its environment included its master, food or a strange dog. For example, it growls most given the strange dog and least given food, but it wags its tail most given its master and least given the strange dog.

One property of behavior left out of early accounts was the emission of responses, responses occurring without eliciting stimuli. How could responses occur if something hadn't caused them? The answer is that there are causes other than eliciting stimuli. When researchers began to study emitted responses, they were given such names as *instrumental* or *operant*, because they were studied in terms of how they were instrumental in changing the environment or how they operated on the environment. They derived their importance not from their relation to eliciting stimuli but from their consequences. Elicited responses were by contrast called *reflex* or *respondent*.

Once that distinction had been made, qualifications were added. In particular, some argued that instrumental or operant behavior consisted of skeletal responses, such as movements of the limbs, whereas reflex or respondent behavior consisted of autonomic responses, such as glandular secretions. This was also seen as paralleling the traditional distinction between voluntary and involuntary action. Such distinctions have since been seriously questioned.

For example, the crouching posture elicited by a sudden loud noise in the startle reflex is a motor response best described as involuntary. But swallowing seems voluntary and yet also involves a reflex relation. It is elicited by stimulation of the back of the throat, which is why you can't swallow if your mouth is dry and you have nothing in it to pass back to that area to stimulate it. On the other hand, driving certainly seems both voluntary and operant. Yet an experienced driver sitting in the passenger seat of a car may involuntarily press hard on the floor even though there is no brake pedal there when something suddenly looms up ahead on the road. We have here all of the possibilities: whether behavior is operant or respondent, we can identify examples that seem either voluntary or involuntary. In other words, **the everyday distinction between voluntary and involuntary actions has nothing to do with our distinction between operant and respondent behavior.** (As we'll note when we later get to verbal behavior, the most important determinant of whether we call our behavior voluntary or involuntary may be whether we can identify the source of our actions.)

Skeletal responses can be elicited and autonomic responses can be emitted. It is important to maintain the distinction between elicited and emitted responding. But the same response may be sometimes elicited and sometimes emitted; we can't classify responding effectively into these two categories on the basis of physiological criteria such as the difference between skeletal and autonomic responses.

Effects of Successive Elicitations

A complication in the effects of stimulus presentations on behavior is that two different presentations of the same stimulus may have different effects. For example, I may startle much more to the first lightning flash in a thunderstorm than to later flashes. Furthermore, the effects of stimuli may depend on how quickly they follow each other. If I'm peeling onions, for example, the tears elicited by the present onion may depend on whether I began working on it right after finishing the last one or just after taking a long break from peeling. And in another effect, called *summation*, a stimulus that is below threshold in eliciting a response if presented once may become an effective elicitor if presented repeatedly. A single whiff of pollen might not be enough to make me sneeze, but several whiffs might do it, even if the amount of pollen in each whiff is about the same. In other words, elicited responding often depends on the number of stimulations and their separation in time.

Habituation. The startle reaction is produced by an unexpected event such as a lightning flash or a sudden loud noise. Even without other events to signal it, a repeated loud noise usually produces successively smaller startle reactions, until eventually there is no startle at all. Many stimuli elicit responses called orienting or observing responses; for example, a dog pricks up its ears in response to a novel sound or begins sniffing in response to an unusual odor. As these stimuli recur, the dog's responding decreases; it occurs with smaller magnitude and longer latency, perhaps even vanishing completely (we'll see later, however, that orienting or observing responses may depend not only on eliciting stimuli but also on their consequences).

This reduction in responding with repeated stimuli is called *habituation*. Another possible term, *adaptation*, sometimes refers instead to changes in behavior in the continued presence of some stimulus or situation, as when an organism is said to adapt to a laboratory setting. Habituation is a characteristic of the elicited responding produced by a variety of stimuli. It occurs with such diverse responses as distress calls of birds to a silhouette of a predator passing overhead, and contractions in earthworms exposed to light, not to mention the startle reactions and orienting responses already mentioned (e.g., Ratner, 1970). It may also be an important component of the dynamics of emotion (Solomon & Corbit, 1974).

Habituation has been an important component of studies of whether preverbal children can distinguish among such features of spoken language as intonations, rhythms, and the basic speech sounds or phonemes of the languages of their caregivers (e.g., Eimas & Miller, 1992; Eimas *et al.*, 1971; Ramus, 2002). One response used in such studies is nonnutritive sucking. Essentially, the infant is given a pacifier connected to a device that records the pressure exerted while sucking. If the infant hears a voice speaking a consonant, say *b*, the infant will typically stop or slow its sucking briefly. If the *b* is repeated, it has less and less effect. In other words, the response to *b* habituates.

Now suppose the speaker says *p* instead of *b*. The difference between *b* and *p* is that the former is voiced (the vocal chords vibrate) whereas the latter is not. Given an infant several weeks old who has been raised by English-speaking caregivers, the first instance of *p* will ordinarily produce a pause in sucking, even though the response to *b* has habituated (this might be called *dishabituation*). These different responses to *b* and *p* tell us that the infant can distinguish between the two consonants even though not yet able to produce them.

Not so, however, if the infant has been raised by caregivers who speak a language in which the distinction between *b* and *p* is unimportant. Consider a child raised in an Arabic household, where a single voiced consonant somewhat like the English *b* has no corresponding voiceless *p*. Once this infant's pause in sucking has habituated to the sound *b*, changing to *p* makes no difference. The infant responds similarly to both consonants; there is no evidence that it distinguishes between them. An important lesson from such studies is that infants begin to learn about the sound features of the languages they hear around them long before they become skillful in differentially producing those features. What the infant hears and therefore learns in the English-speaking household is different from what the infant hears and learns in the Arabic-speaking household.

There is a procedural issue. Measuring pauses in nonnutritive sucking calls for having some sucking in the first place, so procedures of this sort typically require some sucking as a criterion for presenting the speech sounds. Some studies showed increases rather than decreases in sucking after the speech sounds were presented. If sucking is required to produce a sound and sucking then occurs more often, it is reasonable to ask whether the sound has functioned as a reinforcer. Such findings have been explored not only with nonnutritive sucking but also with such other responses as direction of gaze toward one or another sound source. When we address verbal behavior later, we'll see that hearing the sounds of one's native language can have potent reinforcing effects. For the present purposes, the point is that we must distinguish between habituation and other behavioral processes to work out what goes on within such procedures.

Potentiation. But stimuli sometimes have opposite effects. For example, electric shocks elicit squealing in rats; if several shocks are delivered, later ones produce more responding than earlier ones (Badia, Suter, & Lewis, 1966). This effect has been called *potentiation*; another term sometimes used is *facilitation* (e.g., Wilson, 1959). Potentiation is more likely with stimuli regarded as aversive or punishing than with stimuli regarded as appetitive or reinforcing. For example, when I was an undergraduate I was a subject in class demonstrations that involved electric shock (maybe they wouldn't have passed contemporary institutional human-subject review boards). One involved holding my hand against an electrode as shocks were delivered. The first shock just tingled, but it felt stronger with each successive presentation, so it became progressively more difficult not to pull my hand away even though the shock level remained constant. My response to the shock had potentiated.

Potentiation mustn't be confused with another phenomenon called *sensitization* (cf. Ison & Hoffman, 1983). In sensitization, the eliciting effects of one stimulus are enhanced as a result of presentations of some *other* stimulus; one stimulus amplifies the eliciting effect of another stimulus. For example, an electric shock may make it more likely that a later loud noise will produce a startle reaction; the shock is said to *sensitize* the organism to the noise.

Time Since the Last Eliciting Stimulus. If the stimulus is absent for a while after habituation or potentiation, the probability that responding will be elicited returns to earlier values. For example, my startle reaction to loud noise may diminish or even disappear after several noises in succession, but after hours of silence I'm likely again to startle substantially to the next noise. If elicited responding decreases over successive stimuli, it usually recovers to its earlier higher levels after the stimuli are discontinued. Conversely, if elicited responding increases over successive stimuli, it usually returns to its earlier lower levels after the stimuli are discontinued. Habituation and potentiation aren't permanent; a return to earlier levels takes place as time passes.

If habituation were irreversible, it could either occur only once in an organism's lifetime or else successive habituations would drive responding to lower and lower levels until it disappeared altogether. Irreversible potentiation would similarly lead either to a single case of potentiation in the organism's lifetime or to continuing and unlimited increases in elicited responding. Such things may sometimes happen. For example, reactions of the immune system, which are at least analogous to instances of elicited behavior, sometimes seem to show irreversible potentiation (e.g., if someone develops an allergic reaction to bee stings after being stung several times, the reaction to subsequent stings may diminish little if at all as time passes).

REFLEXES, FIXED ACTION PATTERNS, AND OTHER CLASSES OF BEHAVIOR

Probabilities are not the only features of behavior to be considered in distinguishing among kinds of behavior. For example, most reflex relations involve responses the magnitude of which varies with the magnitude of the stimulus, as when, in the patellar reflex, the magnitude of the knee jerk increases with increasing force of the tap to the patellar tendon. Fixed action patterns, however, such as the mating dance of the stickleback fish upon sight of the swollen belly of a female, are not typically treated as reflexes because they do not vary with properties of the stimuli. A fixed action pattern is typically an all-or-none response. Once it has begun, it typically occurs full-blown and runs its complete course.

Organisms orient with respect to features of their environments. They have evolved to do so in different ways. For example, suppose bright light elicits random movement. An insect larva in the light starts moving and continues to do so until, by chance, it moves into darkness; once there it stops. We won't find many of these larvae in the light, but we may find large numbers in dark places (e.g., under the rotting bark of a fallen tree limb). Once we expose them to light, they all begin to move. But their behavior isn't directed toward dark places; they get there by chance, some sooner than others, and they end up congregated there only because that is where they stop.

Orientation that occurs on such a basis is called a *kinesis*, or plural *kineses*; it is distinguished from orientation directed toward or away from some stimulus, which is called a *taxis*, or plural *taxes* (Fraenkel & Gunn, 1961; Loeb, 1918/1973). Examples of taxes include movements toward light (positive phototaxis) and upward movement, away from gravity, as in climbing a tree (negative geotaxis). The details of these and other types of orientation vary (for example, in an organism with two eyes a phototaxis may come about because the organism consistently moves so as to equalize the amount of light received by each eye). Directed movements are sometimes also called tropisms, but that term is more often applied to the behavior of plants (as when flowering plants turn toward the light in phototropisms).

Still other varieties of behavior are available before learning (e.g., Gallistel, 1980; e.g., Von Holst, 1973). Some types of behavior that aren't learned have the characteristics of oscillators (e.g., circadian rhythms). Some have the characteristics of servomechanisms (e.g., the maintenance of balance, during which small displacements produce compensating adjustments). These and other types are worth knowing about because many remain as components of our own human behavior. Some important examples include laughter, yawning, tickle, and the shedding of tears (e.g., Provine, 2000; Provine, Krosnowski, & Brocato, 2009). The various sources of behavior provide our taxonomy, our system for classifying behavior. In chapters to come we'll examine operants, discriminated operants, meanings and rememberings as some of the classes that emerge from this taxonomy.

FROM ELICITED TO EMITTED BEHAVIOR

We've seen that the reflex is just one of the many possible relations among stimuli and responses. Now we examine how stimuli affect responding over more extended periods of time, and not just immediately after the stimulus has been presented. This topic is a prerequisite for considering responses that are emitted rather than elicited, i.e., responses that have other causes than stimulus presentations. This will lead us to a discussion of ways in which the behavioral significance of stimuli may change over time; imprinting as an example relates such effects to establishing operations and the concept of drive or motivation.

One problem in tackling this topic is the very word *response*. It implicitly suggests a response *to* something. Fortunately the term *behavior* does not have a similar connotation, but it would be a poor substitute because the ambiguities of number in behavior and behaviors and kinds of behavior produce other difficulties. We'll later find the coinage *operant* helpful, but it too raises questions, with regard to classes versus specific instances.

As a derivative of *response* without the implication of responding *to* something, I'm sorely tempted to introduce the term *sponse* (cf. Provine, 1988), but it would be presumptuous of me to do so because the term has so little precedent in the literature of behavior and learning. But if the distinction between elicited and emitted responding sometimes seems too esoteric, it might be helpful to try a translation in which *elicited responding* becomes *responding* while *emitted responding* becomes *sponding*.

The Role of Exercise

"There is a variation of an old Latin proverb that states '*repetitio est mater doctrinae*'—repetition is the mother of learning" (Carroll, 2006, p. 154). "Practice makes perfect" is probably a more familiar saying about the role of repetition in behavior. Before the effects of the consequences of responding were appreciated, it was believed that the mere repetition of responding, without regard to its consequences, was sufficient to maintain behavior. Consider the following from Sechenov:

> an infant is able to cough, sneeze and swallow immediately upon birth. The act of sucking also belongs to this category of complex movements.... Indeed, everybody knows that a new-born child is able to suck.... Moreover, it is a well-known fact that the activity of this complex mechanism in the infant is called forth by irritation of the lips; put, for example, your finger, or a candle, or a wooden stick between the child's lips, and it will begin to suck. Try to do the same with a child three months after it has been weaned, and it will no longer do so; however, the ability to produce sucking movements at will is retained by man for life. These are highly remarkable facts: on the one hand, they show that the conduction of sensation from the lips to the central nervous mechanisms which produce the sucking movements apparently ceases in the child after weaning; on the other hand, they indicate that the integrity of this conduction is maintained by the frequent repetition of the reflex (Sechenov, 1863/1965, pp.28–29).

Sechenov here emphasizes the complexity of the sucking response, but more important is his observation on the role of repetition. Not only does repetition maintain the response, according to Sechenov, but the response also becomes independent of the effects of eliciting stimuli. In the infant, sucking is elicited by stimuli ("irritation of the lips"). Later, these stimuli no longer elicit the response, but the organism through adulthood remains able to produce the response even in the absence of these stimuli: "For most altricial species, the ability to suckle from the mother is reduced if suckling is not practiced.... Virtually all of the organized maternally oriented behavior of most infant animals can be shown to be heavily influenced by experienced events" (Hall & Oppenheim, 1987, p. 113; cf. Schoenfeld, 1966).

Based on such phenomena, we might conclude that the repeated elicitation of a response increases the likelihood that the response will be emitted. Early accounts of learning (e.g., Thorndike, 1921; Verhave, 1967) treated effects of response repetition as basic components of learning, described in terms of laws of *exercise* or laws of *practice*. These laws were usually ambiguous on questions such as whether it mattered if the repeated response were elicited or emitted. In any case, they soon were overshadowed by other concerns. As the psychology of learning turned to such phenomena as reinforcement and the signaling functions of stimuli, the possible role of exercise or practice became neglected. The evidence is too scanty to allow firm conclusions about whether exercise or practice might be a basic component of learning, but some tantalizing data exist.

One example is provided by research on the escape responses of a tropical fish, the zebra danio, from a stimulus designed to resemble the rapid approach of a larger predatory fish (Dill, 1974). As the stimulus loomed, the danio began to swim away from it (despite the aquatic medium, this swimming response of the fleeing fish can aptly be called a flight reaction). The latency with which swimming began decreased with repeated exposures; in other words, the stimulus more and more strongly elicited the escape response. After 10 days without further elicitations, the response hadn't returned to its earlier levels. The change produced by the stimulus wasn't reversible (cf. the section on potentiation above).

But not all responses begin with elicitation. Some of the earliest responses in an organism's lifetime occur spontaneously, in the absence of identifiable eliciting stimuli. For example, inside its egg the embryo chick makes uncoordinated movements of its limbs and body. These movements may prevent the developing bones from becoming fixed in their sockets, or may modify the form of growing bones and connective tissue. Later in the embryo's development, eliciting effects of stimuli appear, perhaps simply as the embryo's sensory apparatus matures. This progression from spontaneous responding to elicited responding may be summarized by saying of chick embryos that "they 'act' before they 'react'" (Provine, 1976, p. 210). Stimuli become more important later, as when the chick's rotating movements in pecking its way out of its shell during hatching are affected by whether the chick continues to encounter intact portions of the shell.

Pecking in the young chick depends not only on the conditions that elicited pecks and the consequences of earlier pecks but also on how much the chick has already pecked (Hogan, 1971). Once responding occurs, whatever its origin, it may have consequences, and the chick's survival may depend on whether those consequences in turn affect its behavior. Among gulls, for example, pecking at but missing the parent's beak has different consequences from striking the beak; only in the latter case is the parent gull likely to feed the chick. These differential consequences affect the accuracy of the chick's later pecking only as the response becomes independent of its eliciting stimuli. That development corresponds to one possible formulation of a law of exercise: once a response has been elicited by a stimulus, the response may become more likely even in the absence of the stimulus.

In experiments on salivation, dogs at first salivate only when food is presented, but after several food presentations they begin to salivate even when food is absent (Zener & McCurdy, 1939). Such responding, *spontaneous* salivation, had been attributed to conditioning of the salivary response to features of the setting (e.g., the eliciting stimulus might have been a spot on the wall that the dog just happened to see when food deliveries began). No other eliciting stimuli were identifiable, so including all salivary responses in a reflex relation could be justified only by assuming that salivation could be elicited by such arbitrary features. Yet the problem of identifying stimuli simply vanishes once we admit the possibility of responses that occur without eliciting stimuli; in fact, we can't otherwise conceive of emitted responding at all. The concept of emitted responding will be essential to our treatment of consequential operations such as reinforcement and punishment.

Early behavior theories held that simple repetition was important in its own right to the development and maintenance of behavior. Consider the following from Sechenov: "If a child which has just learned to walk becomes ill and remains in bed for a long time, it forgets the previously acquired art of walking.... This fact testifies once again to the great import for nervous activity of frequent repetition" (Sechenov, 1863/1965, p. 29).

We can't be sure of the evidence upon which Sechenov based his conclusion (e.g., did he observe one child or many children?), but we can consider a more recent example, in Zelazo, Zelazo, & Kolb (1972). The walking reflex appears in newborns and disappears at about eight weeks (McGraw, 1945). The response of the walking reflex has much in common with what the child does later when learning to walk. Coordinated walking movements resembling those of an adult can be elicited by holding the infant under its arms with its feet touching a level surface. When parents exercised their infants' walking reflexes by holding them so as to elicit the walking response, the reflex was less likely to drop out and these infants learned to walk earlier on the average than others whose walking wasn't exercised. Thus, exercising this reflex during the first eight weeks of life not only increased elicited responding that ordinarily decreases during this time; it seemed to shorten the time to its later appearance as a component of emitted behavior, walking.

But the researchers recognized that walking movements produce changes in the infant's world (visual, kinesthetic or tactile) that might reinforce such responses. Even during the eight weeks of exercise, walking was sometimes emitted rather than elicited, and more was going on than simple elicitation: "Walking... seemed to progress from a reflexive to an instrumental response. There is little doubt that learning occurred... Not only were there more responses... but they were better executed" (Zelazo et al., 1972, p. 315). The situation is further complicated as an example of elicitation by changes in the infants' weights and in the supporting capacities of their limbs over the time course of such studies (Thelen *et al.*, 1982).

We've speculated that after a response is produced by a stimulus it may become more likely even in the absence of the stimulus. In other words, eliciting a response may raise its probability of emission. This formulation differs from the classical laws of exercise or practice to which we've related it. It depends only on the simple

experimental operation of presenting stimuli. But we needn't try to formulate laws. Some emitted responses may originate because they are first elicited whereas others may be emitted from the start. The issue isn't the universality of laws of exercise or of such alternatives as spontaneous emission; rather, it is identifying the source of responding in particular cases. We'll treat other problems in behavior analysis in this way in later chapters. Instead of trying to explain instances of behavior in terms of exhaustive formal laws, we'll seek to improve our taxonomy of behavior by classifying behavior in terms of its origins.

Stimulus Presentations in Imprinting

So far, we've concentrated on how stimuli affect responding. Virtually all of the phenomena we've considered in this chapter can be described in terms of response probabilities. For example, in a reflex relation a stimulus raises the probability of a response to near 1.0 and in habituation response probability decreases over successive stimulus presentations. We now examine an outcome of stimulus presentations, *imprinting* (Lorenz, 1937), that must be discussed in terms other than effects on response probability. Imprinting provides a bridge to the treatment of response consequences in the next chapter.

When a duckling hatches, the first moving thing it ordinarily sees is its mother, and even on this first day of life outside the egg the duckling will probably stay close to her. But if the mother is gone and the duckling first sees something else in motion, such as a human, the duckling will behave toward this stimulus as it otherwise would have toward its mother. Such stimuli are said to be *imprinted*, or, in a figurative sense, stamped into the duckling.

Imprinting has been demonstrated in both laboratory and field with varied stimuli, ranging from real and model birds to electric trains (some stimuli, of course, work better than others). The development of imprinting is sometimes said to have a *critical period* of one or a few days, so that it may not occur at all if it doesn't occur during this critical period. Actually, things are more complicated than that (Hoffman, 1996). For example, fear of novel stimuli develops toward the end of the critical period; as the birds get older, they move away from novel stimuli, making distress calls. Thus, older birds don't ordinarily stay near such stimuli long enough for imprinting to occur, but it can occur if this effect of novel stimuli is prevented or reversed.

In any case, the duckling begins behaving in significant ways to a stimulus, whether mother duck, human or some arbitrary moving object, if the stimulus is introduced given appropriate conditions and times in the duckling's life. One response is following the imprinted stimulus as it moves, and it is sometimes said that the duckling's following is elicited by the imprinted stimulus. But speaking of elicitation is misleading.

If the imprinted stimulus is the mother duck, the duckling follows her about and emits distress calls in her absence. But how does the imprinted stimulus produce following? When the duckling walks toward the mother, it finds itself closer to her; when it walks from her, it finds itself farther away. In other words, the natural consequence of walking in different directions is to change its distance from its mother. If closeness to the mother is important to the duckling, it is no surprise that it walks toward her rather than away from her. It follows that if we change the duckling's world so that the mother' closeness requires some response other than walking, the walking should be replaced by that other response.

In such an experiment (Peterson, 1960), a dark compartment on one side of a window contained a moving imprinted stimulus, and on the other side a duckling was given a response that could light up the dark side. The duckling did so even when the response was one incompatible with following, such as pecking at a disk on the wall or standing still on a platform near the window. In other words, the critical property of the imprinted stimulus wasn't that it could elicit responses such as following or pecking or standing still, but rather that it had become important to the duckling and therefore could reinforce such responses as following or pecking or standing still. In natural environments, the duckling's following ordinarily keeps it close to the imprinted stimulus (usually its mother), but a laboratory environment shows that ducklings can learn other responses if they instead of following have the important consequence of keeping the imprinted stimulus close.

In imprinting, the initial presentations of the to-be-imprinted stimulus don't change response probabilities. Rather, they are establishing operations. They change the significance of the stimulus. The imprinted stimulus acquires its significance simply by being presented under appropriate circumstances. It begins as a stimulus toward which the duckling is relatively indifferent but ends as a stimulus that functions as a reinforcer and therefore can shape the duckling's behavior.

Establishing Operations and the Significance of Stimuli

The significance of stimuli can be changed in other ways. Some were treated in Chapter 2 as cases of establishing operations. For example, if a rat is more likely to eat than to run in a running wheel, we'd expect the rat to press a lever more often if its presses produce food than if its presses produce only access to the wheel. But suppose we continue to give the rat free access to food while we lock the wheel, thereby preventing the rat from running. After we deprive the rat of wheel running in this way, it will be more likely to run than to eat when we give it the opportunity to run again, and more likely to press the lever if presses produce access to the wheel than if presses produce food. In other words, we changed the relative significance of eating and running; through the deprivation of one or the other, we made eating more likely than running or running more likely than eating.

Changes in the significance of stimuli that occur with establishing operations are sometimes discussed in terms of *drive* or *motivation* (e.g., Bolles, 1975). Stimuli become more or less reinforcing or more or less aversive, depending upon factors such as the time since their last presentation. Food, for example, becomes more reinforcing as time passes without eating, but it may become aversive right after the eating of an unusually large amount.

The significance of stimuli can be changed through means other than deprivation; for example, as we shall see later, conditioned reinforcers and conditioned aversive stimuli are stimuli that have acquired their reinforcing or aversive properties through their relation to other stimuli.

Physiological studies of motivation are typically concerned with relations between organic factors and the significance of stimuli (e.g., effects of blood levels of glucose on behavior with respect to food, effects of hormonal levels on sexual behavior, etc.). Motivation or drive, therefore, is not a special force to be located somewhere inside the body; rather, it is a term applied to the many environmental and organic variables that make stimuli significant to an organism. The example of imprinting set the stage for a discussion of establishing operations. We're now ready to move on to the consequential operations of reinforcement and punishment.

ADDENDUM 4A: THE TEMPORAL PATTERNING OF BEHAVIOR

Presenting a stimulus may determine the sequence of responses over an extended time. If we give a rat food, for example, it will eat. Once it is finished eating, it will then typically drink if water is available. This relation between eating and subsequent drinking is so strong that by repeatedly delivering food in small amounts we can make the rat drink several times its ordinary daily ration of water (Falk, 1977; Wetherington, 1982). This increase in drinking is called *polydipsia*, and responding that depends in this way on other responding is called *adjunctive* behavior. Adjunctive behavior is behavior in which one response reliably accompanies some other response.

As one type of adjunctive behavior, polydipsia follows at least partly from the rat's normal feeding and drinking pattern. With food and water freely available, the rat ordinarily takes a few large meals daily and drinks after each meal. If we force the rat to take many small meals by delivering food in many

small portions every few minutes, the rat still drinks after each meal but doesn't reduce the size of each drink enough to compensate for its more frequent drinking. Thus, a rat going from 5 large to 50 small meals per day now drinks 10 times as often. But if the rat drinks only half as much after each of the 50 small meals instead of a tenth as much, it will drink 5 times as much as before. Such increases in the rat's water intake are so reliable that polydipsia can be used to get rats to consume substances such as alcohol that they ordinarily reject (e.g., Meisch & Thompson, 1971).

Adjunctive behavior can include other responses besides eating and drinking. For example, with a running wheel instead of water available to the rat, running follows eating much as drinking follows eating in the polydipsia procedure (Levitsky & Collier, 1968).

The presentation of a stimulus may impose temporal structure on behavior in other ways. One experiment examined patterns of behavior generated in pigeons by repeated food presentations (Staddon & Simmelhag, 1971). Each pigeon was placed in a chamber on one wall of which was an opening to a feeder, a tray of grain ordinarily out of the pigeon's reach; food was presented by lighting the tray and lifting it to a position where the pigeon could eat from it. The pigeon was observed through a window and its responding was scored in various categories, such as pecks directed toward the feeder wall or toward the floor and postures such as orientation toward the feeder wall.

Early in each interfood period, one bird usually made quarter-turns, and then, but less frequently, put its head into the feeder opening or pecked toward the floor. Other pigeons showed patterns that included different responses, such as preening or pacing. These responses became less likely and as the time of the next feeder operation approached pecking became the dominant or most likely response.

The repeated stimulus presentations imposed temporal structure on behavior. Responding early in interfood intervals varied from one pigeon to another, whereas later responding was fairly constant across pigeons and usually included some form of pecking. Pecking is also the behavior occasioned by food presentations; the bird takes food into its mouth by pecking. At least in this case, the later behavior seems to have something in common with the responses produced by stimulus presentations (cf. temporal conditioning in Chapter 17). Repeated stimuli seem not only to elicit responses but also to produce behavior at other times that is closely related to the elicited responding. The problem is that it is hard to show that sequences of behavior generated by successive stimuli are unaffected by other variables, such as their consequences or the discriminative effects of stimuli (e.g., Reid & Staddon, 1982).

ADDENDUM 4B: THE EVOLUTION OF BEHAVIOR

What can we say about the evolution of behavior (cf. Skinner, 1984)? We've already noted that behavior has left only indirect evidence in the fossil record. It is safe to assume, however, that response systems evolved before sensory systems. For an organism that can't do anything about what it sees, there is no advantage to seeing. Some organisms remained immobile, but others began to twitch and squirm. Some organisms were passively moved by ocean currents, but others anchored themselves to particular places and still others developed cilia and flagella. The organisms from which we are descended developed ways to get from one place to another.

As motor systems developed, the advantages of responding differentially to environmental events were presumably the basis for the selection of sensory systems. Withdrawal upon being touched might be enough to avoid a predator and would certainly be more advantageous than withdrawing at random. Ingesting things on the basis of their chemical properties would certainly be more advantageous than ingesting things randomly. These properties of behavior are so important that it is hard to imagine a world in which all creatures lacked them. But then we can recall that we shouldn't expect to find examples of ancestral forms within current populations.

The most primitive patterns of movement were probably driven primarily by eliciting stimuli. Earlier we discussed bright light eliciting random movement in insect larvae. They move until by chance reaching darkness and then they stop. We find few in the light but many in the dark. They weren't heading for dark places. They got there by chance and that's where they stopped.

These examples, like reflex relations, share the property that each involves a fixed pattern of responding to environmental events. First there was movement; then the movement came under the control of stimuli. Such patterns were particularly advantageous in stable environments.

For example, given a sharp stimulus to its paw, a dog flexes its leg, drawing its paw up toward its body. For an animal that walks on the ground, this response is advantageous. If the dog steps on a thorn, its flexion pulls its paw away from the thorn. The sloth, however, lives in a very different environment. It hangs from trees and its comparable reflex involves an extension rather than a flexion of its limb. If the sharp stimulus is a thorn, the hanging sloth that pulls its limb toward its body instead of extending it will only drive the thorn in deeper (cf. Hart, 1973, p. 176).

But not all environments are stable. It must have been a major evolutionary step when patterns of behavior became modifiable or, in other words, when some organisms became able to learn. Learning may at first have been selected within restricted domains. For example, upon leaving its nest the digger wasp flies in ever-widening circles; its later return to the nest is based on landmarks, such as the rocks or plants it flew over before its departure (Tinbergen, 1972). Its capacity to learn landmarks is part of its phylogenic heritage and is probably very specific to finding its nest.

The capacity to learn must itself have been selected. The selection of this capacity presumably occurred in different ways in different species at different times. Things about environments that ordinarily remain constant over a lifetime need to be learned only once, whereas things about changeable environments need to be learned and then discarded and then learned again. In our own lives, the names of people we know don't change very often, whereas the items on our shopping lists usually change from one shopping trip to the next. We should therefore expect some kinds of learning to be difficult to reverse while others remain transient and easily modified.

For example, the survival of a newborn calf may depend on how well and quickly it learns about the features that distinguish its mother from others; in a stable environment, such learning (sometimes called imprinting) may be effectively permanent. Under such conditions, the time of life during which learning occurs may also be limited. But as the calf grows and becomes more independent, it will also learn about many things that change from day to day, such as the passage of predators or the places where food or water are found. The capacity to learn and to remember may follow a developmental course. For example, Peterson (1962) discovered that the effects of shaping a key peck with newly hatched ducklings were at first short-lived, in that shaping had to be restarted from scratch if one came back to it later that day or on the next day, but that its effects became longer and longer lasting as the chick became older.

Those who look for mechanisms of learning in the nervous system must recognize that evolutionary contingencies have selected different kinds of learning, with some producing fairly permanent changes in behavior and others producing changes that are easily reversed, with some restricted to relatively narrow situations and others broadly general, and with some operating at different times in an organism's life than others. In other words, the selection of different kinds of nervous systems depends upon the different kinds of behavior they produce. If that evolution worked like the evolution of anatomical structures, we should not be surprised if some evolved learning systems are modular (cf. Chapter 3)

PART III
LEARNING WITHOUT WORDS

Chapter 5
Consequences of Responding: Reinforcement

The Indo-European root, sekw-, *to follow, links* consequence *to* signal *and* designate *(from the Latin* signum, *something that one follows) and to* social *and* association *(from the Latin* socius, *a companion or follower). It shares its prefix,* con-, *with* conditioning, contingency *and* contiguity. Conditioning, *through the Indo-European root* deik-, *to show or pronounce, has many relatives:* dictation, *from the Latin* dicere, *to say;* teach, *from the Old English* taecan, *to show or instruct;* judgment, *from the Latin* judex, *he who pronounces the law; and* paradigm, *from the Greek* para, *beside, plus* deiknunai, *to show.* Contingency, *from the Latin* contingere, *to touch on all sides or to happen, has varied meanings: a possibility; a condition of depending on chance; something incidental to something else. Like* contact, *it combines the roots* com-, *together, and* tangere, *to touch.* Contiguity, *the condition of touching or being in contact, has the same origins. Curiously,* contingency *and* contiguity *are usually contrasted in learning theories:* contingency, *in its technical use, stresses how the likelihood of one event is affected or caused by other events, whereas contiguity implies the juxtaposition of events in space or time without regard to causation.*

Reinforcement
  Mazes and Learning Curves
  Experimental Chambers and Cumulative Records
  The Vocabulary of Reinforcement
Extinction
  Extinction Versus Inhibition
  Response Reinforcer Contingencies and Reinforcer Deliveries
  Side Effects of Extinction
  Extinction Versus Free or Noncontingent Reinforcement
Addendum A: Extinction and Superstition

First, a reminder. Now and for a while we'll concentrate only on behavior without words. We'll consider human examples but mainly in cases where our human talk is irrelevant, and for that reason it is very easy to not to notice the human examples. They are embedded among examples involving nonhuman organisms such as pigeons and rats, and readers often miss the use of human behavior to illustrate the relevance of findings from nonhuman research. For example, one exam question I used for several semesters of teaching these topics asked how many instances of human behavior were mentioned in the text through the chapters corresponding to the present Chapter 8. The four choices were: *a*) none; *b*) fewer than

KEY TERMS: Reinforce, Reinforcer, Reinforcement; Contingency; Extinction, Spontaneous Recovery; Free or Noncontingent Reinforcement; Superstition; Side Effects.

10; *c*) about 20; *d*) more than 40. By actual count, more than 50 human examples appeared in those chapters in the first and second editions (there are even more now than there were then). The modal answer, however, was typically *a* or *b*. In other words, some students thought there weren't any at all, and most of the rest didn't notice more than ten of them. If you don't think *d* is the correct answer after reading those chapters, try actually counting the human examples.

The advantage of starting with nonhuman organisms is precisely that they don't talk: unlike us, they are nonverbal creatures. We must resist the temptation to treat what they do as if they talk to themselves about it the way we talk to ourselves. The absence of talk makes them so different from us that we would be well advised to take nothing about them for granted and to deal with them as if we had just come upon them as aliens, which in a way is what they are relative to us. Yet all of our behavior that involves talk is built upon nonverbal foundations that are covered over by the talk, and nonverbal organisms provide our best path to understanding those foundations. When we get to our human learning with words we'll be in a far better position to deal with what is special about it because we first surveyed learning without words.

Behavior has consequences, and an important property of behavior is that it can be affected by its consequences. We study this phenomenon by arranging consequences for behavior, but doing so involves more than just presenting stimuli. The stimuli must occur in some relation to behavior. We have to arrange the environment so that responses make something happen.

Natural environments already provide consequences for behavior. Even before we intervene organisms change their environments by doing things or by going from one place to another. But we can better study how consequences affect behavior by arranging consequential environments in the laboratory. For example, I can build a maze in which a water-deprived rat finds water after making an appropriate sequence of turns, or a chamber in which a food-deprived pigeon produces food by pecking a key on the wall. Then I can find out whether water affects the turns the rat takes as it runs through the maze or food affects how often the pigeon pecks the key.

In this and the next chapter we'll explore the historical development of research on the consequences of behavior and we'll see how reinforcement is relevant not only to behavior maintained by physiologically significant consequences such as food and water but also to behavior with more subtle consequences, as when, in sensory-motor interactions, my eye movements affect what I see.

REINFORCEMENT

Chapter 2 introduced Thorndike's experiments, in which animals learned to escape from problem boxes by operating a device that released the door. Typically, a food-deprived animal was placed inside the box with food available outside. In its varied activity, the animal sooner or later operated the device and was free to leave the box. At first this was a low probability response, but because it opened the door its probability went up over repeated trials.

Thorndike described how the consequences of responding affected later responding with a principle he called the *Law of Effect*. The law went through many revisions, but its essence was that response probability could be raised by some consequences and lowered by others. In language closer to Thorndike's, responses with satisfying effects were stamped in whereas those with annoying effects were stamped out. The earliest version of Thorndike's law was called the *strong* Law of Effect. Later, he repudiated the second half of the law, keeping the probability increase or stamping in but discarding the probability decrease or stamping out. What remained was then called the *weak* Law of Effect. This historical point will be relevant later, when we deal with punishment.

Figure 5–1 shows data from one of Thorndike's cats. To escape from its box, this cat had to pull a string running from a wire loop at the front of the box to a bolt holding the door. The first time in the box the cat took 160 seconds to escape (from here on, seconds will often be abbreviated *s*). Its time decreased gradually and irregularly over successive trials until, during the last few trials, the cat reliably escaped in less than 10 s. This gradual decrease in the time taken to complete a task came to be called *trial-and-error learning*. Köhler later contrasted this gradual change with the sudden or insightful solutions he observed with chimpanzees.

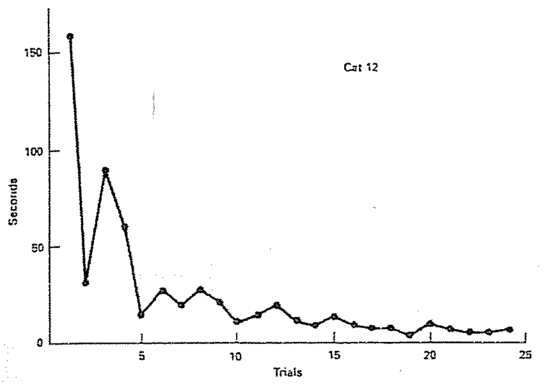

Figure 5-1 A learning curve: a cat's time to escape from a problem box as a function of trials. (From Thorndike, 1898, Figure 1)

Mazes and Learning Curves

In later years, trial-and-error learning was studied with many different organisms in many different situations. Experimenters believed that the intelligence of different species could be compared by seeing how rapidly learning went in problem boxes, mazes, runways and other apparatuses (e.g., Hilgard, 1951). Apparatus design began to be dictated by theoretical questions: Does learning take place in discrete steps, on an all-or-none basis, or instead gradually and continuously? Do organisms learn movements (response learning) or properties of the environment (stimulus learning)? Do the consequences of responding lead directly to learning or only make the organism perform so as to show what it had learned in other ways?

A common feature of these experiments was that responding became more probable when it had certain consequences. The change in probability was measured differently depending on the apparatus and the experimental aims. Graphs showing how behavior changed during an experiment were called *learning curves*: time to complete a response as a function of number of trials (e.g., Figure 5–1); percentage of correct responses; proportion of animals reaching some criterion of successful performance. Sometimes these measures were transformed to ease comparisons among them. When rats ran through a maze, for example, time to run from startbox to goalbox usually decreased, whereas percentage of correct turns and proportion of rats making errorless runs increased. Converting time to run the maze to speed (the reciprocal of running time) made all three measures increase with learning. But the shapes of learning curves depended so much on apparatuses and on measures taken that the progress of learning couldn't be described in any unitary way.

The problem was that these experiments produced complicated behavior. For example, measuring the time course over which a rat stopped entering blind alleys as it learned its way through a maze didn't show how learning proceeded at a single choicepoint. This consideration led to the gradual simplification of mazes, as illustrated in Figure 5–2.

Diagram A shows the plan of the earliest maze used to study animal learning (Small, 1899–1900), a 6-by-8-foot modification of the hedge maze at Hampton Court in England. (Curiously, such mazes also provided settings for the hydraulically operated statues that inspired Descartes to formulate his concept of the reflex; cf. Chapter 4.) When a cage door at the start was lifted, rats could enter the maze; food was in the goal area at the center. As their experience grew they reached the goal area more rapidly and with fewer wrong turns along the way. But it was difficult to examine learning at any given choicepoint. Learning might have been quicker at choicepoint 1 in diagram A than at 7 either because 1 was earlier in the maze than 7 or because their floor plans differed; learning might have been quicker at 4 than at 5 either because it could be approached in two ways, from 3 or from 5, or because a rat that usually went directly from 3 to 4 only rarely encountered 5.

Gradually, mazes evolved into more systematic forms, as in diagram B. In this maze, sometimes called a *U-maze* after the form of the successive units, choicepoints were essentially the same as the rat approached each one; they differed only in where they were in the sequence and in whether left or right turns were correct. This kind of systematic arrangement made it easy to specify the correct sequence (in B, right-left-right-left-left-right) and to keep track of errors. Even here, however, position and sequence interactions complicated the analysis. For example, is the rat's choice of left at 4 affected by the preceding right turn at 3 or the following left turn at 5? Would it matter if the rat approached 4 after coming back from the blind alley at 3, having made an error there, instead of after a correct right turn at 3? Does it matter that 4 is in the middle rather than near the beginning or end?

It was perhaps inevitable that the maze would be reduced to a single choicepoint, as in the *T-maze* shown with a goalbox on the right in diagram C in Figure 5-2. Here, when leaving the startbox the rat has only to make a single choice of right or left. But complications were still possible. For example, suppose one rat in its first trial in this maze turns right while a second rat turns left. Should the second one be allowed to retrace its steps after reaching the empty box at the end of the left arm? If it is returned to the startbox instead, should it be forced to the goalbox by blocking the left arm, to give it time in the goalbox equal to the first rat's? The next logical step was to eliminate choicepoints completely, leaving nothing but a simple runway, as in diagram D. Now no errors are possible, and measures of behavior can be reduced to just the speed of the rat's run from startbox to goalbox.

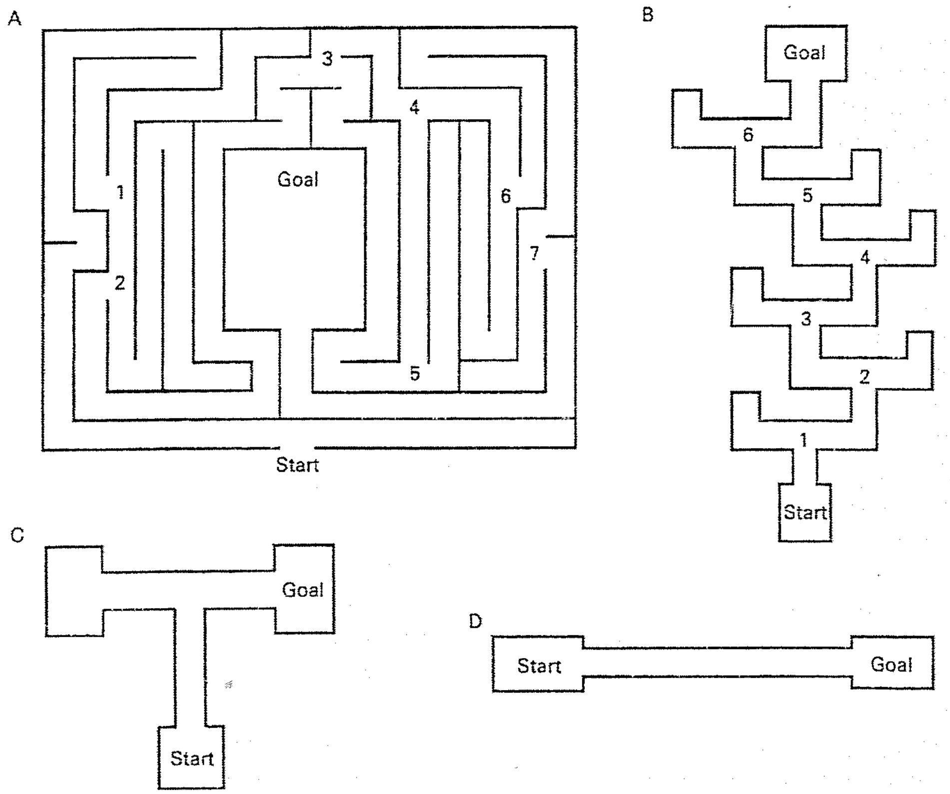

Figure 5-2 Stages in the evolution of mazes in studies of animal learning: A, the Hampton Court maze, as adapted by Small (1899–1900); B, a U-maze with six choicepoints; C, the single choicepoint T-maze; and D, the runway or straight alley.

There were other problems. Average measures of the performance of a group didn't necessarily represent performances of the individuals in the group. Suppose each rat running in a T-maze usually changes abruptly from making frequent errors to making consistently correct choices, but this change occurs on different trials for different rats. In a large group of rats, 65% might make correct turns by trial 5, 72% by trial 6, 79% by trial 7, 83% by trial 8, and so on, until performance becomes stable at 98% to 100% by trial 20. This group performance, giving the appearance of a gradual increase in correct choices, would completely obscure the abrupt change in errors for all individual rats (Sidman, 1952).

Even the simple runway wasn't an ultimate solution, because the speed of a rat's running down a straight alley can be affected by many trivial factors. If trials began with the opening of a startbox door, it depended on which way the rat was facing when the door opened. Other factors included the experimenter's handling of the rat when moving it between trials from goalbox to startbox, odor trails left by other rats, and even whether the goalbox had enough room to allow a running rat to slow down before banging its head against the goalbox back wall (Killeen & Amsel, 1987).

With mazes or runways, the experimenter had to return the rat from goalbox to startbox to begin a new trial. Thus, the experimenter rather than the organism determined when behavior occurred. Furthermore, measuring how long the rat took didn't specify what it was actually doing during that time. Two experimental innovations helped to solve these problems. The first was an apparatus designed so that the organism could repeatedly emit an easily specified response without the experimenter's intervention; the second was a recording method based directly on the rate or frequency of responding rather than on indirect measures derived from response sequences or groups of organisms. These innovations, inspired partly by an interest in reducing the handling of the organism and thereby simplifying the experimenter's work, were important features of a direction of research initiated by B. F. Skinner (1930, 1938, 1950); see especially Skinner (1956) for a history of these developments.

Experimental Chambers and Cumulative Records

Figure 5–3 illustrates two representative apparatuses: on the left a standard rat chamber with a single lever, and on the right a three-key pigeon chamber. They share response devices, mechanisms for delivering reinforcers such as food or water, and stimulus sources.

In a typical arrangement, a rat that has been food deprived is placed in the chamber. A lever protrudes from one wall. Near the lever is a food cup into which food pellets can be dispensed from a delivery system on the other side of the wall; a distinctive sound accompanies each pellet delivery. The houselight provides general illumination, and noise can be broadcast from the speaker to mask sounds from outside the chamber.

The first step is feeder training. Pellets are delivered into the food cup. Sooner or later, the rat finds and eats them. Once this happens, pellet deliveries continue until the rat comes quickly to the food cup from any location upon each delivery; 10 or so pellets are usually enough. Once feeder training is done, the apparatus is changed so that pellet deliveries depend on lever presses. Eventually the rat presses the lever, the press produces a pellet, and the pellet occasions eating. The rat will then probably go back to the lever and press it again. (We'll consider alternatives to waiting for the lever press in a later chapter on shaping.) The outcome of interest is the rate at which the rat presses the lever. If it increases, we call the food pellet a *reinforcer*. In the type of chamber shown in Figure 5–3, other kinds of reinforcers can be substituted. For example, the pellet dispenser could be replaced by a dipper that delivers small amounts of water or milk.

A pigeon chamber differs from one for rats in that keys substitute for levers and the feeder works by bringing a tray of pigeon food (mixed grain or commercially available pellets) within the pigeon's reach for a few seconds. The feeder opening is

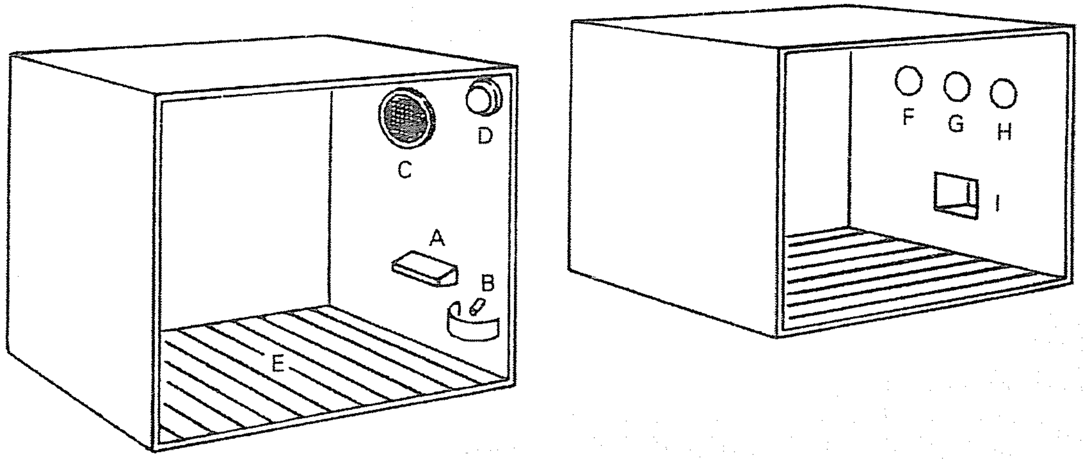

Figure 5-3 A rat chamber (left) and a three-key pigeon chamber (right). The rat chamber includes a lever (A), a food cup and pellet delivery tube (B), a speaker (C), and a lamp or houselight (D); some rat chambers include a grid floor through which shock can be delivered (E). The pigeon chamber includes three keys (F, G, H) and the opening to a food hopper (I); lamps or projectors behind each key allow colors or patterns to be displayed on them.

centered below the keys. It is common practice to light the feeder and turn off all other lights whenever the feeder is operated. The chamber typically includes other features, such as a houselight for dim general illumination, masking noise or other auditory stimuli, and so on.

A pigeon key is a piece of plastic mounted behind a hole in the chamber wall. It is attached to a switch that records the pigeon's pecks if they're forceful enough (keys are routinely sensitive to forces of less than 0.1 N or Newtons, about 10 g or one-third of an ounce). The plastic is usually translucent or transparent, so that lamps or miniature projectors or a computer monitor behind the key can present patterns or colors on it. The chamber in Figure 5–3 contains three keys, arranged horizontally about 23 cm (9 in) above the chamber floor. A given experiment might use only one key, some combination of two, or all three. Keys are typically lit when they're in use. As with the rat, if a food-deprived pigeon's key pecks produce food, the rate at which the pigeon pecks the key will ordinarily increase.

The rat and the pigeon are common laboratory organisms. They each have idiosyncratic species-specific behavior patterns that must be taken into account, and we mustn't assume that what we observe with rats or pigeons can be generalized to other organisms (though it often can be). Nevertheless, the diet, housing, susceptibility to disease and other characteristics of these animals are reasonably well understood, and their size, relatively long lifespan and economy make them particularly convenient. Thus, they have often served in research on the consequences of responding.

Responding in apparatuses like those in Figure 5–3 has sometimes been called *free-operant* responding: *free* because the organism is free to emit the response at any time rather than waiting for the experimenter (the rat in a goalbox can't run through the maze again until it is placed back in the startbox and the experimenter opens the startbox door); and *operant* because the response operates on the environment.

Free-operant responding lends itself to a recording method, the cumulative record, that gives a convenient, detailed picture of how responding changes over time. Most contemporary cumulative records are produced by computer,

but in the original cumulative recorder, illustrated in Figure 5-4, a roll of paper was threaded around a roller. A motor drove the roller at a constant speed, feeding out the paper. A pen or other writing device rested on the paper as it passed over the roller, and each response (e.g., a pigeon's key peck) moved the pen a small distance at right angles to the movement of the paper. Thus, at any time during a session this record shows the total responses accumulated.

Figure 5-5 shows examples of cumulative records. Because the record advances at a constant speed, its slope is steeper the higher the rate of responding, as illustrated in records A and B. In the scale for Figure 5-5, rates are roughly 30 and 12 responses/min in records A and B respectively. Record C includes only a few responses; the horizontal portions indicate periods of time without responses (note that a cumulative record can't have a negative slope, because the pen can record responses only by moving in one direction across the page).

In record D, a magnified section of record C that includes a few responses accompanies an event record on the same time scale. For each response in the event record, a small step occurs in the cumulative record; this property of cumulative records isn't always obvious because typical response and time scales are often too small for such fine resolution. Even so, different patterns of responding are easily distinguished. For example, response rates in E and F are roughly the same, but E is steplike whereas F is relatively smooth. This means that E was produced by short high-rate bursts of responding (steep segments) separated by pauses (flat segments), whereas F was produced

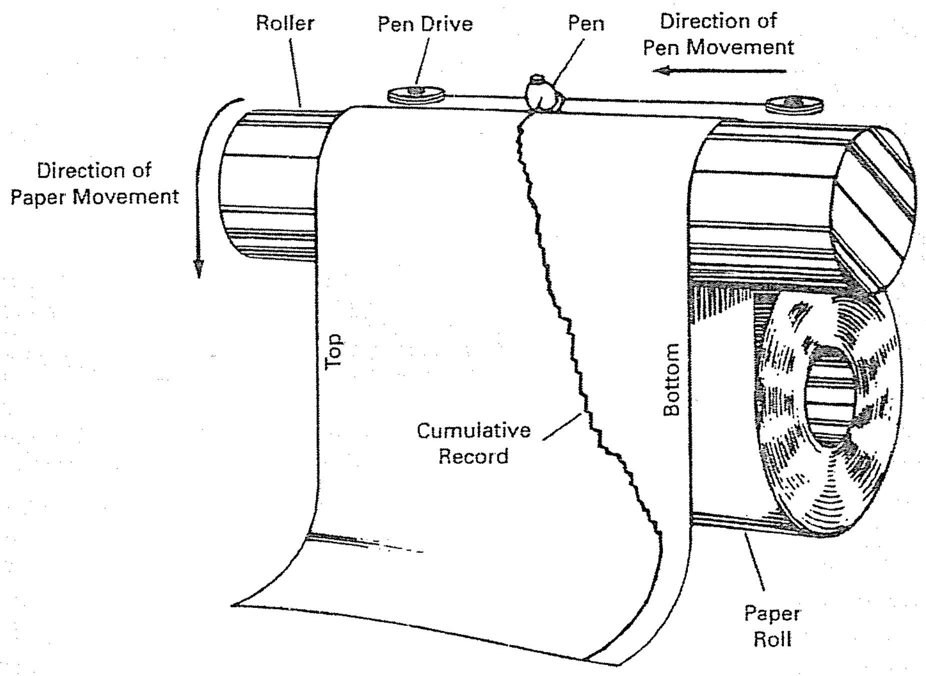

Figure 5-4 Main components of a cumulative recorder. A roller drives the paper at a constant speed and each response moves the pen a fixed distance across the paper. Paper speed and step-size per response vary with the behavior under study. A common scale is about 1 cm/min (about 2.5 min/in) and 1100 responses across the paper width (about 200 responses/in). At this scale, a slope of 45° represents a rate of about 40 responses/min. When the pen moves near the top of the paper, it resets automatically to its starting position near the bottom.

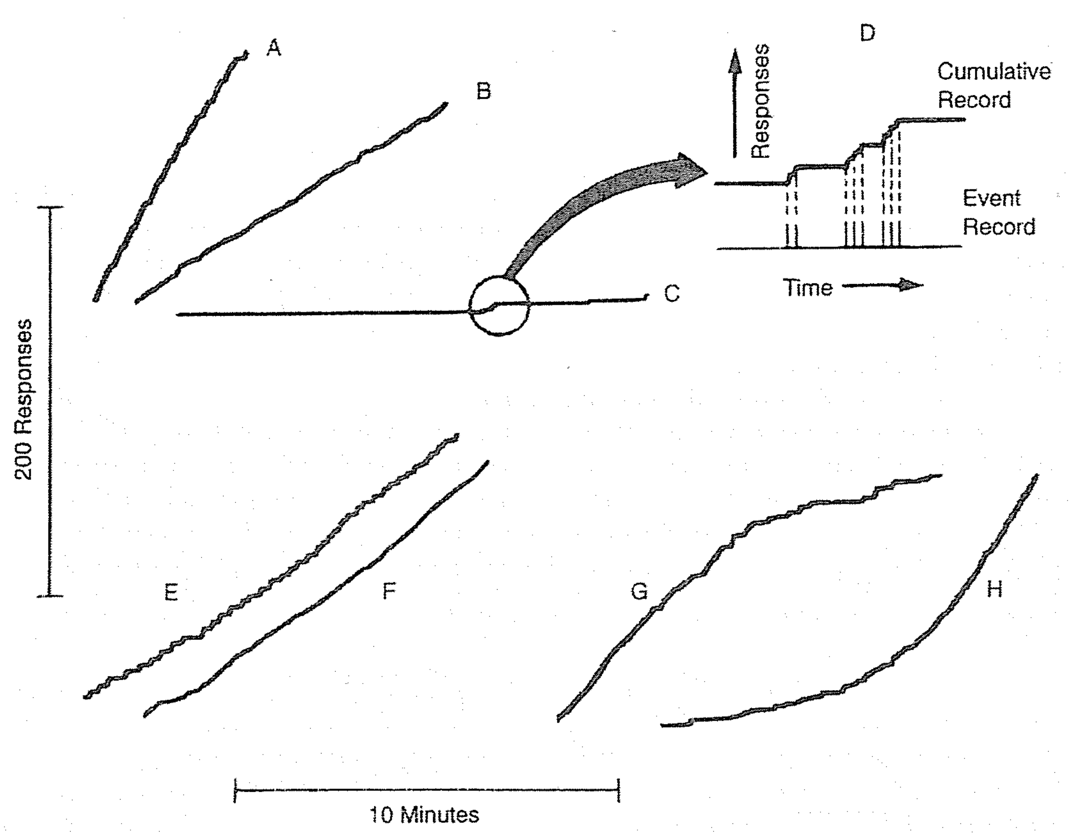

Figure 5-5 Sample cumulative records. A 45° slope represents about 20 responses/min. Records A and B differ mainly in response rate, higher in A than in B. Rate is zero through most of C; a magnified segment of C with a few responses is shown in D in relation to an event record. Records E and F are about equal in rate but show different patterns of responding: E is steplike, showing periods of responding alternating with pauses; the smoother-grained F shows relatively steady responding. Records G and H show rates that change over time, decreasing in G (negative acceleration) and increasing in H (positive acceleration).

by more uniform responding. This property is sometimes called *grain*; of the two records, E has a rougher grain than F.

Records G and H illustrate other properties of behavior made visible in cumulative records. In G, the rate begins at roughly 25 responses/min but gradually decreases as time passes; in H, it changes in the opposite direction, increasing from a relatively low rate to roughly 30 responses/min (records in which slopes decrease over time are called *negatively accelerated*; those in which they increase are called *positively accelerated*).

Figure 5-6 shows other features sometimes incorporated into cumulative records. Records A and B show how pen displacements can indicate other events besides responses. Here only some responses produced food, irregularly in A (as at *a*, *b* and *c*) and regularly in B (as at *d* and *e*). The repeated concave pattern in B, as between *d* and *e*, is sometimes called *scalloping*. In C, responding that began at *f* produced food at *g*, as indicated by the pen displacement. The pen then reset to *h* and the sequence repeated at *h* to *i* and so on. This type of record makes successive segments easy to compare (e.g., the segment ending at *g* contains more responses than the one ending at *i*). Record D shows how sustained pen displacements can distinguish different conditions. Here responding occasionally produced food only in the presence of a tone; the pen stayed in its normal position during the tone, in segments *j*, *l* and *n*, but was displaced downward in its absence, in segments *k*, *m* and *o*.

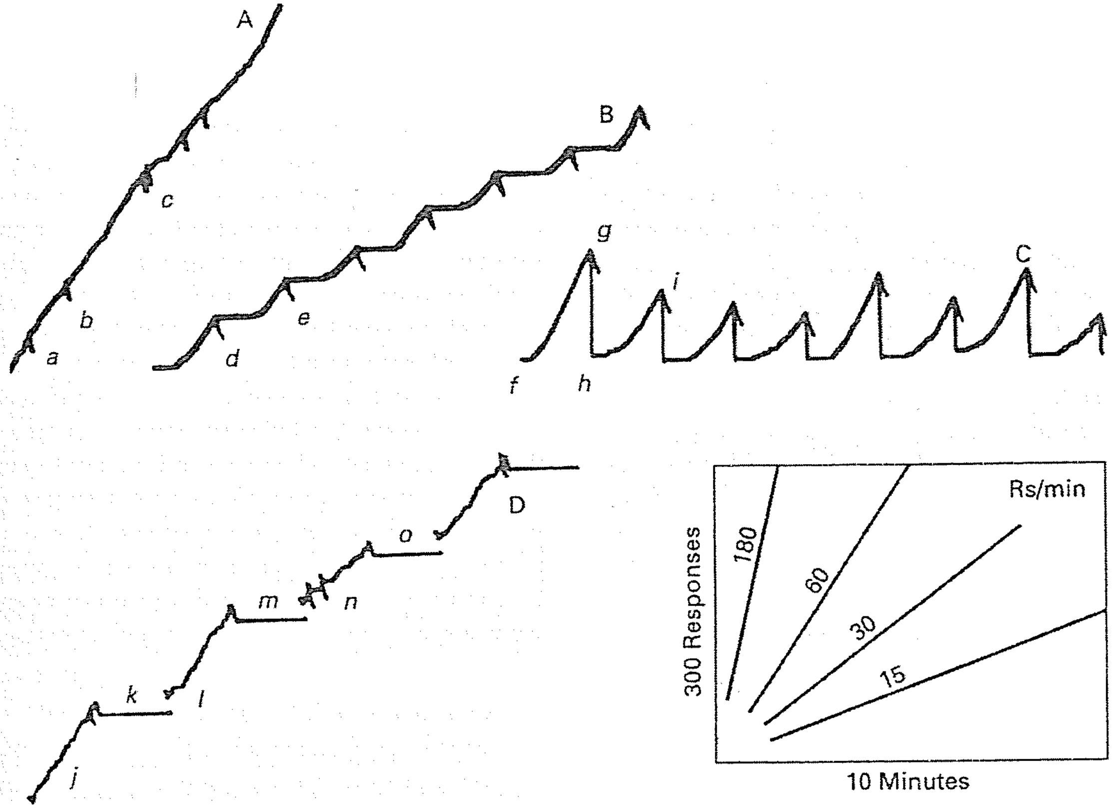

Figure 5-6 Additional features of cumulative records. In A and B, pen displacements superimpose a record of other events, such as food deliveries, on cumulative responses (as at *a* through *e*). In C, pen resets make it easy to compare successive segments of the record (*f* to *g*; *h* to *i*). In D, sustained pen displacements distinguish responding during a stimulus (at *j*, *l*, *n*) and nonresponding in its absence (at *k*, *m*, *o*). A slope of 45° represents about 40 responses/min (the scale differs from that in Figure 5-5).

With this treatment of free-operant behavior and cumulative records, we've explored part of the technology of the science of behavior. Before we can effectively consider the findings made available through such analyses, we must attend to some properties of the language of behavior.

The Vocabulary of Reinforcement

Lever pressing becomes more probable when a water-deprived rat's lever presses produce water than when they don't. Key pecking becomes more probable when a food-deprived pigeon's key pecks produce food than when they don't. And perhaps a child's cries become more probable when they produce a parent's attention than when they don't. These cases illustrate reinforcement: Responding increases when it produces reinforcers. The principle is simple, but as it evolved from Thorndike's initial versions of the Law of Effect to its contemporary status, it carried problems of language and logic with it. **Reinforcement is not a theory.** It is something that happens in behavior, and we must learn to spot it when it happens and not to confuse it with other things that happen that might seem superficially related to it.

Table 5-1 summarizes some properties of the contemporary vocabulary of reinforcement. This vocabulary includes the term *reinforcer* as stimulus and the terms *reinforce* and *reinforcement* as either operation or outcome. For example, when a rat's lever presses produce food pellets and lever pressing increases, we say that the pellets are reinforcers and that the lever presses are reinforced with pellets. **The response that increases must be the one that produces the consequences.** For example, if a rat's lever press produces shock and only the rat's jumping increases, it would be inappropriate to speak of either pressing or jumping as reinforced.

A reinforcer is a type of stimulus, but reinforcement is neither stimulus nor response. As a procedure or operation, reinforcement is presenting a reinforcer when a response occurs; it is carried out on responses, and so **we speak of reinforcing responses rather than organisms.** We say that food reinforced a rat's lever press or that a pigeon's key peck was reinforced with water, but not that food reinforced the rat or that the pigeon was reinforced for pecking or that a child was reinforced. The main reason for this restriction is illustrated in the last examples: When we speak of reinforcing organisms, it is too easy to omit the response or the reinforcer or both. The restriction forces us to be explicit about what is reinforced by what. Nor must we omit the organism; we can always say whose response it was (e.g., the child's crying).

TABLE 5-1 The Vocabulary of Reinforcement.
This vocabulary[a] is appropriate if and only if three conditions exist: (1) a response produces cosequences; (2) the response occurs more often than when it doesn't produce those consequences; and (3) the increased responding occurs *because* the response has those consequences.

Term	*Restrictions*	*Examples*
reinforcer (noun)	A stimulus	Food pellets were used as reinforcers for the rat's lever presses.
reinforcing (adjective)	A property of a stimulus	The reinforcing stimulus was produced more often than the other, nonreinforcing stimuli.
reinforcement (noun)	As an operation, the delivery of consequences when a response occurs	The fixed-ratio schedule of reinforcement delivered food after every tenth key peck.
	As a process, the increase in responding that results from the reinforcement	The experiment with monkeys demonstrated reinforcement produced by social consequences.
to reinforce (verb)	As an operation, to deliver consequences when a response occurs; responses are reinforced and not organisms	When a period of free play was used to reinforce the child's completion of school work, the child's grades improved.
	As a process, to increase responding through the reinforcement operation	The experiment was designed to find out whether gold stars would reinforce cooperative play among first-graders.

[a] A parallel vocabulary is appropriate to punishment (including *punisher* as a stimulus and *punish* as a verb; cf. Chapter 6), with the difference that a punishing consequence makes responding occur less rather than more often.

The term *reinforcement* has also often served as a name for an outcome or process, the increase in responding that follows from reinforcing consequences. This dual usage as both operation and outcome complicates things, but it has been around a long time and is unlikely to change. For example, the statement that a response was reinforced can be interpreted in two ways: Either the response produced a reinforcer (procedure) or responding increased because it produced a reinforcer (outcome). We'll favor the usage of reinforcement as a procedure. The outcome can be described so easily in terms of changes in responding (i.e., responding increased) that substituting other terms for a direct description of what happens to responding seems hard to justify. But the outcome or process usage has so much precedent that it can't be avoided. In addition, the overlap between procedure and outcome vocabularies extends to many other terms, such as *extinction*, *punishment*, and *conditioning* (Ferster & Skinner, 1957).

The vocabulary of reinforcement leads to some logical difficulties even when restricted to procedures. When a response becomes more likely because it has produced a stimulus, we say the response has been reinforced and we call the stimulus a reinforcer. If asked how we know the stimulus was a reinforcer, we point to the increase in responding. If we're then asked why the increase occurred, we might say it did so because the response was reinforced. Soon we begin to repeat ourselves. Once we define a reinforcer by its effect on behavior, we create a problem of circular definition if we simultaneously define the effect by the reinforcing stimulus (Meehl, 1950).

One solution is to recognize that the term *reinforcement* is descriptive, not explanatory. It names a relation between behavior and environment; it doesn't explain the relation. The relation includes at least three components. First, responses must have consequences. Second, their probability must increase (i.e., they must become more probable than when not having these consequences). Third, the increase must occur *because* they have these consequences and not for some other reason. For example, if we knew only that responding increased, it wouldn't be appropriate to say that the response must have been reinforced (maybe it was elicited). It wouldn't even be enough to know that the response was now producing some stimulus it hadn't been producing before. We'd still have to know whether responding increased *because* the stimulus was its consequence.

Suppose a parent attends to an infant whenever the infant makes cooing sounds, but suppose also that the infant is more likely to make them with the parent present than with the parent absent. The infant coos and the parent comes, and now the cooing increases. How do we decide whether the infant is now cooing because cooing has been reinforced or because the parent is now present and makes cooing more likely? It is even more complicated, because the infant's cooing may also reinforce the parent's attention. In working out these interactions, it is important to distinguish reinforcing effects of consequences from other effects (cf. Poulson, 1984).

Consider another and less pleasant example. Assume an abusive parent gets annoyed whenever an infant cries and tries to suppress the cries with slaps. The infant cries and then gets slapped and this produces even more crying. In this case, the consequence of crying is getting slapped and getting slapped produces more crying, but we wouldn't want to say the slapping reinforces the crying. Two criteria for reinforcement are satisfied but the third isn't. Crying doesn't increase because slapping is a consequence; the slapping brings on crying even if the infant isn't crying at the time of the first slap. Stimuli may have eliciting or other effects along with or instead of their effects as consequences of responding. (Under these unhappy circumstances, by the way, the infant may learn eventually to suppress the crying; as we'll see in Chapter 7, it would then be appropriate to say that the slapping punishes the crying.)

The vocabulary of reinforcement requires that responding has consequences, that responding increases, and that this increase occurs because responding has its consequences and not for other reasons. Once these conditions are met, we say the response was reinforced and the stimulus was a reinforcer.

We might also assume the stimulus will continue to be an effective reinforcer in the future and will reinforce other responses in other situations. But either assumption may be incorrect. The effectiveness of reinforcers changes over time, and any consequence may reinforce some responses but not others. For example, money is more likely than a smile to reinforce the services of a plumber or an electrician, but the opposite is likely to be the case if the behavior to be reinforced is a lover's embrace. Despite these reservations, the reinforcers used in many standard experimental situations (e.g., food with food-deprived organisms) are likely to reinforce a variety of responses; the experimenter who chooses a stimulus that will reinforce some responses but not others will sooner or later have to cope with the difference. More about the relativity of reinforcement will come up in the next chapter.

A cautionary note is in order. Positive reinforcement can be used to change a developmentally delayed child who engages extensively in self-injurious behavior into one who has learned communicative skills and therefore has been empowered to deal in more constructive ways with his or her caregivers. We might therefore think that positive reinforcement would be especially important to the practitioners of an approach called positive psychology. Unfortunately, they reject it along with the establishing events that make it effective; their rhetoric implies that contingent acts of kindness should always be replaced by random ones (Catania, 2001a; Seligman & Csikszentmihalyi, 2000, 2001).

EXTINCTION

The consequences of many responses remain reasonably constant throughout life. For example, we usually touch the objects we reach for and we usually get from one floor to another when we climb a flight of stairs. But for other responses, consequences change. Responses reinforced during childhood may no longer be reinforced in adulthood. Educational systems may arrange consequences such as praise or grades for solving arithmetic problems or answering factual questions, but sooner or later these artificial consequences are discontinued, with the hope that more natural consequences will maintain the responses when the student moves on to other settings. When a response is reinforced, its probability increases. But the increase isn't permanent: Responding decreases to its earlier levels when reinforcement is discontinued.

The discontinuation of reinforcement is called *extinction*; when this reduces responding to its earlier level the responding is said to be *extinguished*. This mainly demonstrates that the effects of reinforcement are temporary. Responding is maintained while reinforcement continues but not after it stops. Thus, the decrease in responding during extinction isn't a special process requiring a separate treatment; it is one of the properties of reinforcement.

At one time, responding during extinction was a primary measure of reinforcement. *Resistance to extinction* was expressed as the time elapsed until responding dropped to some specified level. Two hypothetical records of extinction of a rat's lever presses are shown in Figure 5-7. Response rate decreases over time in both (negative acceleration), but depending on the extinction criterion either could represent more resistance to extinction. If the criterion is the time until 2 min go by without a response, then A shows more resistance to extinction than B; A doesn't include 2 min without a response but such a period begins halfway through B. If instead the criterion is total responses emitted, then resistance to extinction is greater for B than for A. Because its definition permitted such ambiguities, resistance to extinction lost favor as a measure of the effects of extinction.

But resistance to change, of which resistance to extinction is a special case, remains an significant property of behavior (Nevin, 1992). Two responses might each be maintained at similar rates, but one might be persistent in the face of changes in other events, such as the introduction of reinforcement arranged for different responses, whereas the other might be fragile. For example, arithmetic facts, spelling and other academic skills are said to be *fluent* when they have been practiced and reinforced to the point where they occur with high accuracy and short latency (Johnson & Layng, 1992). Once such skills become fluent, they also become less likely to be disrupted by changes in settings or other distractions. Thus, fluent skills are far more resistant to change than those that have not been learned to a fluency criterion.

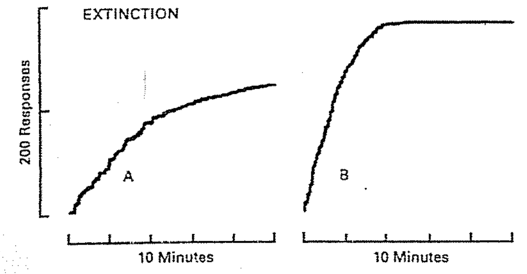

Figure 5-7 Hypothetical cumulative records of extinction of a rat's lever presses after food reinforcement. Either A or B might be said to demonstrate greater resistance to extinction, depending on whether it is measured by the time taken until 2 min pass without a response or by total responses during the session.

As a rough approximation, we can say that the total activation of responding over its reinforcement history determines its resistance to change, whereas its rate is determined by its context, including such current variables as the reinforcement of other response classes (cf. Nevin, McLean, & Grace, 2001b). Whatever its starting point, extinguished responding tends to decrease by a constant proportion over equivalent times, roughly according to the equation:

$$R = k \log t$$

where R is the response rate, *t* is the time in extinction, and *k* is a dimensional constant (e.g., Catania, 2005b; Catania & Keller, 1981; Katz & Catania, 2005; Nevin & Grace, 2005). We will later see, however, that the course of extinction can be influenced by schedules of reinforcement (Chapter 15).

Extinction Versus Inhibition

If extinction didn't occur, the effects of reinforcement would be permanent. Any responding engendered by reinforcement would last through a lifetime. Clearly that cannot be generally so. For example, if you wear a watch you probably often turn your wrist to look at it; the consequence of looking is seeing the time. But if you stop wearing the watch for some reason, you'll eventually stop looking; seeing a bare wrist isn't an effective reinforcer.

The history of the concept of extinction, however, wasn't so simple. It was long assumed that extinction actively suppressed responding. Extinction was said to have *inhibitory* effects, in contrast to assumed *excitatory* effects of reinforcement. This treatment went back to a language that had been applied to data from Pavlov's conditioning experiments (cf. Chapter 17). The effectiveness of Skinner's operant language depended in part on escaping from the implications of those earlier usages (Skinner, 1938, pp. 96–102). Once those usages had carried over to the language of consequences, they were kept because they seemed consistent with other effects that often accompanied extinction. Thus, texts on learning tended to devote separate chapters to reinforcement and extinction rather than treating them as two aspects of a single phenomenon.

Consider *spontaneous recovery*. In a typical extinction session, responding decreases as the session continues. But the rate at the beginning of the next extinction session is usually higher than it was at the end of the last one. Some hypothetical cumulative records illustrating spontaneous recovery are shown in Figure 5-8. Responding at the start of a session was said to have recovered spontaneously from inhibition built up by the end of the last session; it was assumed that this inhibition, actively suppressing responding, increased within sessions but dissipated between sessions.

Phenomena such as spontaneous recovery were taken to mean that responding that had been reduced by extinction was somehow "there all the time but inhibited" (R. L. Reid, 1958). Various

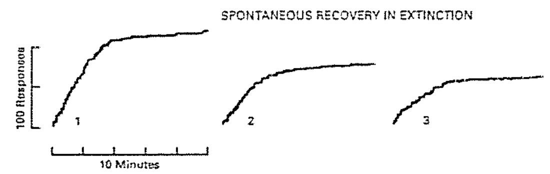

Figure 5-8 Spontaneous recovery in hypothetical cumulative records of a rat's lever presses in sessions of extinction after food reinforcement. The response rate at the start of session 2 is higher than it was at the end of session 1; similarly, the rate at the start of session 3 is higher than it was at the end of session 2.

accounts of extinction were formulated in terms of inferred processes such as inhibition, frustration, interference or fatigue (Kimble, 1961), but when a response was said to be inhibited in extinction, there was no way to measure what was doing the inhibiting. Later we'll consider a different variety of inhibition in which both what is inhibited and what does the inhibiting are clearly specified.

It wasn't necessary to assume that extinction required active suppression. For example, the effects of presession conditions such as handling may make the start of a session different from later times. If so, effects of extinction late in one session might not transfer to the start of the next session. On these grounds, Kendall (1965) reasoned that the usual pattern of response rates in extinction sessions could be reversed under the right conditions. The key pecks of three pigeons were first reinforced in 1-hr sessions. Repeated 1-min sessions of extinction followed. The first long extinction session came only after responding had reliably decreased to zero in the brief sessions. Within a few minutes, each pigeon began to respond. Until this session, responding had never extinguished at times later than 1 min into a session; responding occurred at these later times when the opportunity was finally available. In a sense, Kendall had demonstrated spontaneous recovery within a session rather than at its start.

Another example of recovery of extinguished responding has been called *regression* or *resurgence* (Epstein & Skinner, 1980; Keller & Schoenfeld, 1950, pp. 81–82). Suppose a rat's chain pulls are extinguished and then lever presses are reinforced. If the lever presses are later extinguished the previously extinguished chain pulls are likely to reappear. By analogy to clinical terminology, the phenomenon suggests regression from current behavior (lever presses) to older behavior that was once effective (chain pulls). It is of considerable practical significance. For example, in training show dogs it is sometimes necessary during a show to withhold reinforcers that had been available during training, but this must be done carefully to prevent regression to behavior that had been common earlier but would now spoil the dog's performance.

Response-Reinforcer Contingencies and Reinforcer Deliveries

Contingencies expressed in terms of probability relations between responses and their consequences can be summarized in graphic form much like those between stimuli and the responses they elicit (Figure 4–2). The coordinate system is illustrated in Figure 5–9. The y-axis shows the probability of a stimulus given a response, p(S/R); the x-axis shows its probability given no response, p(S/no R). Relative to Figure 4–2, the S and R terms have been reversed. The earlier figure showed effects of stimuli on responses; this one shows effects of responses on stimuli.

At A, the probability of the stimulus is high given a response but is otherwise low, as when a rat's lever presses produce food. At B, stimulus probability is independent of responses, as when food is delivered without regard to lever presses. At C, stimulus probability is zero whether or not a response has occurred, as when food is discontinued in extinction. We'll later consider other kinds of contingencies in other contexts. For example, cases in which responses reduce stimulus probability, as at D, illustrate avoidance (Chapter 8), and those in which responses produce a stimulus with a probability of less than 1.0, as at E, illustrate reinforcement schedules (Chapter 15).

Discontinuing reinforcement has not one but two effects: (1) the contingency between responses and reinforcers ends, so (2) reinforcers are no longer delivered. In this context, the term *contingency* simply describes the consequences of responding; here it is *the effect of a response on stimulus probability*. For example, in a contingency in which a rat receives food if and only if it presses a lever, a lever press raises the probability of food from zero to 1.0, but in a contingency in which lever presses do nothing, the probability of food is independent of lever presses. (Strictly, a response-stimulus contingency is virtually always part of a three-term contingency, but we needn't address that issue here; cf. Chapter 11.)

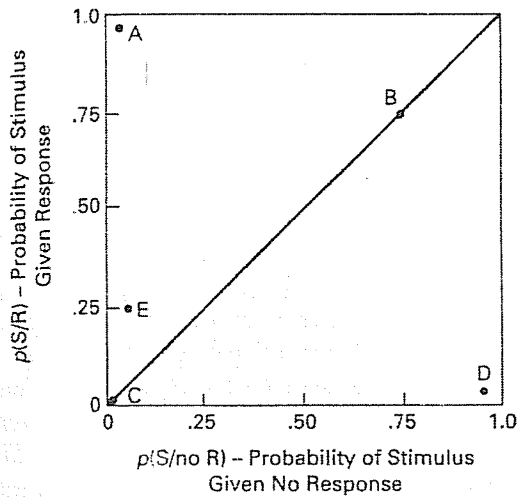

Figure 5-9 Response-stimulus contingencies represented in terms of stimulus probability given a response, p(S/R), and stimulus probability given no response, p(S/no R). The graph includes reliable production of stimuli by responses (A), response-independent stimuli (B), extinction (C), prevention of stimuli by responses, as in avoidance (D: see Chapter 8), and intermittent production of stimuli by responses, as in reinforcement schedules (E: see Chapter 15). Cf. Figure 4–2.

Side Effects of Extinction

If the standard extinction procedure terminates both a contingency and reinforcer deliveries, how much of what happens depends on the end of the contingency and how much on the end of food deliveries? In extinction, the rat's responses no longer do anything, but the rat is also no longer eating. Taking away reinforcers will affect more than just the reinforced response. If food is suddenly taken away from a food-deprived rat that has been eating, for example, the rat becomes more active and perhaps urinates or defecates. If food was produced by lever presses, the rat might bite the lever (Mowrer & Jones, 1943). If other organisms are in the chamber, the rat might attack them (Azrin, Hutchinson, & Hake, 1966). And the opportunity to engage in such aggressive responses may reinforce other responses. For example, an organism might pull a chain if chain pulls let it get at something it can sink its teeth into (Azrin, Hutchinson, & McLaughlin, 1965). These effects, though observed in extinction, aren't results of terminating the reinforcement contingency. They occur when response-independent or free food deliveries stop as well as during extinction. We call them side effects because they are indirect products of the change in contingencies. Whether after free food or food produced by responses, a rat that has been eating can no longer do so. In extinction, these side effects get superimposed on decreases in the previously reinforced responding because the termination of reinforcers is necessarily a part of extinction.

These side effects had been thought to show that extinction was more than evidence for the temporary effect of reinforcement. Yet many, such as the aggressive responding generated by terminating reinforcer deliveries, could have been observed in situations that removed reinforcers but didn't involve response consequences. These observations have crucial implications for the application of reinforcement and extinction procedures. For example, those who work with children sometimes use free or response-independent reinforcers rather than extinction to avoid the side effects of terminating reinforcer deliveries (e.g., Hart *et al.*, 1968, on the social reinforcement of a child's cooperative play).

Extinction was long regarded as the most appropriate way to get rid of problematic behavior, but as the side effects of extinction were increasingly recognized that view has fallen from favor. You can see the difference even across the earlier editions of this book. While I was preparing the third edition, my colleague Eliot Shimoff urged me to use more human examples and argued that it would be helpful to include not just anecdotes and hypothetical cases but also actual human data presented in figures. I agreed that it was a good idea and began searching out candidate data sets, first in the *Journal of the Experimental Analysis of Behavior* and then in the journal where it might be assumed they'd be more likely to appear, the *Journal of Applied Behavior Analysis*. Here and there I found an case of the extinction of human responding, but typically embedded in a complex data presentation. A clear description of the procedure and outcome seemed to require a distracting detailed elaboration and even those examples were rare, so you won't find any human extinction data in this edition either.

At the time I was becoming heavily involved in a new Masters track in Applied Behavior Analysis in my department. It was a joint effort with the behavioral programs at the Kennedy-Krieger Institute in Baltimore (Catania, DeLeon, & Cataldo, 2001). My colleagues there were much involved in work with self-injurious and other serious problem behavior in children with various developmental disabilities, and so I consulted them about their use of extinction. What I learned was that their procedures had evolved over time to minimize its use and its attendant side effects in favor of effective alternatives that did not produce such effects.

Dealing with problem behavior presented by children with severe developmental disabilities is hard enough. Why deal also with aggressive behavior or other emotional side-effects of extinction, especially when other techniques are available (cf. Lalli, Casey, & Kates, 1997; Lerman, Iwata, & Wallace, 1999; Thompson *et al.*, 2003; see also DRO in Chapter 9)? The reason I hadn't been able to find appropriate extinction data in the applied literature was that extinction was not a preferred procedure, so few data were collected and presented. What, then, are the alternatives to extinction?

Extinction Versus Free or Noncontingent Reinforcement

As already noted, discontinuing reinforcement in extinction has two components: (1) it terminates a contingency between responses and reinforcers and (2) reinforcers are no longer delivered. The latter is responsible for some unwelcome side effects of extinction. The contingency can be terminated without terminating the delivery of reinforcers by arranging the free or noncontingent delivery of reinforcers (noncontingent reinforcement is sometimes abbreviated NCR). Figure 5–10 illustrates differences between contingent reinforcement, extinction, and free reinforcement or NCR. Baseline conditions, left of the dashed vertical line, show contingent reinforcement (for illustrative purposes, the three cases are shown as identical). This hypothetical example assumes a rich intermittent schedule in which reinforcers are delivered after a small but variable number of responses.

On the right, the top frame shows continued contingent reinforcement, the middle shows extinction, and the bottom shows the free (noncontingent) delivery of reinforcers, with their timing matched or yoked to those in the top frame. When free reinforcers are delivered, they may follow responses closely but they occur at varying times after responses. Because these reinforcers do not depend on responses, they do not typically maintain responding. The top and bottom right cases are alike in that they both include reinforcers, whereas the middle and bottom right cases are alike in that neither includes a response-reinforcer contingency.

Now consider the emoticons, which illustrate an important side effect: in all other conditions, food-deprived rats who had been eating continue to eat; in extinction they no longer do so. Discontinuing the food presentations has behavioral effects sometimes described as emotional. For example, during extinction the opportunity to attack another organism may function as a reinforcer (Azrin *et al.*, 1966). This figure shows that **the discontinuation of reinforcers is what produces the emotional behavior; the response-reinforcer contingency has nothing to do with that side effect.**

The contingency between responses and reinforcers is terminated in both extinction and free or noncontingent reinforcement in Figure 5–10. Relative to the top frame both the middle and bottom frames show responding that decreases over time. In other words, **responding decreases once the response-reinforcer contingency is broken, whether or not reinforcers are delivered.** For convenience of exposition, the response decrements in the middle and bottom right frames are shown as equal, but in practice responding usually declines somewhat more slowly with free reinforcers than with extinction, especially on the first as opposed to later exposures to these procedures.

Why the slower decline with free reinforcers than with extinction? First, a rat's lever presses are reinforced with food; then, when presses no longer produce food, food deliveries continue independently of behavior. As lever pressing continues, some are still followed very closely by food. Lever pressing declines slowly because these

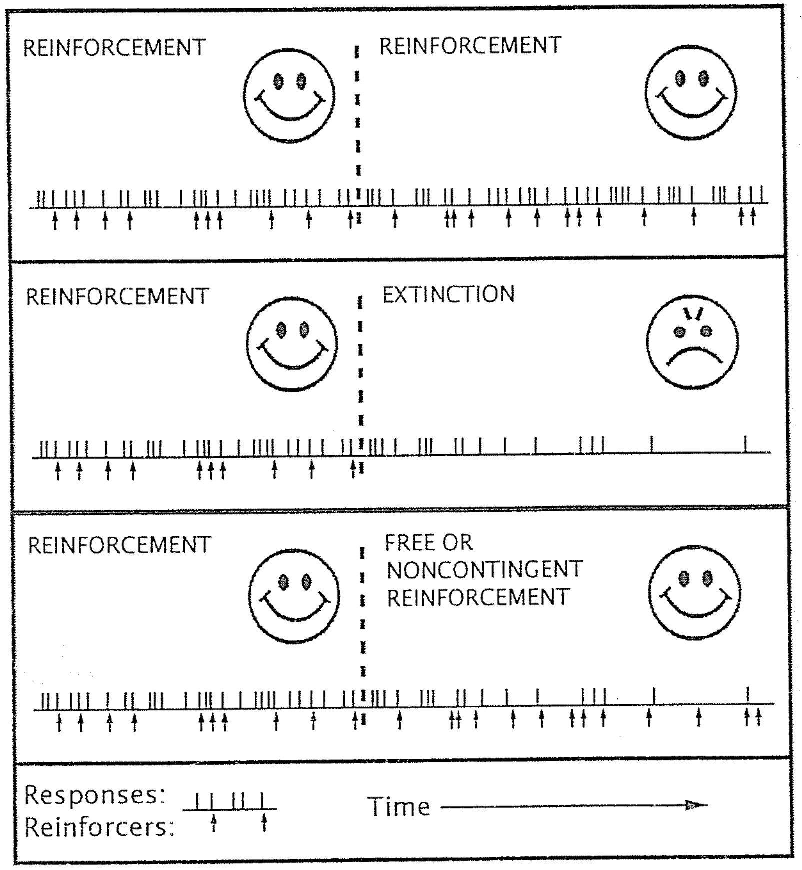

Figure 5-10 Schematic display of hypothetical responding maintained by intermittent reinforcement (left) followed after the dashed vertical line by one of three conditions: continuation of intermittent reinforcement, extinction, or free reinforcement (also called NCR or noncontingent reinforcement). With both extinction and free reinforcement, rate of responding decreases because the response-reinforcer contingency is discontinued. The emoticons represent different side effects of the contingencies: The organism is happy as long as it continues to eat whether the food is response-contingent or not, but it is not happy if the food stops.

frequent incidental successions of responses and reinforcers counteract the effects of terminating the reinforcement contingency. Eventually lever pressing drops out, but extinction would probably have been quicker. For this reason, arranging a transition from a reinforcement contingency to response-independent reinforcer deliveries has too often been regarded as a poor way to examine the effects of terminating reinforcement contingencies (e.g., Boakes, 1973).

Again, there are practical implications, and free reinforcers are not the only option. Suppose an institutionalized developmentally disabled boy often engages in self-injurious behavior such as head-banging or eye-poking. Suppose further that we discover that his behavior is in large part maintained by staff attention as a reinforcer. Extinction is not well-advised because of the harm he might do to himself if we ignore the self-injurious behavior. Giving him attention independently of the self-injurious behavior is one possibility, but that might reduce the behavior only slowly. An alternative procedure is to use attention to reinforce an alternative response, and especially one that is incompatible with the self-injurious behavior (Repp & Deitz, 1974). The self-injurious behavior will inevitably decrease as the alternative response increases. Such procedures, sometimes called *differential reinforcement of other behavior*, have been widely applied to problem behavior (e.g., see Skiba, Pettigrew, & Alden, 1971, on thumbsucking, and DRO in Chapter 9).

Reinforcers are important, after all. If attention has contributed to increases in the self-injurious behavior of a developmentally delayed child, that tells us how important attention is to the child. The child will be better served if we shape alternative and more functional ways of commanding attention than if we simply remove attention (e.g., Carr, Severtson, & Lepper, 2009; Goh, Iwata, & DeLeon, 2000; Hagopian *et al.*, 2000). Unfortunately, generations of students given cursory accounts of behavioral methods in introductory textbooks seem to have taken away as the primary message that if you see a child doing something you don't approve of you shouldn't reinforce it. Instead, you should just ignore it. But free reinforcers coupled with the shaping of other behavior are the better alternative, and though it might seem more complicated to arrange these contingencies, the absence of the emotional side effects of extinction usually more than makes up for the difference. Even better, of course, is not getting to the point of being tempted to impose extinction. One way to reduce a child's bad behavior is to reinforce good behavior, and that is one reason why "Catch 'em being good" is such an effective slogan for either parents or teachers.

ADDENDUM 5A: EXTINCTION AND SUPERSTITION

This is an addendum because superstition, in its behavioral sense, is no longer regarded as the ubiquitous and enduring phenomenon it was once thought to be. It happens, but it is typically transient. Ordinarily, organisms are exquisitely sensitive to the difference between events that are produced by their behavior and events that by coincidence follow the behavior. We humans are sometimes an exception, but that is because our human superstitions involve much more than coincidence; they are typically intricately interwoven with verbal behavior as well as natural contingencies. This addendum will not consider those kinds of superstitions.

In extinction, the effects of terminating response-reinforcer contingencies are combined with the more general effects of terminating reinforcer deliveries. Why then has extinction remained the primary way to study the effects of terminating contingencies for so long? Disconnecting the lever from the feeder is more convenient than disconnecting the lever and at the same time substituting a clock that operates the feeder periodically, but it is unlikely that the answer lies in a mere change in apparatus. It more likely lies with some other properties of behavior.

In the phenomenon called *superstition* (Skinner, 1948), food is repeatedly presented to a food-deprived pigeon at short intervals (e.g., every 10 or 15 seconds). Responses that occur just before food deliveries are likely to be repeated and therefore to be followed closely by still more food deliveries. The effect of the accidental succession of responses and reinforcers

> is usually obvious. The bird happens to be executing some response as the hopper appears; as a result it tends to repeat this response. If the interval before the next presentation is not so great that extinction takes place,... [this] strengthens the response still further.... The bird behaves as if there were a causal relation between its behavior and the presentation of food, although such a relation is lacking. (Skinner, 1948, pp. 168–171)

Skinner noted that the topography or form of each pigeon's responding usually changed gradually as the procedure continued, as accidental relations developed between responding and food deliveries. Skinner referred to these changes as *topographical drift*. A pigeon responded temporarily as if its responses were producing food as a consequence, but no particular response consistently remained as superstitious behavior. Over long periods of time with this procedure, Staddon and Simmelhag (1971) observed that pecking often dominates as the response just preceding food deliveries (cf. Chapter 4 on elicitation).

But drift itself demonstrates that the effects of accidental correlations between responses and reinforcers are typically transient. The evidence that free reinforcers do not ordinarily maintain behavior has become substantial. Nevertheless, worries about creating and maintaining superstitious patterns have been persistent. Superstitious responding generated by accidental successions of responses and reinforcers was seen as a recurrent problem because such accidental sequences can occur whether reinforcers are independent of responses or are their consequences.

If one response is followed by a different response that is reinforced, the reinforcer may affect both even though its delivery depends only on the second one (Catania, 1971; Kazdin, 1977). Even when responses have consequences, properties unrelated to reinforcement may become stereotyped if they typically accompany reinforced responses. For example, a bowler's gestures after releasing the ball may persist because they are so often closely related to earlier responses and to the impact of the ball (Herrnstein, 1966). Furthermore, if reinforcement contingencies change so that features of responding that once were relevant become irrelevant and if these old features aren't incompatible with the newly relevant ones, the old ones may persist simply because they continue to be followed very closely by reinforcers (Stokes & Balsam, 1991). Still another difficulty is that superstition is too easily invoked to explain behavior not accounted for in other ways (Guthrie & Horton, 1946; Moore & Stuttard, 1979).

Accidental contiguities of responses and noncontingent reinforcers should sometimes have effects similar to those of scheduled contiguities in which reinforcers are contingent upon responding, because the organism contacts only the temporal relations between responses and reinforcers and not the contingencies that generated them. It is not obvious that organisms can be sensitive to the sometimes subtly different temporal distributions of response-reinforcer intervals produced by contingent and noncontingent procedures, and yet they are (Catania & Keller, 1981; Katz & Catania, 2005).

If free reinforcers affect behavior early in a transition from contingent to free reinforcers, then responding will decrease more slowly than in extinction. This is not a big enough concern to rule noncontingent reinforcement out of consideration in either experimental or applied settings. If higher or lower rates of noncontingent reinforcement are available as an option, this concern favors the lower rates. But if behavior persists for long periods of time under such arrangements it is more appropriate to look for other sources of the behavior than to attribute it to superstition arising from the adventitious correlations of responses and reinforcers.

Chapter 6
Reinforcers as Opportunities for Behavior

Reinforcers are sometimes discussed in terms of properties such as strength and resistance to change. As a relative of strong *and* tough, *strength has Old English roots. Resistance, in the sense of a staying in place, is one of a large family of words related to the Indo-European* sta-, *to stand. Its many relatives include* constant, contrast, substitute, superstition, exist *and* system.

We've identified reinforcers only in terms of their effects, so without making a stimulus a consequence of responding we can't say if it will serve as a reinforcer. But even demonstrable reinforcers such as food vary in effectiveness depending on deprivation. Food as a consequence of lever presses might not do much to the rat's rate of pressing if the rat has already had plenty to eat. Yet reinforcers are inevitably oversimplified if we treat them merely as stimuli. Presenting any reinforcer involves a transition from one situation to another, and we'll see that an important property of a reinforcing situation is the responding for which it sets an occasion.

Chapter 2 introduced establishing operations, the procedures that make events more or less effective as reinforcers. Chapter 4 discussed some effects of establishing operations as examples of *motivation* or *drive*. We can now define these terms more precisely. When we study motivation or drive we're concerned with what makes consequences more or less effective as reinforcers or as punishers. In the taxonomy of establishing operations, deprivation and satiation are important ways but not the only ways to change the effectiveness of stimuli as reinforcers or punishers. Let's now survey the variety of events that may serve as reinforcers.

Reinforcers are sometimes distinguished on the basis of the types of operations that established them. For example, a *conditional reinforcer* is one that has become effective by virtue of its relation to some other reinforcer (e.g., the light that comes on when a pigeon feeder is operated will eventually become a conditional reinforcer because of its relation to food delivery). The establishing operation here is that of arranging the relation between the stimuli (i.e., setting up the feeder so that food deliveries are accompanied by the light). In human behavior, money often works as a conditional reinforcer, and because of its relation to so many other possible reinforcers (all of the things we've bought with it), it is sometimes called a *generalized* conditional reinforcer. A reinforcer that doesn't depend on a relation to other reinforcers is called an *unconditional reinforcer*.

KEY TERMS: Relativity of Reinforcement; Premack Principle; Conditional or Conditioned Reinforcers; Extrinsic and Intrinsic Reinforcers; Sensory-Motor Learning.

Many events that are regarded as unconditional reinforcers are events with obvious biological significance, such as food, water, and sexual contact (cf. Richter, 1927). But reinforcers are not limited to events of obvious biological significance. For example, sensory stimuli such as flashing lights can powerfully reinforce the behavior of some children on the autism spectrum (Ferrari & Harris, 1981). And when developmentally disabled children are briefly deprived of music or social praise, these events may become more effective reinforcers of other behavior such as operating a switch; conversely, when the children are satiated with music or social praise, the effectiveness of these events as reinforcers decreases (Vollmer & Iwata, 1991).

Reinforcers have also been distinguished on the basis of their relation to responses. An *intrinsic* reinforcer (sometimes also called an *automatic* reinforcer) is one that has a natural relation to the responses that produce it (as when a musician plays because of the music that the playing produces), whereas an *extrinsic* reinforcer (sometimes also called a *contrived* reinforcer) has an arbitrary relation to those responses (as when the musician plays for money). The term *extrinsic* has also been applied to stimuli presumed to function as reinforcers because their function has been instructed (as when a child is told that it is important to earn good grades); despite their label, such stimuli are often ineffective reinforcers.

We earlier discussed eliciting, discriminative and reinforcing functions of stimuli. The presentation of stimuli can also have establishing functions (cf. Michael, 1982). Consider two examples: Tasting unsalted soup doesn't make it more likely that salt will be passed when you ask for it, and arriving in front of a locked door doesn't make it more likely that you'll find the key in your pocket. But you will be more likely to ask for the salt or to reach into your pocket. In each case something that had been neutral (the salt or the key) has become reinforcing. Such effects have sometimes been called *incentive* functions; the term *incentive*, however, has been applied both to the establishing and to the discriminative functions of stimuli (e.g., Bolles, 1975; Logan, 1960).

THE RELATIVITY OF REINFORCEMENT

Reinforcers exist in great variety. Some are consumed. Others aren't. Some appear effective on the organism's first experience with them. Others acquire their reinforcing properties during the organism's lifetime. No common physical properties allow us to identify reinforcers independently of their effects on behavior. For example, it is difficult for us to say what it is about teacher attention that reinforces student behavior, but we do know that when a teacher moves to a first-grader's desk with perhaps a pat on the shoulder or an encouraging comment contingent on the student's reading, the student's study behavior increases and other nonstudy behavior such as dawdling decreases (Hall, Lund, & Jackson, 1968). We also know that changes in student behavior can reinforce the behavior of the teacher (Sherman & Cormier, 1974).

It is tempting to equate reinforcers with events that colloquially are called rewards, but that would be a mistake. Reinforcers don't work because they make the organism "feel good" or because the organism "likes" them. Our everyday language doesn't capture what is important about reinforcers. For example, in a study of the reinforcers that might be effective in managing the behavior of people with profound handicaps, predictions based on staff opinion of what would work for each individual were inconsistent with the reinforcers identified by systematically assessing each individual's preferences among those events (Fisher *et al.*, 1992; Green *et al.*, 1988).

Some events that superficially seem "rewarding" may not function as reinforcers; others that seem the opposite may have powerful reinforcing effects. For example, falling from high places or being violently twisted and shaken hardly seem like potential reinforcers. Yet they surely contribute to the reinforcing effects of roller coasters and other amusement-park rides. Restraint also seems an unlikely reinforcer, but an analysis of the self-injurious behavior of three children with severe developmental disabilities showed that the placing of restraints that prevented the children from poking or biting themselves could reinforce arbitrary responses such as putting marbles in a box (Favell, McGimsey, & Jones, 1978). Once such a reinforcer is identified, it can be used to reinforce behavior that is incompatible with self-injury.

Laughter seems like a reinforcing consequence for the telling of jokes. Suppose an instructor tells some jokes, the class laughs, and as a result the instructor tells jokes more often. We can say that the laughing reinforced joke telling, but just on this evidence we can't say that laughter in general is a reinforcer (cf. Provine, 2000). Suppose now the instructor puns, the class laughs, and as a result the instructor puns less often. The laughing didn't reinforce the punning (in fact, we should say it punished the punning: Chapter 7). Whether laughter was a reinforcer or a punisher depended on whether it was contingent on telling jokes or contingent on punning. Actually, punning is more likely to be reinforced by groaning than by laughing. Suppose the instructor puns, the class groans, and as a result the instructor puns more often. Now we can say that the groaning reinforced the punning. Depending on whether the consequences are laughter or groans, punning is either reinforced or punished. In fact, laughing at puns can be bad enough to make a groan man cry. (In previous editions of this text, I had to correct copy editors who repeatedly changed the spelling of *groan* to *grown*.) **The effectiveness of a reinforcer depends on its relation to the responses that produce it.**

When a rat's lever press produces food, the food affords the rat an opportunity to eat. If we make a lever and food simultaneously available to the rat, the rat is more likely to eat than to press. What if the probability of a response goes up if it provides an opportunity to engage in another response more probable than itself (Premack, 1959, 1971)? Another way of putting it is to say that if response A is more probable than response B, an opportunity to engage in A will reinforce B. If this is so, food is an effective reinforcer for a food-deprived rat's lever presses simply because eating is usually more probable than pressing.

Consider an experiment that reversed the reinforcing effects of two stimuli by varying the probabilities of the responses occasioned by each (Premack, 1962). A rat's running in a wheel was controlled by engaging or releasing a brake on the wheel. The rat's drinking was controlled by moving a drinking tube into or out of an opening in a stationary wall on one side of the wheel; licking was recorded with an electrical device called a drinkometer. As tested during brief periods when both responses were available, running became more probable than drinking after the wheel was locked with water freely available, but drinking became more probable than running after the drinking tube was removed with the wheel freely available. In each case, the opportunity to engage in the more probable response reinforced the less probable response. When running was more probable than drinking (after the wheel was locked), licking became more likely if it released the brake and allowed the rat to run than if it didn't allow an opportunity to run. Conversely, when drinking was more probable than running (after the drinking tube was removed), running became more likely if it produced the drinking tube and allowed the rat to drink than if it didn't allow an opportunity to drink.

Thus, reinforcers cannot be defined independently of the responses that they reinforce. In Premack's experiment, drinking reinforced running when drinking was more probable than running, but running reinforced drinking when the probabilities were reversed. According to this account, reinforcers are relative and their reinforcing properties are determined by the responses for which they provide an opportunity.

This relativity had long been unrecognized. Most learning experiments had been restricted to responses of relatively low probability (e.g., a rat's lever press) and to reinforcers that occasioned highly probable responses (e.g., food and eating). These were common and convenient but nonetheless special cases. Few thought to ask whether situations might be arranged in which opportunities to press a lever or peck a key could be used to reinforce eating (e.g., Sawisch & Denny, 1973). The question isn't so far-fetched. The opportunity to eat a good meal can be an effective reinforcer, but how often are children persuaded to finish their dinners by making other activities depend on that eating? Eating can reinforce, as when a child's dessert depends on whether the child has completed homework, or it can be reinforced, as when the opportunity to go online depends on whether the child has finished dinner.

We can illustrate the relativity of reinforcement further by expanding the last experiment to three responses. Let's add a feeder to the running wheel and the drinking tube. By restricting access appropriately, we can make eating more probable than running and running in turn more probable than drinking. We then find that running can be reinforced by the opportunity to eat but the opportunity to run can reinforce drinking; in other words, running can simultaneously reinforce and be reinforced. These relations are shown on the left in Figure 6-1. If with water deprivation we now make drinking the most probable response, the reinforcement relations change, as shown on the right in Figure 6-1. By changing the relative probabilities of the three responses, we can make an opportunity to engage in any one an effective reinforcer with respect to either or both of the others. The relative likelihoods with which different children prefer coloring in books, building with blocks or playing in the playground can be useful information to someone who has to manage the behavior of children in daycare centers or the early grades of elementary school (e.g., Wasik, 1970).

Deprivation makes reinforcers effective because the probability of a response ordinarily goes up when we restrict the organism's opportunities to engage in it (e.g., Timberlake, 1980). But the detailed operation of Premack's principle has engendered controversy, mainly because its operation depends on how probabilities are calculated. Choice among simultaneously available responses may be more satisfactory than the proportion of time occupied by each (e.g., Dunham, 1977; Eisenberger, Karpman, & Trattner, 1967). A further complication is that some responses are more likely than others to substitute for each other (Bernstein

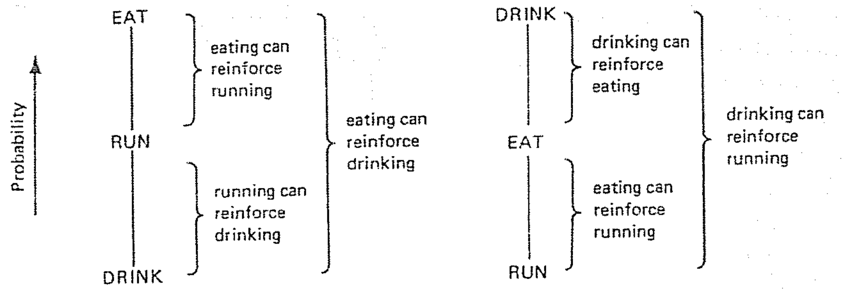

Figure 6–1 Reinforcement relations given different response probabilities in a behavior hierarchy. When eating is most probable and drinking is least probable (left), an opportunity to eat reinforces running or drinking but an opportunity to run reinforces only drinking. At another time (right), when drinking is most probable and running is least probable (e.g., after water deprivation), eating still reinforces running but both eating and running can now be reinforced by an opportunity to drink.

& Ebbesen, 1978; Rachlin & Burkhard, 1978). For example, deprivation of the opportunity to eat one food may not make its eating effective as a reinforcer if another food is available, but it might do so if water was available instead of the second food. Here eating one food and eating the other are substitutable responses (either will replace the other as a reinforcer) but eating and drinking aren't.

We introduced reinforcers as kinds of stimuli, but we are now speaking of reinforcers in terms of responses. **A less probable response can be reinforced by the opportunity to engage in a more probable response.** This gives us a behavioral criterion for deciding whether some stimulus might function to reinforce some response and avoids the issues of circular definition that troubled the Law of Effect. The treatment shifted because we found that an important property of a reinforcer was the responding that it occasioned. Premack discussed this as follows:

> I have been accused of ignoring the stimulus; in fact, all the response probabilities we have measured have involved explicit stimulus operations, e.g., pellets, sucrose solutions, pinball machines, manipulanda, activity wheels, etc.... Even though it is desirable to take the position that a response is a means of calibrating the value a subject sets on a stimulus, and that the precise contribution which the response factors make to value is a matter for research, we should not swing so far in the other direction as to return to the traditional stimulus error. For example, tradition talks about the reinforcement value of food. But food has no value; we must talk about the value of food when it can be smelled (but not seen), seen (but not smelled), both seen and smelled, eaten, etc. It is part of the advantage of the response language to avoid this kind of error automatically. (Premack, 1971, pp. 125-127)

Like the reflex, reinforcement is a relation and not a theory or hypothesis. This relation includes responses, their consequences, and the behavior occasioned by those consequences.

THE ANALYSIS OF REINFORCEMENT

In discussing extinction in Chapter 5, we considered both rate of responding and resistance to change as possible measures of response probability or strength. Rate and probability of responding have long seemed closely linked. For example, if we calculate probabilities by looking at the proportion of time periods within which responses appear, we'd expect to find them there roughly twice as often if the rate of responding doubles. But rates and probabilities do not tell us how persistent or fragile responding will be. A high-rate response might drop off precipitously given a small change in reinforcement while a low-rate response might stubbornly continue despite a major change in reinforcement. The term strength as a label for the effects of reinforcement on responding had had a long and varied history (e.g., Killeen & Hall, 2001; Killeen *et al.*, 2002; Nevin & Grace, 2000, 2005; Skinner, 1938).

Strength: Resistance to Change Versus Rate of Responding

When we later treat shaping and schedules of reinforcement (Chapters 10 and 15), we'll see that a problem with response rate as a general measure is that it can be differentially reinforced. For example, we can make reinforcers contingent only upon slow responding or only upon rapid responding. If reinforcement contingencies can make responding either slow or fast, how can we use those rates to judge how strong the responding is? Resistance to change seems a more plausible measure because it is a property of responding that might vary independently of rate. Either slow or fast responses rates might be either durable or fragile given changes in the reinforcement contingencies.

Biology often provides useful analogies, and the issue of strength offers a case in point. If we were studying sensory processes, as in the visual system, we would find it useful to distinguish between effects of stimuli at the level of individual receptors as well as effects produced further along by interactions among those receptors. For example, in the visual system we could study how light acts on rods or cones, and how their reaction to light varies with duration or intensity of stimulation. These light and dark adaptations would depend on the depletion and recovery of photosensitive pigments during periods of light and dark. We can study these adaptations by examining single cells. But the firing of each photoreceptor will also inhibit and be inhibited by the firing of neighboring cells. These interactions will determine patterns of firings that can be studied only by examining the activity of cells in the context of the activity of the cells that surround them. (These interactions, by the way, will be relevant to some of the behavioral phenomena we'll consider in later chapters.)

If we think of the effects of reinforcers on responses as a variety of stimulation, then that stimulation might produce more or less resistance to extinction, but the particular rate of responding produced by the reinforcement contingencies will depend not only on the reinforcers that have been produced but also by the interactions of this response with other responses for which reinforcers might have been arranged, or, in other words, by the context in which the responding occurs. Thus, just as visual adaptations might depend on the direct stimulation of receptors whereas patterns of firing of visual receptor cells might depend on interactions among them, so also resistance to change might depend on direct effects of reinforcers whereas rates might depend on interactions among reinforced response classes. The language of response strength is most useful when the particular dimensions of responding that it invokes are adequately specified.

Activation and Coupling

We've already seen that we can break the effects of reinforcement down into components, such as effects on resistance to change and on response rate. Are there other ways to analyze what happens in reinforcement? A survey of the many quantitative treatments of reinforcement and behavior is beyond the scope of this chapter (e.g., Killeen, 1992; Killeen & Bizo, 1998; Nevin, 1992; Nevin & Grace, 2000), but some of them provide useful hints about how to proceed. For example, we've already seen how environmental changes can have multiple effects, as when extinction produces not only decreases in responding because a contingency has been broken but also emotional side effects because reinforcers are no longer delivered.

Reinforcers too can produce multiple effects. Some accounts suggest that, like extinction, reinforcement outcomes may also be the product of two different effects. In this case they would be effects not of the termination of reinforcers but rather of their delivery. One is a general activation or arousal effect, as produced for example when a rat that had not been eating is now able to eat again, and the other is a more specific effect of the response-reinforcer contingency, a coupling that determines the particular response class that comes to be strengthened by the reinforcer (e.g., Killeen, 1994; Killeen & Sitomer, 2003). To this point our analyses have mainly served for interpreting complex behavior in terms of basic processes, as when we showed how imprinting should be treated in terms of acquired reinforcers rather than in terms of elicited behavior (Chapter 4). This example shows that even the basic processes themselves can sometimes be subject to analysis.

THE ACQUISITION OF BEHAVIOR

Let's return now to learning, by examining how an organism may acquire responses through reinforcement. Record A of Figure 6-2 shows a hypothetical cumulative record of a rat's very first session of reinforced lever pressing. The first few responses are separated by long pauses. Within 5 min or so, the long pauses disappear and responding increases for the rest of the 15-min session. The acquisition of

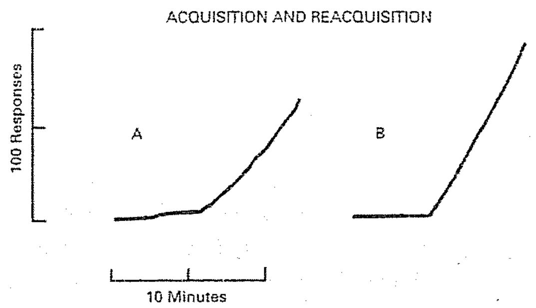

Figure 6–2 Hypothetical cumulative records of a rat's initial acquisition of lever pressing during the first session with each response reinforced with food (A), and reacquisition of lever pressing with reinforcement introduced again after a period of extinction (B).

lever pressing seems gradual. To repeat these observations, we extinguish lever presses until responding decreases to previous levels and then conduct another reinforcement session. Record B of Figure 6-2 shows what it might look like. Because of the prior extinction, there is no responding at first. When responding finally occurs and is reinforced, it immediately rises to a rate roughly equal to that at the very end of the first reinforcement session. This time, acquisition of lever pressing was abrupt rather than gradual. How do we reconcile these two very different performances?

Sidman (1960) has presented an alternative to the view that the initial learning produced an irreversible change.

> The animal learned not only those responses that succeeded in depressing the lever but also learned to go to the tray, to pick up the small pellet, to bring it to its mouth, etc. And these responses were learned in their correct sequence, because their reinforcement was correlated with the appropriate stimuli both from the environment and from the preceding behavior. The tray approach, for example, could be reinforced only after the sound from the food magazine; reaching for the pellet could be reinforced only after the pellet had dropped into the tray, etc.... What did we extinguish when we disconnected the feeding mechanism?... There is no magazine sound, no pellet sound, no visual pellet, no tactual pellet, etc. Tray-approach is still possible, but only in the absence of some of its controlling stimuli. The responses involved in picking up and ingesting the pellet can no longer occur in their originally learned context. While our extinction procedure may have returned the lever-pressing response to its preconditioning level, other components of the total learned sequence could not have undergone complete extinction. (Sidman, 1960, pp. 101–103)

From this analysis, Sidman summarized the reasons for the difference in acquisition in the two sessions: "When reinforcement was again introduced..., the animal did not have to relearn the whole sequence, because the whole sequence had not been extinguished" (Sidman, 1960, p. 103).

Reinforcement, then, doesn't produce learning; it produces behavior. In looking to see whether a rat presses a lever when a reinforcement contingency operates and not otherwise, we're concerned with the extent to which the rat has learned the consequences of its lever pressing. The consequences of responding are critical to learning not because learning follows from them but because **they are what is learned**. Contingencies involve the ways in which the environment is affected by behavior and are therefore important features of the environment for organisms to learn.

Sensory-Motor Learning

The consequences in the rat experiments so far have mainly included stimuli with biological significance for the organism's survival (e.g., food or water). But many presumably lesser consequences are important in day-to-day interactions with the environment. Response-reinforcer contingencies are all around us, but are easily overlooked (see Parsons, 1974, for a human example). We open a book to read it; we listen to hear what someone says; we reach toward a pencil to pick it up. Each consequence sets the occasion for new responses. When we finish reading one page, we turn it and read the next; when we've heard what has been said, we ask a question or make a comment; when we pick up the pencil, we write something with it. To the extent that each case involves behavior maintained by its consequences, each may be discussed in the vocabulary of reinforcement. Seeing reinforces looking; hearing reinforces listening; contacting the pencil reinforces reaching for it.

The interaction of sensory processes with behavior has been a longstanding source of controversy in the psychology of learning. Theorists took sides in debates over whether learning was sensory or motor. Did organisms learn responses or did they learn relations among stimuli? Did they learn response-stimulus or stimulus-stimulus associations? One issue is whether sensory processes should be treated as behavior. Doing so is consistent with the view that behavior must be dealt with in terms of relations among stimuli and responses rather than in terms of stimuli alone or responses alone. We can't measure seeing and hearing as unambiguously as we can measure a rat's lever presses or a pigeon's key pecks, but they are still kinds of behavior. They depend not only on whether visual or auditory stimuli are present but also on what the organism does. Looking makes seeing more likely and listening makes hearing more likely. The organism isn't passive in its environment.

Sounds and lights and other basic events have sometimes been described as neutral relative to potentially strong reinforcers or punishers (e.g., food and shock). But the label *neutral*, though convenient, is a misnomer. Events that are consequences of behavior are unlikely to be truly neutral; it is unlikely they'll have absolutely no effect on behavior. Nevertheless, reports of the reinforcing effects of stimuli such as lights and sounds were received with skepticism before the relativity of reinforcers was appreciated. Repeated demonstrations, however, led to the acceptance of *sensory reinforcement* (Kish, 1966). For example, if a rat's lever pressing in darkness briefly turns on a light, its pressing transiently increases. In other words, the light serves temporarily as a weak reinforcer. Phenomena like these were discussed in terms of *exploratory behavior* and *curiosity*, and experiments were extended to varied sensory consequences. For example, a monkey alone in an enclosure will operate a switch if that gives the monkey an opportunity to look at other monkeys outside (Butler, 1957).

In this research, responses chosen for ease of measurement were used to assess the effects of sensory consequences. But in any environment behavior inevitably has such consequences. The organism changes its environment simply by moving from one place to another and the things it sees and touches change as it moves; spatial relations among the components of its environment are a fundamental part of what it learns (Gallistel, 1990).

An experiment by Held and Hein (1963) illustrates these relations between behavior and sensory consequences. Pairs of kittens were raised in darkness; their first visual experience was in the apparatus shown in Figure 6-3. Both kittens received the same kind of visual stimulation: Each wore a shield that prevented it from seeing its own feet and body; each was prevented from seeing the other by the wide central post; and each saw the same pattern of black and white vertical stripes uniformly covering the walls of the circular enclosure. A sort of miniature carousel linked them, but one kitten (A) moved actively whereas the other (P) was moved passively. The active kitten stood on the floor; the passive kitten stood inside a carrier suspended above the floor. As the active kitten walked around the post, the passive one in its carrier moved a corresponding distance on the other side. If the active kitten turned around in place to walk in the other direction, the pulley system turned the passive one's carrier so that it too turned to face the new direction.

The kittens were exposed to similar visual stimuli, but those for the active kitten were consequences of its own behavior whereas those for the passive kitten were not; they depended on the active kitten's movements rather than its own. The kittens were then both given standard tests

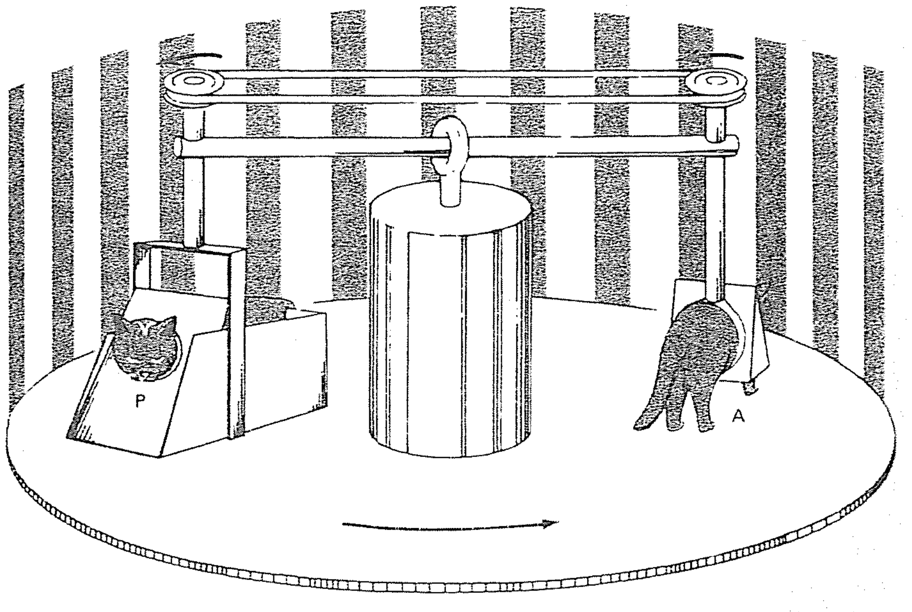

Figure 6–3 An apparatus for studying the relation between behavior and visual stimuli. Both the active kitten (A) and the passive one (P) are harnessed to a pulley system with its fulcrum at the central column. Kitten A stands on the floor and Kitten P rides in a carrier. As Kitten A moves about, the pulley system duplicates its changes in position for Kitten P (see arrows). (From Held & Hein, 1963, Figure 1)

of visual-motor coordination, such as visual paw placement: Normal kittens held in the air a short distance away from a horizontal surface such as a table edge extend their paws toward it. Even though their exposure to visual stimuli was equivalent, only the active kitten responded appropriately in these tests; the passive kitten became able to respond appropriately only later, after it was allowed to walk about freely in a lighted room.

This experiment has much in common with a classic one by Stratton (1897), who for eight days wore prisms that inverted and reversed his visual fields. At first his world looked upside-down and flipped over and his movements weren't coordinated with his surroundings. For example, in walking he looked down to see where he was going, but with the inverting prisms he was looking at the ceiling instead of the floor. Similarly, he had difficulty in pointing at or reaching for objects, because things once seen below eye level were now seen above and things on the right were now seen on the left, and vice versa. As time passed, however, his coordination improved and Stratton reported that the world no longer even looked upside-down.

The consequences of behavior are again crucial. Looking and moving in the visual field have different consequences with and without inverting prisms. Adjusting to the prisms requires learning the new consequences. For example, seeing the floor as one walks is important, but when one begins wearing inverting prisms, seeing the floor, once a consequence of looking down, becomes a consequence of looking up; in these situations, of course, up and down can be defined either relative to the visual field or relative to the body (cf. Harris, 1965). Thus, if seeing the floor while one is walking is a reinforcer and one walks while wearing inverting prisms, seeing the floor will reinforce the response of looking up instead of the response of looking down.

In discussing Stratton's inverted vision, we moved from locomotion to responses of smaller magnitude, such as eye movements. Even such

small-scale responses can have profound consequences. If you notice something out of the corner of your eye, you're more likely to see it clearly if you look toward it than if you look away from it (except in dim light, when you see an object best by looking a little bit away from it). Suppose that in a uniform visual field a contour such as the edge of an object can reinforce eye movement. We'd expect eye movements to become coordinated with events in the visual field. Data from the eye movements of newborns are consistent with this view. For example, with continued visual experience infants shown a simple figure such as a triangle in a uniform visual field tend to fixate more and more accurately on the edges and vertices of the triangle (e.g., Salapatek & Kessen, 1966).

We've much to learn about how arbitrary such relations between responses and consequences can be (cf. Hein *et al.*, 1979). For example, imagine setting up an optical system that projects visual stimuli in an infant's field of view and alters the natural consequences of eye movement. It presents stimuli only when the infant looks straight ahead. Now a stimulus appears in the right visual field but if the infant looks right, the stimulus disappears. If the infant looks left, however, the stimulus moves left to where the infant is now looking (and vice versa for stimuli appearing in the left visual field). In other words, this optical system creates a world in which the infant can fixate on an object only by looking away from it (cf. Schroeder & Holland, 1968).

The infant would probably learn how to look at things in this environment, but we'd want to think twice about doing it. Visual areas of the brain develop dramatically in infancy in humans as well as in cats (e.g. Blakemore & Cooper, 1970; Freeman, Mitchell, & Millidot, 1972). Some early changes may later be modifiable, but others may be relatively permanent. For example, a child with early problems of binocular vision may never acquire proper depth perception if the problems remain uncorrected until adulthood. Some things are more easily learned, whether in general or at some times rather than others, and some things learned are more easily changed than others. On the one hand, we may be reluctant to invoke reinforcement in accounts of behavior that is learned early and relatively permanently; on the other, we must take care not to rule it out when the permanence of the behavior might depend on the constancy throughout life of the contingencies that originally created it. Contingencies and consequences are part of the description of what we learn even with respect to our simplest interactions with events in the world. Phenomena such as sensory-motor learning demonstrate that reinforcement doesn't explain learning; rather, it is part of the description of what is learned. Organisms learn the consequences of their own behavior.

ADDENDUM 6A: SOME METHODOLOGICAL AND MEASUREMENT ISSUES

The cumulative record made it easy to see moment-to-moment changes in behavior, but it is not always feasible to arrange such recording in schools or treatment centers or other nonlaboratory environments. Sometimes behavior is recorded in checklists or charts or with counters or other devices. In such cases, observers must have well defined criteria for what counts as an instance of relevant behavior and what does not; without such criteria, data collected at one time by one observer may not be comparable to behavior collected at another time or by a different observer (consistency across time and observers is referred to as *reliability*). The measures must also be relevant to the behavior of concern. For example, the time a child spends with an open book may not be well correlated with the time the child spends actually reading (the relevance or appropriateness of measures is referred to as *validity*). Recording instances of behavior that can be counted is often more productive than recording time spent engaging in the behavior. For example, there is a big difference between completing five arithmetic problems in a minute and completing only one.

Table 6-1 Sample values on a linear scale vs. a logarithmic scale.

linear	10	20	30	40	50	60	70	80	90	100
logarithmic	2	4	8	16	32	64	128	256	512	1024

These issues are of special significance in education, but schools rarely track what students learn in a behaviorally significant way. Doing so is not just a matter of measuring with reliability and validity. As with cumulative records, important properties of the data must be made easily visible and the units recorded must be unambiguous (cf. Johnston & Pennypacker, 1993).

One method for recording the progress of learning is in *celeration charts* (Lindsley, 1992). In a standard celeration chart, the horizontal or x-axis (typically time or number of opportunities for behavior) is the usual linear scale, but the vertical or y-axis (the behavior of interest, usually expressed as a rate) is scaled logarithmically. A logarithmic scale is one in which a given interval is equivalent to a constant ratio rather than to a constant difference. For example, as Table 6-1 illustrates, on a linear scale the space between 5 and 10 equals that between 15 and 20 or 80 and 85 or 622 and 627 or any other pair where one equals the other plus 5; on a logarithmic scale, however, the space between 5 and 10 equals that between 15 and 30 or 80 and 160 or 622 and 1244 or any other pair where one equals twice the other.

The advantage of this scale for examining behavior that changes over time, such as the rate at which a student solves arithmetic problems, is that equal percentage changes over constant time periods produce lines of constant slope. Such charts have been useful in assessing the acquisition of academic skills and fluency (see Binder, 1996; Johnson & Layng, 1992). Another important component of such measures is determining minimal units of learning, sometimes called learn units (Greer & McDonough, 1999; Greer & Ross, 2004). A treatment of these and related educational innovations is beyond the scope of this section, but they promise advances in educational methods that are quantifiable as well as effective. They may also clarify the role of contingencies in the development of behavior over the life span (cf. Gewirtz & Pelaez-Nogueras, 1992).

ADDENDUM 6B: LATENT LEARNING

Questions about what is learned were implicit in a controversy based upon a phenomenon called *latent learning* (Thistlethwaite, 1951). Consider the experiment illustrated in Figure 6-4 (Tolman, 1948; Tolman & Honzik, 1930). Food-deprived rats in each of three groups negotiated a maze. In one group, the rats found food in the goalbox. Over successive daily trials their entries into blind alleys gradually decreased. In a second group, the rats found no food in the goalbox. Their entries into blind alleys decreased but remained much higher than those of the first group. In a third group, food was introduced in the goalbox only after 10 sessions. This group, starting like the second group (no food), quickly became like the first group (food). Rats that previously negotiated the maze without food in the goalbox now began to run with as few entries into blind alleys as did those that had always found food there. Until food was introduced, the learning of this group had been latent; what they had learned was demonstrated when food was introduced in the goalbox.

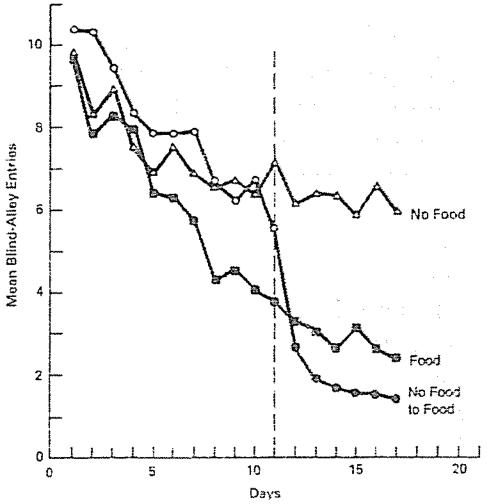

Figure 6-4 A latent-learning experiment. Rats were allowed one trial per day in a 14-choicepoint maze. One group (filled squares) always found food in the goalbox; a second (unfilled triangles) never did. A third group found no food in the goalbox through day 10 (unfilled circles) but then began to find food there (filled circles). This group, which until day 10 had performed like the second group, quickly became like the first one. They'd been learning the maze all along, so food in the goalbox was needed only to get them to demonstrate what they'd learned. (From Tolman & Honzik, 1930)

At first the argument was that because rats learned the maze equally well whether or not food was in the goalbox, learning could not to be attributed to food reinforcers. That argument was soon extended to reinforcers in general and latent-learning experiments were said to have shown that learning could occur without reinforcement. But then came the counterargument that food in the goalbox isn't the only possible reinforcer of a rat's maze-running. The rat's removal from the maze at the end of the run, its escape from the constrained spaces of blind alleys, or its return to its home cage where it is fed each might function as an effective reinforcer. So experiments varied the rat's handling at the end of the run, the width of the alleys it ran, and its home cage feeding. Each time an experiment showed that some reinforcer could generate maze learning, another one showed latent learning where that reinforcer couldn't have been effective. And so it went.

Even in principle the argument couldn't have been resolved, and latent learning gradually faded away as a critical theoretical issue. The reason was that a rat's negotiation of a maze inevitably involves the consequences of responding. At any choicepoint, one turn is followed by entry into a blind alley and the other by an opportunity to move farther through the maze; at the final choicepoint, only one turn is followed by entry into the goalbox, whether or not food is there. The rat's sniffing, touching, looking and moving through the maze are always consequential responses, even if they aren't as easily accessible to us as correct turns or entries into blind alleys. These consequences are what the rat learns. Whether to call them reinforcers is mostly a matter of preference. The language of latent learning, however, led to a blind alley.

Chapter 7
Consequences of Responding: Punishment

Punishment *and* reinforcement *have fairly straightforward histories.* Punishment *stems from the Latin* poena, *pain or penalty, with* impunity *and* subpoena *as relatives.* Reinforcement *stems from the Latin* fortis, *strong, which is related to such structural words as* fort, burg *and* belfry, *and perhaps, via the Indo-European root,* pri-, *to love, with* free *and* friend *as distant relatives.*

So far we've emphasized a relation, *reinforcement*, in which the consequences of responses make responding more likely. In another relation, *punishment*, the consequences of responses make responding less likely. Furthermore, a stimulus that reinforces responding when responses produce it may serve the opposite function when those responses remove it: Its removal may punish the responding. Inversely, a stimulus that punishes responding when responses produce it may reinforce responding when those responses remove it. For example, money may reinforce, as when a child is paid for completing a chore, but its removal may punish, as when the child's allowance is canceled because of a misdeed. Similarly, a painful burn may punish, as when I learn not to touch a pan just taken from a hot oven, but its removal or prevention may reinforce, as when I learn to treat a burn with appropriate medication or to put on kitchen gloves while handling things around a stove.

Except for positive reinforcement (reinforcement by presenting a stimulus: cf. Chapter 5), these relations are often grouped together as instances of *aversive control*. In other words, aversive control includes both *punishment* and *negative reinforcement* (reinforcement by removal or prevention of aversive stimuli). This chapter treats punishment; the next one turns to negative reinforcement, in escape and avoidance procedures.

Biological parallels are again relevant. In Chapter 3 we discussed the gradual disappearance of legs in the ancestors of whales as a variety of extinction; that disappearance is analogous to the disappearance of unreinforced responding in the extinction of operant behavior. Whether in phylogeny or ontogeny, use it or lose it applies: What becomes nonfunctional is no longer used and is lost (e.g., Carroll, 2006). But both realms also include the possibility of harmful functions, as when in phylogeny a mutation that causes a distinctive marking makes an organism more likely to be spotted by predators, or as when in ontogeny some response causes a painful burn or some other tissue damage. "Much thought and attention has been directed to the 'creative' dimension of natural selection and how new traits evolve, but this is only one aspect of the evolutionary process. Natural selection also acts to remove, in Darwin's words, 'injurious change' (Carroll, 2006, p. 72). Punishment in ontogenic selection is an analog of the effects of harmful variations in phylogenic selection.

KEY TERMS: Punishment; Recovery from Punishment; Relativity of Punishment; Positive and Negative Punishment; Eliciting Side Effects; Discriminative Side Effects.

THE VOCABULARY OF PUNISHMENT

As a procedure, *punishment* is arranging a response consequence that makes the responding less likely. The stimulus arranged as a consequence is called a *punisher*. For example, if a rat's lever presses produce electric shock and this procedure reduces lever pressing, the lever press is said to be punished and the shock is said to be a punisher. In these respects, the vocabulary of punishment parallels the vocabulary of reinforcement (cf. Table 5-1).

Like *reinforcement*, the term *punishment* has been applied to both procedures and outcomes. Thus, saying that a response was punished may mean either that the response produced a punisher or that responding decreased because it produced a punisher. As with reinforcement, our preferred usage will be to restrict the term *punishment* to the vocabulary of procedures or operations and to describe the outcome directly in terms of changes in responding. As with reinforcement, however, the outcome usage has so much precedent that we can't avoid it.

The vocabulary of punishment also parallels that of reinforcement in its object: **We say that responses, not organisms, are punished.** If a rat's lever pressing produces shock and lever pressing decreases, it is appropriate to say that the rat was shocked and that the lever press was punished; it goes against colloquial usage, but it is *not* appropriate to say that the rat was punished. As with reinforcement, this grammatical distinction encourages us to be precise when we observe and describe behavior. One reason it differs so dramatically from everyday usage is that the everyday concern is too often with retribution rather than with changing behavior.

Consider a child misbehaving. A parent calls the child and administers a spanking when the child comes. It may be convenient to say simply that the parent punished the child, but this usage makes it far too easy to omit the responses that might be affected. The immediate consequence of the child's misbehaving was that the parent called; the spanking occurred after the child obeyed the call. Although the child might misbehave less in the future because of the spanking, the spanking will probably also make the child less likely to come when called. By the way, this doesn't mean that a parent intent on spanking should go to the child instead; reinforcing responses incompatible with the misbehavior and other, better alternatives don't require spanking at all.

The point isn't just grammatical. We're more likely to see what is happening if we state the punished response explicitly (spanking punished the child's approach) than if we settle for a less precise description (spanking punished the child). In endorsing the grammar of reinforcing responses and punishing responses, we needn't prejudge how these operations affect behavior; we'll assume that they'll often affect other responses besides those for which they're arranged (e.g., a spanking may also elicit crying). A vocabulary that states the consequences of behavior helps us to describe such effects without ambiguity.

COMPARING REINFORCEMENT AND PUNISHMENT

The effect of punishment is simply the opposite of that of reinforcement. The relation between the two is illustrated in Figure 7-1, which presents hypothetical reinforcement and punishment data. The top graph shows changes in a rat's lever pressing during food reinforcement and then during extinction. We could as easily substitute scenarios involving the effects of reinforcing help with household chores or of punishing errors in school work; coming up with plausible everyday human alternatives is a useful exercise. In any case, during *baseline*, when lever pressing has no consequences, responding occurs infrequently. When *reinforcement* begins, responding increases over the first few sessions, after which it remains at a fairly stable level. In *extinction*, responding gradually returns to its former low baseline level.

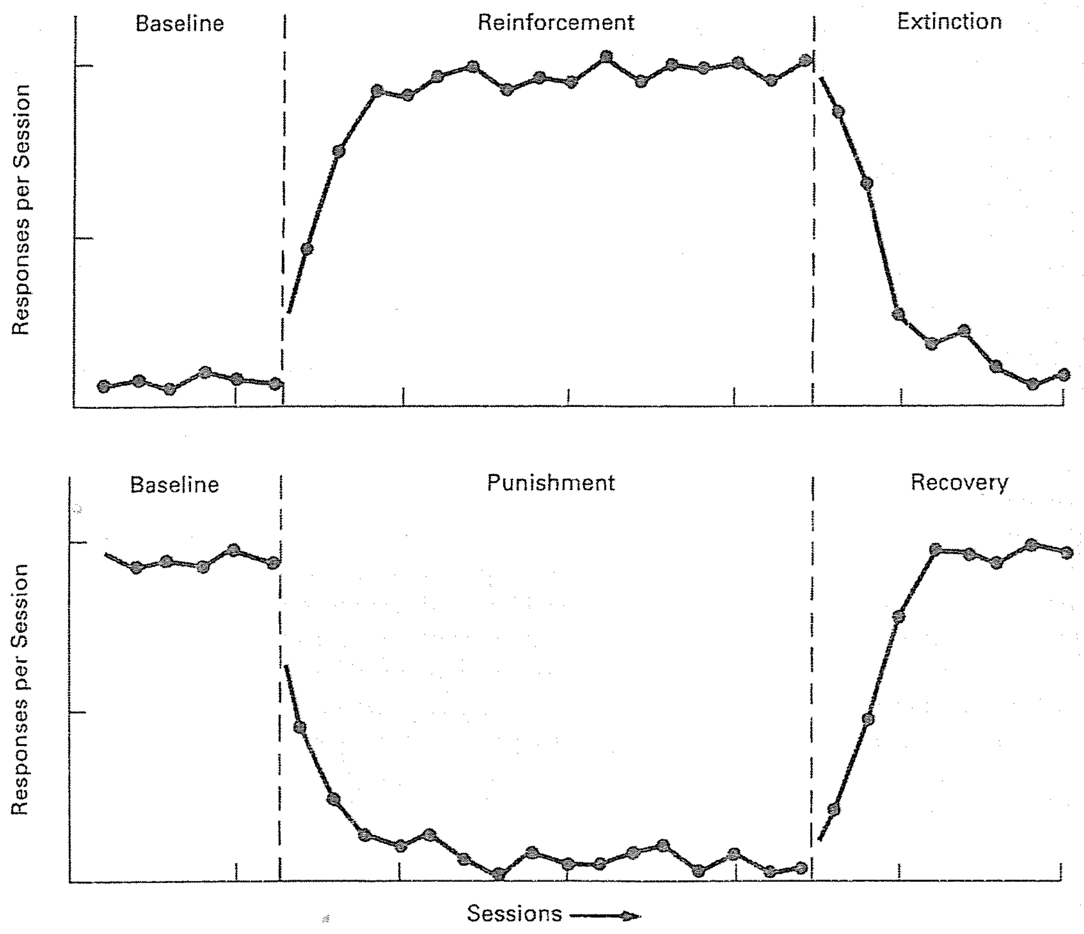

Figure 7-1 Effects of reinforcement and punishment on the hypothetical lever pressing of a food-deprived rat. Top graph: unreinforced pressing (baseline), the increase when it produces food (reinforcement), and the return to earlier levels when reinforcement ends (extinction). Bottom graph: pressing maintained by reinforcement (baseline), its decrease when response-produced shock is added (punishment), and the return to earlier levels when punishment ends (recovery). Pressing maintained by reinforcement is the baseline for illustrating the effects of punishment because a decrease in responding can't easily be seen unless responding is already frequent.

The bottom graph shows changes in the rat's lever pressing during and after punishment of lever presses with electric shock. Because punishment reduces responding, some responding must exist to begin with or we couldn't observe a decrease. In this example, responding is already maintained by food reinforcement, which continues throughout all sessions; the effects of punishment are then assessed by superimposing it on this baseline. The *baseline* shows the responding maintained before lever pressing is punished. When *punishment* begins, lever pressing decreases to a maintained low level. In *recovery*, punishment is discontinued and responding gradually returns to its former high baseline level.

Reinforcement and punishment are symmetrical: The former increases responding whereas the latter decreases it, but their effects continue as long as the procedures are maintained and disappear after they end (responding returns to earlier levels). Actual rather than hypothetical effects of punishment are shown in Figure 7-2 (Estes, 1944).

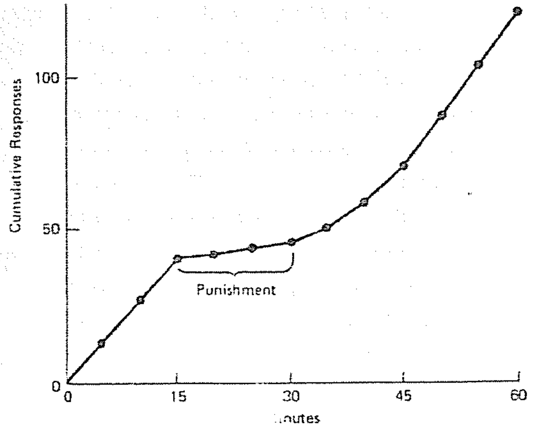

Figure 7-2 Cumulative record of the effect of punishment superimposed upon a baseline of food reinforcement that continued throughout the session. The response was a rat's lever press and the punisher was shock. The rate of pressing decreased during punishment and recovered after punishment was discontinued. (from Estes, 1944, Figure 10).

The effects in Figure 7-2 are clear. Nevertheless, the effectiveness of punishment has long been controversial. Punishment was incorporated into Thorndike's early versions of his Law of Effect (cf. Chapter 5). Thorndike argued then that behavior could be stamped out by annoyers as well as stamped in by satisfiers. Statements of Thorndike's law that included the punishment component were called the *strong* Law of Effect. Later, Thorndike withdrew the punishment component and the version that remained, which included only the stamping in of behavior, was called the *weak* Law of Effect. Thorndike based his conclusion on experiments on human verbal learning in which saying "right" to the learner enhanced responding whereas saying "wrong" had less effect than saying nothing. Thorndike accepted the finding as general evidence against the effectiveness of punishment.

Thorndike's conclusion had so much impact that even the data in Figure 7-2 were interpreted to mean that punishment was an ineffective procedure. The argument, based on recovery after punishment, was that punishment wasn't to be taken seriously as a way of changing behavior because it suppressed responding only temporarily. Yet by this criterion reinforcement should also be judged ineffective. Perhaps punishment was judged by different criteria than reinforcement because people were looking for a way to get rid of behavior that would work even if applied only once or only briefly. Yes, an extreme punisher can suppress behavior for an extremely long time, but as a rule punishment continues to work only as long as the punishment contingency continues. Once you've started punishment, you must remain ever alert for the reappearance of the punished behavior. We have reasons to conclude that techniques other than punishment should be used when possible, but if this is correct it is because Thorndike and his successors were right for the wrong reasons.

In more recent times, investigators turned again to punishment and studied conditions that modify its effectiveness in suppressing behavior (e.g., Azrin & Holz, 1966; Church, 1963). For example, experiments with electric shock as the punisher of a pigeon's food-reinforced key pecking showed that the more intense and immediate the punisher, the more effective it is. A punisher introduced at its maximum intensity suppresses responding more effectively than a punisher introduced at low intensity and then gradually raised to maximum intensity. Also, the effectiveness of the punisher may change over extended periods of punishment, as when a punisher of low intensity gradually becomes ineffective after many presentations (cf. Azrin & Holz, 1966, pp. 426-427). And, as with extinction, it is easier to reduce the likelihood of a response when some other response that produces the same reinforcer is available than when there are no other alternative responses that produce the same reinforcer.

Experiments on the properties of punishment not only changed the criteria for the effectiveness of punishment; they also raised questions about whether punishment had been judged adequately according to the old criteria. On occasion a single punisher or a few punishers, if intense enough and the consequence of a weakly maintained response, might make a response disappear for most or all of a lifetime. Such exceptional effects might as well be regarded as permanent, but why should we be surprised that recovery can sometimes be so slow that it exceeds an individual life span? Consider the analogous argument for reinforcement. Would we be surprised if winners of million-dollar lottery prizes continued to buy occasional lottery tickets throughout their lives even if that behavior was never again reinforced by a win? We wouldn't reject the existence of extinction on those grounds, so we shouldn't reject the existence of recovery after punishment because punishment occasionally has very durable effects.

Some incidental features of punishment probably contributed to its unusual treatment. A reduction in responding can be studied only if some responding already exists. A response that is never emitted can't be punished. That is why experiments on punishment usually superimpose punishment on reinforced responding. But the effects of punishment then depend on what maintains responding. For example, punishment by shock will probably reduce food-reinforced lever pressing less if a rat is severely food-deprived than if it is only mildly food-deprived.

Another difficulty is that punishing stimuli are likely to have other effects that occur whether or not they're produced by responses. As with reinforcement, the punishment effect must depend on the *relation between responses and punishers* (contingency) and not simply on the *delivery of punishers*. For example, shocks may reduce the rate at which a pigeon pecks a key even when they're delivered independently of key pecks, so before we speak of them as punishers we should be sure that they reduce responding more when they're produced by pecking than when they occur independently of pecking (cf. Azrin, 1956).

Prejudices against recognizing punishment were at times so strong that effective procedures were even given a different name, *passive avoidance*. For example, consider a rat on a platform above an electrified grid that is set up so the rat is shocked when it steps down to the grid; once shocked, the rat is less likely to step down again later. It is appropriate to say that stepping down is punished by shock, but some would say instead that this rat is passively avoiding the punisher, by holding back from stepping down. This usage allowed effective punishment procedures to be discussed in the vocabulary of passive avoidance while other procedures that didn't reduce responding were used to defend the claim that punishment was ineffective.

Punishment is a name for a relation between responding and consequences. The issue is mainly when it is appropriate to apply this name. Early analyses of punishment emphasized the impermanence of its effects. For this reason, punishment long went unacknowledged as a fundamental behavioral procedure. But the existence of consequences that reduce responding is no longer questioned. Except that the effects differ in sign, punishment parallels reinforcement: Reinforcement increases reinforced responding and punishment decreases punished responding. Both procedures have temporary effects; when they're discontinued, responding returns to earlier levels. Punishment can modify human behavior, so questions inevitably arise about the ethics of its application. But we cannot hope to resolve such questions without an adequate analysis of its properties.

THE RELATIVITY OF PUNISHMENT

In experiments on punishment, punishers are usually chosen for their reliable effect on a variety of responses, because such stimuli reveal the effects of punishment most clearly. One such punisher is electric shock, which can be measured accurately and can be presented at levels that are effective and yet do not damage tissue. Such stimuli, however, are only extreme instances of punishers. For example, as we've already seen, some developmentally disabled children chronically engage in head banging, hand biting and other self-injurious behavior. A brief squirt in the face with water mist from the sort of spray bottle that is used to moisten indoor plants is at worst a minor annoyance. When applied contingent on such behavior, however, it is an effective punisher (Dorsey *et al.*, 1980). As such it can be relatively innocuous, especially when compared with the serious damage these children can do to themselves (unfortunately, some who oppose any use of punishment find even this application unacceptable).

Punishment is inevitable because it is arranged by many natural contingencies. The child who teases a barking dog may get bitten and the child who plays with fire may get burned. Furthermore, even stimuli that ordinarily serve as reinforcers can become punishers. For example, food that is reinforcing at the beginning of a holiday feast may become aversive by the time the meal has ended. On the other hand, events that superficially seem aversive, such as falling from a height, may be reinforcing under certain circumstances (consider sky diving and bungee jumping). Like reinforcers, punishers can't be defined in absolute terms nor in terms of common physical properties. Rather, they must be assessed in terms of the relation between punished responses and the responses occasioned by the punishers.

The Premack principle of reinforcement (Chapter 6) stated that an opportunity to engage in more probable responses will reinforce less probable responses. That analysis has also been extended to punishment (Premack, 1971). Let's return to the apparatus that controls a rat's opportunities to run in a running wheel or drink from a drinking tube. It has been modified by a motor that can either lock the running wheel in position and prevent the rat from running or rotate it at a fixed speed and force the rat to run.

In this apparatus, depriving the rat of opportunities to run while giving it free access to water makes running more probable than drinking; depriving it of water while giving it opportunities to run makes drinking more probable than running (cf. the probabilities of running and drinking in the two parts of Figure 6-1). Rotation of the wheel can now be made a consequence of drinking: each time the rat drinks, the wheel begins to turn and the rat is forced to run. When running is more probable than drinking, this procedure increases drinking and it is appropriate to say that drinking is reinforced by running. But when running is less probable than drinking, this procedure has an opposite effect: now drinking decreases when running is its consequence and it is appropriate to say that drinking is punished by running. If their relative probabilities can be reversed, as illustrated in this example, any given response can be either reinforced or punished by any other response.

The stimuli and responses in typical experiments on reinforcement and punishment were chosen to make those procedures work (e.g., with food-deprived rats, eating is typically far more probable than lever pressing). They thereby obscured the potential reversibility of consequences as reinforcers and punishers. Responding can be raised or lowered by changing its consequences, and these outcomes are determined by the behavioral and not the physical properties of the consequences.

SIDE EFFECTS OF PUNISHMENT

As with reinforcers, punishers can have effects independent of their contingent relation to responses. If an organism is shocked or burned or pinched, some of its responses may have little to do with whether these events were brought on by the organism's own behavior. Difficulties arise in analyzing punishment because such effects must be distinguished from those that depend on the relation between responses and their consequences. As with reinforcement and extinction, once again we must deal with side effects.

Some effects of shock may be primarily physiological, as when successive shocks systematically reduce a rat's skin resistance. The effectiveness of later shocks may then vary with the rat's resistance. Other effects are primarily behavioral, as when the apparatus fails to prevent responses by which the rat can reduce its contact with the shock source. For example, fur is an insulator and rats have sometimes minimized punishing effects of shock by pressing the lever while lying on their furry backs (Azrin & Holz, 1966). In either case such side effects must be taken into account.

Eliciting Effects of Punishers

Figure 7-3 is from an experiment (Camp, Raymond, & Church, 1967) that compared effects of response-produced and response-independent shock. Lever pressing was maintained by food reinforcement in three groups of rats. Measured against a no-shock control group, shock reduced responding in both groups, but response-produced shock suppressed responding much more than response-independent shock (see also Church, 1963; Church, 1969). This difference makes it appropriate to call the response-produced shock a punisher. Events affect behavior most when behavior affects those events in turn (cf. Rachlin, 1967, p. 87). Just as we must distinguish between effects of reinforcer deliveries and effects of the response-reinforcer contingency, so also we must distinguish between effects of punisher deliveries and effects of the response-punisher contingency.

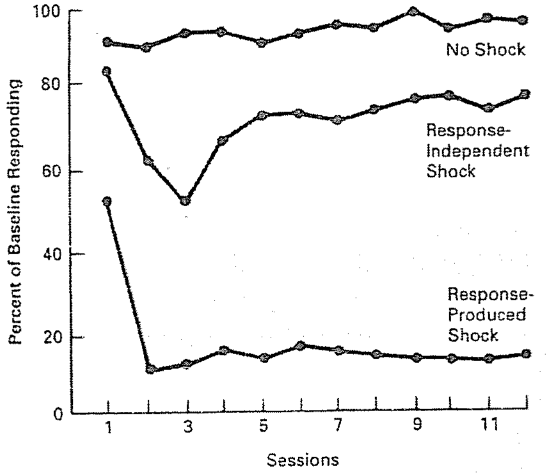

Figure 7-3 Effects of response-independent and response-produced shock on rats' lever pressing maintained by food reinforcement. Shock rate in the response-independent shock group was matched to that in the response-produced shock group, but response-produced shock reduced responding more than response-independent shock (from Camp et al., 1967, Figure 5).

The punishment of two classes of species-specific behavior in the Mongolian gerbil (Walters & Glazer, 1971) provides another example (see also Shettleworth, 1978). Sand digging, part of the gerbil's nesting behavior, consists of scooping sand and kicking it backwards; alert posturing, a defense reaction occasioned by sudden or aversive stimuli, consists of standing erect on the hindpaws with ears cocked. Delivering shock was difficult in the sandbox where the experiment was conducted, so a tone was established as an aversive stimulus by repeatedly pairing it with shock in a different setting. When contingent upon digging, the tone was an effective punisher. Digging decreased on producing it and recovered after the contingency ended; while digging decreased, alert posturing increased. But contingent on alert posturing, the tone wasn't effective. Alert posturing increased and didn't return to earlier levels for several sessions after the contingency ended; digging did not appreciably change. The eliciting effect of the tone on alert posturing was more powerful than its punishing effect. In punishment as in reinforcement, we must recognize the separate effects of response-stimulus contingencies and stimulus deliveries.

As another example, consider a monkey in a restraining chair with shock electrodes on its tail (e.g., Morse & Kelleher, 1977, pp. 193-198). At 5-min intervals, the monkey's lever press delivers shock to its own tail. A while after being placed in the chair, it begins to press. Eventually 5 min pass and its next press delivers a shock (this arrangement is called a 5-min fixed-interval schedule; see Chapter 15). The monkey briefly jumps and for a while stops pressing. But soon it starts again, pressing more and more rapidly until shocking itself once more at the end of the next 5-min interval. It repeats this performance throughout daily sessions. When shock is discontinued, lever pressing virtually ceases; when shock is reinstated, it returns. If shock level is raised, lever pressing increases; if shock level is lowered, it decreases. Shocks depend completely on the monkey's behavior; it would receive no shocks if it didn't press the lever. Why doesn't it just stop pressing?

The paradox is that the same shock that maintains lever pressing when produced by pressing according to a 5-min fixed-interval schedule suppresses lever pressing when produced instead by every press; the same shock can also be used to initiate and maintain escape and avoidance behavior (e.g., Barrett & Stanley, 1980). So how do we decide whether the language of punishment is appropriate? If every lever press produces shock, we call the shock a punisher because it reduces responding. But when a lever presses produces shock only at 5-min intervals, should we call the shock a reinforcer because it generates responding? It wouldn't help to yield to the temptation to call the monkey a masochist. Masochism is just a name we use when a stimulus that we think should be a punisher serves as a reinforcer; it doesn't explain anything.

Other experiments showed that even though a schedule in which lever presses produce shock every 2 min usually maintains higher rates of pressing than one in which they produce shock every 6 min, monkeys switch to the one with the longer time between shock when allowed to choose between them (Pitts & Malagodi, 1991). In other words, the monkeys preferred shocking themselves less often to shocking themselves more often. We might be reminded of the effectiveness of restraints as reinforcers with some children who engage in self-injurious behavior; they preferred the situation with restraints, in which they couldn't injure themselves, to situations without restraints in which they could (and did).

As long as human behavior includes such problems as self-injury, our concern with such phenomena is justified. We've seen that it is sometimes more appropriate to compare response-produced shock with response-independent shock rather than with no shock at all. Electric shock elicits manual responses such as lever pressing in monkeys. These eliciting effects of shock can be strong enough to override the punishing effects, so that lever pressing occurs in spite of and not because of the punishment contingency. If so, the case is analogous to that of the parent who tries to stop a child from crying by punishing the crying and has trouble because the punisher elicits the very response that the parent is trying to suppress. We may not like such contingencies, but if we ignore them that doesn't mean they'll go away.

Discriminative Effects of Punishers

Another side effect of punishment comes about because punishers can acquire discriminative properties, as when a response is reinforced only when it is also punished. An experiment designed to make response-produced shock signal the availability of food arranged two alternating conditions (Holz & Azrin, 1961). In one, a pigeon's key pecks had no consequences; in the other, every peck produced shock and some pecks also produced food reinforcers. A low rate of pecking was maintained when pecks produced no shock, because then they never produced food either; but pecking increased once pecks began to produce shock, because only then did they sometimes also produce food. Sample records with food reinforcers discontinued are shown for two pigeons in Figure 7-4. With no shocks, the rate of pecking was low. When responses began to produce shock, the rate increased. When shocks stopped, a brief rate increase (arrows) was followed by a decrease to the earlier low levels.

Again we can ask whether the shocks should be called punishers. In fact, from Figure 7-4 we should conclude that shock was a reinforcer. The main difference between the shock and other more familiar reinforcers is that here the shock acquired its power to reinforce through its relation to food; it will lose its power if that relation ends. Perhaps these procedures are relevant to human behavior. For example, a battered child might provoke a parent to the point of a beating because the beatings are often followed by more attention from the then-remorseful parent than ever follows less traumatic parent-child interactions. In this example, the beating is analogous to the shock in Figure 7-4 and the parent's attention is analogous to the food. A parent's attention can be a potent reinforcer and may sometimes override the effects of consequences that would otherwise serve as punishers. Thus, a behavioral analysis may be relevant to human problems such as child abuse. (By the way, I hope you noticed that this is still another human example attached to the discussion of a pigeon experiment.)

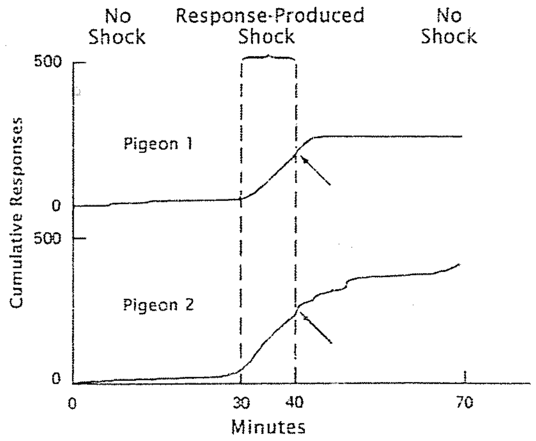

Figure 7-4 Discriminative effects of shock. First, two conditions alternated: Pigeons' pecks produced neither food nor shock, or pecks always produced shock and occasionally produced food. In the sessions shown here, no food was presented. For both pigeons, pecking increased when pecks began to produce shock but decreased to earlier levels when response-produced shock ended (from Holz & Azrin, 1961, Figure 3).

TIMEOUT AS PUNISHMENT

Suppose two levers are available to a monkey and presses on one of them produce food whenever a light is on. We can expect some presses on the other lever, but we can punish them by making each one produce a time period during which the light turns off and presses on the first lever won't produce food. Such periods are called *timeout* or, more fully, *punishment by timeout from positive reinforcement* (Ferster, 1958).

Timeout originated in experiments with pigeons and rats and monkeys and was common in the laboratories in which I worked as a graduate student and early in my career, but those origins are unfamiliar and it is no doubt now best known in its human applications (e.g., Wolf, Risley, & Mees, 1964). For example, a period of time in an isolation room has sometimes been used to punish the problem behavior of institutionalized children. But as we've seen in other examples, judgments about whether particular events will be reinforcers or punishers can be difficult. When timeout in isolation was used in an attempt to punish the tantrums of a 6-year-old girl on the autism spectrum, her tantrums increased substantially instead of decreasing. This child often engaged in self-stimulation (e.g., creating visual flicker by waving her fingers closely in front of her eyes), but this behavior was often interrupted by the staff. For her, time in the isolation room was a reinforcer because there she could engage in the self-stimulation without interruption (Solnick, Rincover, & Peterson, 1977).

As with any form of punishment, the main function of timeout is to reduce behavior, but it is too often applied without attention to alternative behavior that might be reinforced (Winett & Winkler, 1972). Some who use timeout without understanding its behavioral basis apply it counterproductively. For example, some years ago I conducted research with preschoolers in a local daycare center and occasionally had opportunities to observe how the staff managed the behavior of the children in their care. Timeout was one of their procedures. By then, isolated timeout rooms were already a thing of the past; instead the staff used designated timeout areas. A child who disturbed other children in a playgroup might be sent to sit for five minutes on a timeout carpet near the play area, in view of the other children. I watched as one boy in timeout sat quietly for about four minutes and then began to cry and act up (four minutes of sitting quietly is a very long time for a preschooler). The staff person nevertheless rigidly adhered to the five minutes of timeout and let him return to the playgroup at the very time when his disruptive behavior was most dramatically increasing. It would have been better to "catch the child being good" and to have allowed his return at the end of three or four minutes. Unfortunately I could not intervene: a condition of my research arrangements was that I not intrude upon staff practices. But it is good to know that timeout procedures have been falling out of favor and are gradually being replaced by other practices based on reinforcement rather than punishment.

THE ETHICS OF PUNISHMENT

Every day animal shelters put down unadopted dogs and cats and other pets that exceed their capacities. Sometimes it is behavior that makes the difference between life and death. If an owner delivers to a shelter a dog that barks excessively or that damages furniture, the dog is far less likely to be placed with a new owner than one that does not present these problems (e.g., Wilkes, 1994, 2009). Is it unethical to use punishment if that might eliminate such behavior so that instead of being put down the dog can successfully live in a human home?

The behavioral properties of aversive control have implications that are consistent with ethical arguments against aversive control. For example, a parent who arranges frequent aversive consequences for a child's behavior may acquire aversive properties. To the extent that the child then learns to escape from or avoid the parent's company, contingencies other than those available to the parent will increasingly influence the child's behavior. But if punishment seems the only effective technique for reducing the self-mutilating behavior of a brain-damaged child, punishment might be a lesser evil than the permanent damage the child might self-inflict. Self-injurious children have been known to disfigure and even blind themselves. Ethical precepts are concerned with the acceptable and unacceptable outcomes of our actions, which implies that the consequences of behavior cannot be ignored.

One of the ethical standards for interventions in applied behavior analysis is to impose the least restrictive treatment. This means that a treatment that empowers a child by shaping the child's communication skills will always be preferable to one that calls for restraints or other environmental restrictions. Assessments and applications of reinforcement and extinction and punishment contingencies have become increasingly sophisticated (e.g., Fisher & Bouxsein, 2011; e.g., Lerman & Toole, 2011; Vollmer & Athens, 2011). We may look forward to the day when this level of understanding of basic contingencies is shared not only among behavior analysts but also among parents and other caregivers.

In a discussion of the status of our culture and its progress in finding alternatives to aversive control, B. F. Skinner came to the following conclusion:

> Even in politics and government the power to punish has been supplemented by a more positive support of the behavior which conforms to the interests of the governing agency. But we are still a long way from exploiting the alternatives, and we are not likely to make any real advance so long as our information about punishment and the

> alternatives to punishment remains at the level of casual observation. (Skinner, 1953, pp. 192-193.)

Skinner included techniques of reinforcement among his "alternatives to punishment." Unfortunately, if we look at the behavior of those who deliver reinforcers or punishers, we see that effects of reinforcement often show up only long after the reinforcer is delivered (the effect of reinforcing a child's diligent schoolwork each day might not be obvious until several days have gone by), whereas the effects of punishment often show up right away (one effect of slapping or scolding a child who is teasing a sibling is that the teasing stops immediately). Thus, delivering a punisher is much more likely to produce immediate consequences than delivering a reinforcer. That means that people will more readily learn techniques of aversive control than techniques of reinforcement, but it doesn't mean that the aversive techniques are better.

Some argue against any kind of modification of behavior, whether involving aversive stimuli or positive reinforcers. But those who make such arguments should recognize that our behavior is being modified all the time, both by natural contingencies and by the artificial ones created by those around us. Politicians and businesses and educators and religious groups and governments arrange all sorts of contingencies that affect our behavior in many different ways. Denying them won't make those contingencies go away. Ignorance rarely provides good grounds for ethical judgments, so a counterargument is that our best defense against the misuse of behavioral techniques is to learn as much as we can about how they work.

ADDENDUM 7A: PUNISHMENT AS PRIMARY RATHER THAN DERIVATIVE

As discussed earlier, the research by Azrin and colleagues that treated the direct effects of punishers (e.g., Azrin, 1956; Azrin & Holz, 1966) was conducted in a context in which punishment was typically regarded either as ineffective or, when effective, as the indirect product of other processes. For example, Skinner (1953) argued that the behavior leading up to a punished response generated aversive proprioceptive stimuli that ended when the individual turned to doing something else. According to this interpretation, the punished response only indirectly became less likely, because switching to something else avoided those aversive proprioceptive cues. Punishment was then said to happen because everything else was reinforced by avoidance of the aversive stimulus, i.e., the punisher.

The argument that punishment is reducible to avoidance is curious, because punishment is a robust phenomenon in which the abrupt introduction of a strong punisher typically reduces responding very quickly, whereas it is usually difficult and time-consuming to establish avoidance, though it is robust once in place (we shall treat avoidance in more detail in the next chapter). Why invoke hypothetical behavior presumed to arise from relations that are difficult to establish to explain behavior that is easy to establish?

There are additional grounds for arguing that punishment is not a derivative process. For example, as we've already seen, not only are less probable responses reinforced when they occasion more probable responses, but also more probable responses are punished when they occasion less probable responses. These relations are reversible, in the sense that reversing the probabilities reverses the reinforcement and punishment effects. When eating is more probable than running, stimuli that occasion eating will reinforce running whereas stimuli that occasion running will punish eating, but when eating is less probable than running, stimuli that occasion running will reinforce eating whereas stimuli that occasion eating will punish running. The differences between reinforcement and punishment lie along a continuum in which some consequences raise probability of behavior and others lower it. Consequences can move from one side of the continuum to the other, as when food is sometimes reinforcing and sometimes aversive. To invoke some indirect process for just one side of that continuum is far less parsimonious than to assume that the continuum has symmetry.

Evolutionary considerations add to the plausibility of the argument for the direct effects of punishers. If the phylogenic origins of reinforcement are rooted in the effects of consequences on response probabilities it would be strange indeed if operant behavior evolved so that those probabilities were limited to only one direction, so that consequences could raise them but could never reduce them except by indirectly raising the probabilities of incompatible responses. But even if punishment did work that way, so that we learned not to stick our hands into fires only because we avoided the aversive proprioceptive stimuli occasioned by approaching fires, it would make little practical difference. For those who have to make decisions about whether or when to use punishers, punishment contingencies will work pretty much the same way whichever theoretical position one assumes, but nevertheless we will here treat punishment as a behavioral process in its own right and not as a derivative of something else.

Chapter 8
Consequences of Responding: Escape and Avoidance

Escape, *as* ex-, *out of, plus* cappa, *cape, seems to be derived from the Old North French* escaper, *to take off one's cloak or, by extension, to free oneself from restraint.* Avoidance *shares its sense of getting out of, as in making empty, with several of its relatives:* vacant, evacuate, vanish, waste. Aversive *is derived from the Latin* a, *away, plus* vertere, *to turn.* Vertere *has a Germanic relative, the suffix* -ward *or* -wards, *which has had senses of warding off, guarding or regarding; the suffix appears in* reward, *a word often inaccurately treated as a substitute for* reinforcement.

Organisms get rid of as well as produce stimuli. For example, a rat doesn't ordinarily expose itself to shock, and if shock does occur the rat escapes from it, given the opportunity. If presenting an aversive stimulus punishes a response, removing or preventing that stimulus may reinforce the response. When a response terminates or prevents an aversive stimulus and becomes more probable for that reason, the stimulus is called a *negative reinforcer* and the procedure is called *negative reinforcement*. Positive and negative reinforcement are distinguished by whether a response produces or removes a stimulus.

Later we'll encounter some problems in the vocabulary of positive and negative reinforcement. But the vocabulary has substantial precedent. It is standard usage that *positive* and *negative*, as modifiers of the term *reinforcement*, refer to the consequence produced by responding (whether the response adds something to the environment or takes something away), and that *negative reinforcer* refers to the stimulus itself and not to its removal (if removal of shock reinforces a rat's lever press, shock, not the shock-free period that follows the response, is the negative reinforcer). This vocabulary was established gradually. At one time, negative reinforcement was defined as above in some textbooks but as equivalent to punishment in others; such misunderstandings still appear occasionally. As indicated by the etymologies introducing each chapter of this text, language evolves. Thus, the evolving language of reinforcement may eventually make the distinction between positive and negative reinforcement of marginal value (cf. Michael, 1975). Nevertheless, when we do invoke the vocabulary of positive and negative reinforcement and punishment, we'll adhere to standard usage, summarized as follows:

KEY TERMS: Aversive Control; Escape and Avoidance; Deletion versus Postponement; Negative and Positive Reinforcement; Negative Reinforcement and Extinction; Species-Specific Defense Reactions.

- *Reinforcement* makes the reinforced response *increase*.
- *Punishment* makes the punished response *decrease*.
- The modifier *positive* means that the consequence of responding is the *addition* of a stimulus to the organism's environment.
- The modifier *negative* means that the consequence of responding is the *subtraction* of a stimulus from the organism's environment.
- Given that it is possible to make an unambiguous distinction, *reinforcers* and *punishers* are stimuli and not the absence of stimuli.

The last item above is parenthetically qualified because the distinction is sometimes difficult. For example, is it more appropriate to think of a traffic ticket in terms of the piece of paper presented to you or your loss of money when you pay the fine? Similarly, is it more appropriate to think of an exam in terms of getting a high grade or avoiding a poor one?

ESCAPE

Escape procedures are the simplest examples of negative reinforcement: An organism's response *terminates* an aversive stimulus. They differ from avoidance procedures, in which the response prevents or delays an aversive stimulus. This vocabulary is consistent with everyday usage: We *escape* from aversive circumstances that already exist, but we *avoid* potential aversive circumstances that haven't yet happened. For example, you might leave a party to escape from the company already there or to avoid someone expected to arrive later.

In institutional settings for developmentally disabled children, the children sometimes behave aggressively because in that way they escape from simple demands placed upon them, such as working on tasks designed to teach them how to fasten and unfasten clothing buttons . For two such children, aggression dropped to near-zero levels when they could escape from demand situations by engaging in other behavior that was incompatible with aggression (Carr, Newsom, & Binkoff, 1980). Such cases of escape might of course imply that typical demand situations in such settings don't provide enough reinforcers. The more important point, however, is that some problem behavior is maintained not because it produces attention or other positive social consequences but rather because it is a way to escape from or avoid some circumstance (cf. Fisher & Bouxsein, 2011; Vollmer & Athens, 2011).

Conditions for escape can be arranged for a rat by constructing a compartment with an electrified grid floor. Movement from one place to another as the escape response is illustrated in Figure 8-1, from a shock-escape experiment with rats in a runway (Fowler & Trapold, 1962). Running speed was fastest when the shock was turned off as soon as the rat reached the end of the runway. The longer the delay between reaching the end of the runway and the shock turning off, the slower the rat ran. This is one of many examples of quantitative effects of reinforcement. For example, with both positive and negative reinforcement, immediate reinforcement is more effective than delayed reinforcement (*delay* parameter) and large reinforcers are more effective than small ones (*magnitude* or *intensity* parameter; see *parameter* in the glossary).

Movement from one place to another makes an effective escape response, but discrete responses such as lever presses are easier to record. A rat's

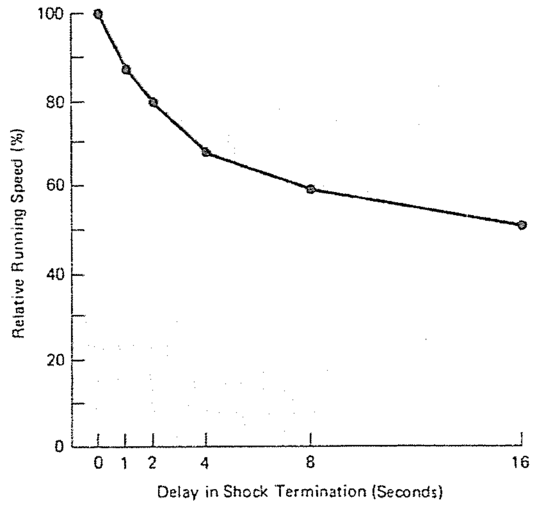

Figure 8-1 Relative running speed as a function of delay of shock termination in the escape responding of rats in a runway. The longer the delay between reaching the end of the runway and shock termination, the slower the rats ran. (from Fowler & Trapold, 1962, Figure 1).

lever press in the presence of shock can turn off the shock, or in the presence of bright light it can turn off the light (e.g., Keller, 1941). If in either case responding increases because of its consequences, we say the response has been reinforced. The parallel between positive and negative reinforcement is straightforward. In the absence of food, responding that produces food increases; in the presence of shock, responding that removes shock increases. Yet escape procedures receive less attention than more complex procedures. Research on negative reinforcement is dominated by avoidance, in which aversive stimuli are prevented or delayed by responses that occur in their absence.

Elicited Responding and Escape

One reason for the relative neglect of research on escape is that it is usually easy to raise the probability of a rat's lever presses or a pigeon's key pecks using positive reinforcement, but it is often difficult to do so using negative reinforcement in escape procedures (e.g., Hoffman & Fleshler, 1959). This is at least in part because the temporal relation between reinforced responses and responses produced by the reinforcer in positive reinforcement differs from that in negative reinforcement. The two conditions are diagrammed in Figure 8-2.

In positive reinforcement, the reinforcer is absent when the reinforced response is emitted. After the response, the reinforcer is presented and occasions other responses. For example, if a rat's lever press is the reinforced response and food is the reinforcer, food is absent while the rat presses; eating doesn't occur until food becomes available after the press. Lever pressing and eating occur at different times and so don't compete directly with each other.

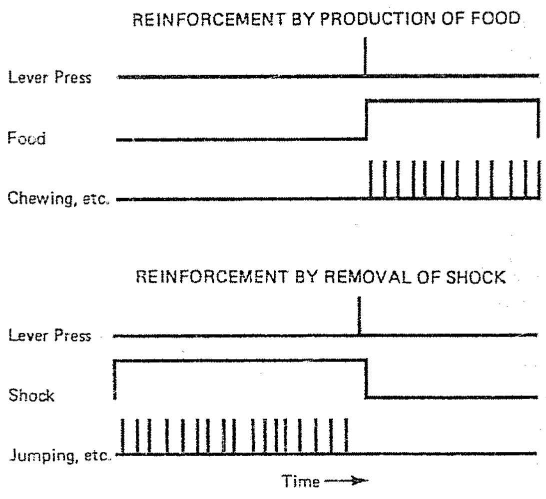

Figure 8-2 Temporal relations between reinforced responses and other responses produced by the reinforcer in positive reinforcement (top) and in negative reinforcement (bottom). In food reinforcement, the reinforced lever press has already occurred when food delivery produces behavior (e.g., handling the food), so those responses don't compete much with pressing. In shock escape, responses produced by shock (e.g., jumping) occur at the same time that reinforced lever pressing must occur, so those responses do compete with pressing.

In negative reinforcement, however, the negative reinforcer is present before the reinforced response is emitted; only after the response is it removed. For example, if a rat's lever press is the reinforced response and shock is the negative reinforcer, shock is present before lever presses occur. The shock generates responses such as elicited jumping or behavior that reduces contact with the shock source. As long as shock is producing these responses, they can compete with lever pressing. Those responses disappear only after a lever press turns off the shock.

Similarly, if the negative reinforcer is bright light from which the rat can escape by pressing a lever, the rat may instead close its eyes and hide its head in a corner. Any movement from that position is punished by greater exposure to the light, so the rat isn't likely to come out of the corner and press the lever. Getting a rat to escape from light by lever pressing requires procedures that reduce the likelihood of such competing responses (Keller, 1941).

The Ambiguous Distinction between Positive and Negative Reinforcement

The comparison in Figure 8-2 suggests that **whether stimuli are presented or removed may be a less important criterion for distinguishing between positive and negative reinforcement than whether responses generated by the reinforcer occur at times when they can compete with the reinforced response.** An experiment on escape from cold illustrates the point (Weiss & Laties, 1961). A rat was placed in a cold chamber in which its lever presses were reinforced by the operation of a heat lamp. On the one hand, this procedure could be called positive reinforcement: it involves adding energy in the form of heat when a lever press occurs. On the other hand, cold may function as a stimulus through its effects on temperature receptors in the rat's skin. Cold is the absence of heat, but it is also a significant and potentially aversive environmental event. By this interpretation the procedure should be called negative reinforcement, because turning on the heat lamp after a response terminates the effects of cold.

In escape from cold, a case can as easily be made that reinforcement involves presenting a stimulus as that it involves removing a stimulus. It is easy to find ambiguities in other instances of reinforcement. For example, we can argue that water reinforcers terminate aversive stimuli generated by a dry mouth, or that food reinforcers terminate aversive stimuli generated by the depletion of nutrients in the bloodstream (cf. Hull, 1943). Why then did the distinction between positive and negative reinforcement take on behavioral significance?

Let's reconsider the rat in the cold. Before the reinforced lever press occurred, the rat huddled in a corner and shivered. Those responses reduced the likelihood of lever presses. Once a lever press occurred, the heat lamp turned on and those competing responses became less likely, but a rat that is no longer cold can't escape from cold. Responses that competed with the reinforced response occurred before rather than after reinforcement, so this example is more like escape from shock than like production of food or water (cf. Figure 8-2). Thus, it is probably more appropriate to call this negative instead of positive reinforcement.

We cannot completely eliminate the ambiguity. Reinforcement always involves a change in an organism's situation, so it inevitably leads to differences in responding before and after the change. At best, we can regard such changes as producing a continuum of effects, ranging from those in which other responses are highly likely to precede and compete with reinforced responses to those in which they don't do so or perhaps even raise the likelihood of reinforced responses.

One other dimension of behavior may enter into the positive-negative distinction. In workplaces, for example, managers may get employees to complete tasks either by threatening consequences for noncompletion or (probably more rarely) by giving praise or recognition when the work is done. Aside from removal of a threat versus presentation of praise and even if both procedures were equally effective in motivating employees, it is reasonable to guess that workers given a choice would prefer the praiseful workplace over the threatening one. Given two reinforcement contingencies, the one that is closer to the positive end of the positive-negative continuum will probably be preferred to the other.

Positive and Negative Punishment

The distinction between positive and negative reinforcement is easily extended to positive and negative punishment (though here, too, ambiguous cases are possible). Responses can be punished by some events, such as shock or forced running in a running wheel. Responses also can be punished by the termination of events. For example, removing food contingent on a food-deprived rat's lever presses is likely to reduce the rate of pressing. The problem is that if this is negative punishment it might be hard to demonstrate. If the rat is food deprived and food is available, it will probably eat rather than press, so we won't have many opportunities to punish lever pressing by removing the food (consider Figure 8-2 with food and shock reversed). For this reason, studies of negative punishment haven't usually removed the positive reinforcer itself; paralleling the emphasis on avoidance rather than escape in studies of negative reinforcement, they have instead removed a stimulus in the presence of which responses are reinforced.

For example, suppose two levers are available to a monkey and that presses on one lever produce food whenever a light is on. We can expect some presses on the other lever, but we can punish them by making each one produce a time period during which the light turns off and presses on the first lever won't produce food. As we have seen, such periods are called *timeout* or, more fully, *punishment by timeout from positive reinforcement* (Ferster, 1958).

Identifying Contingencies: Procedural and Behavioral Criteria

The distinction between positive and negative reinforcement has a long history. There was a time when it was debated whether the terminology of positive and negative should apply to the procedures of presenting or removing stimuli or to the effects of these procedures in reinforcement and punishment. Table 8-1 summarizes the distinction as it is currently observed with regard to procedures. The first word (positive versus negative) distinguishes between procedures based on presentation versus removal; the second word (reinforcement versus punishment) distinguishes between outcomes based on whether responding goes up or goes down.

As discussed in connection with Figure 8-2, however, behavioral criteria may be more relevant than procedural ones. For example, when the lever presses of a rat in the cold turn on a heat lamp, the contingencies can be entered in different cells depending on whether we describe the procedure in terms of presenting or removing heat or presenting or removing cold. But what is more important is that cold produces shivering and huddling, which is likely to compete with lever pressing. Procedures that produce behavior not likely to compete with the response to be reinforced are usually categorized as positive reinforcement. Those that produce behavior that competes with the response to be reinforced are usually categorized as negative reinforcement.

The effectiveness of punishers can also be organized in terms of relations between the behavior that is punished and the behavior that is produced by the punisher. For example, when food is available to a food-deprived pigeon and a contingency arranges that key pecks remove the food (negative punishment), eating will make it unlikely that many punishable key pecks will be emitted.

Table 8-1 Procedural criteria for distinguishing among positive and negative reinforcement and punishment

	Responding Increases	*Responding Decreases*
Response Produces a Stimulus	Positive Reinforcement	Positive Punishment
Response Removes or Prevents a Stimulus	Negative Reinforcement	Negative Punishment

Touching the Hot Stove: Natural Aversive Contingencies

In themselves, reinforcement contingencies are neither good nor bad. Reinforcers can be used by parents to shape appropriate social behavior or by teachers to shape academic skills, but they can also be used by gambling institutions or drug dealers or pimps to shape other less functional kinds of behavior. The same is true of aversive control. Without punishment, a child who had been burned upon touching a hot stove or bitten upon approaching an unfamiliar barking dog would remain undeterred from doing so later on. Without negative reinforcement, a child would not learn to escape from the rain or the cold by coming indoors or to avoid those others who often cheat at games or bully or lie. On the other hand, many governments punish varieties of speech or consumption or unusual behavior, and many engage in forms of coercion that are most appropriately interpreted as instances of negative reinforcement. If that is how much of our world works, we should at least be clear about what happens.

The nonhuman examples remain relevant. The relativity of reinforcers should remind us that any given state of affairs may be reinforcing or aversive depending on its context. Suppose a tone is correlated with shock deliveries to a rat but nothing happens when a buzzer is sounding. If chain pulls turn off the tone and turn on the buzzer, the buzzer onset will reinforce chain pulls. By pulling the chain, the rat escapes from the tone and its accompanying shock deliveries.

Suppose, however, that the tone is correlated with food deliveries but again nothing happens when a buzzer is sounding. If chain pulls turn off the tone and turn on the buzzer, the buzzer onset will punish chain pulls. By pulling the chain, the rat produces a timeout from the tone and its accompanying food deliveries. In other words, the buzzer can serve as reinforcer or as punisher depending on its context. Similarly, as gauged by absenteeism a school environment may be reinforcing or punishing depending on the conditions that prevail at home. For one child, going to school is reinforced because it is a convenient way to escape from an abusive parent; for another, going to school is punished because it means having to deal with school bullies or arbitrary teachers.

In later chapters we'll see that organisms differentially attend to stimuli depending on whether the stimuli are correlated with relatively reinforcing or relatively aversive conditions. Meanwhile it is important to recognize that, as shown by these examples and as pointed out by Perone (2003), aversive control is ubiquitous and indeed unavoidable.

AVOIDANCE

In avoidance, the aversive stimulus isn't present when the reinforced response occurs. The two major varieties of avoidance procedures are called *deletion* and *postponement*. Deletion procedures are analogous to swatting a mosquito before it gets to you: Once you've swatted it, you've permanently prevented that particular mosquito from biting you. Postponement procedures are analogous to putting coins in a parking meter: You postpone the violation flag as long as you put in coins and reset the meter, but once you stop putting in coins, the meter eventually runs out.

As an example of a deletion procedure, imagine a rat in a chamber with a lever and a floor grid through which brief shocks may be delivered. Shocks are delivered once a minute, but if the rat presses the lever before the next one is due, that shock is omitted. In this procedure, the rat can avoid shock completely by pressing at least once a minute. Deletion procedures are sometimes conducted in discrete trials. For example, a light comes on for 30 s. If the rat presses the lever during the 30 s, no shock is delivered at the end of the trial; it is delivered only if the rat fails to press.

Now consider an example of a postponement procedure (Sidman, 1953). Two clocks control shock deliveries. Which clock runs depends on whether the last event was a shock or a lever press.

The first clock times the *shock-shock* or *SS interval*, the time between shocks if the rat doesn't press. Whenever a shock is delivered, this clock resets to zero and starts timing a new SS interval. Whenever the rat does press, however, control shifts to the second clock. This clock times the *response-shock* or *RS interval*, the time by which each lever press postpones the next possible shock. While this clock runs, each press resets it to zero and starts a new RS interval. In these circumstances, the rat can postpone the shock indefinitely by always pressing the lever before the current RS interval ends. If no response occurs and a shock is delivered at the end of an RS interval, the SS-interval clock takes over again. With this procedure, called *Sidman avoidance* or *continuous avoidance*, avoidance responding can be studied independently of escape responding; shock can be prevented by avoidance responses, but once delivered the shock is so brief that there is little if any opportunity to escape from it.

Data for one rat's lever pressing are shown in Figure 8-3, which shows response rate as a function of RS interval, with SS interval as a parameter. Across functions, the RS interval that produced the highest rate depended on the SS interval. Consider an avoidance schedule with a 6-s RS interval and a 2-s SS interval. Any responding at all postpones shock, and a rate of one response every 5 s or so avoids shock completely. When the RS interval is shorter than the SS interval, however, some patterns of pressing increase rather than decrease shock. For example, consider a 2-s RS interval and a 6-s SS interval. A rat that never presses receives shock every 6 s, or 10 per minute. But if the rat presses every 3 s, a shock is delivered 2 s after each press and the rat will receive about 20 shocks per minute (strictly speaking, if this increase in shock rate reduced responding, it would be appropriate to say that responding was punished). The rat can avoid shock completely by responding so fast that 2 s never go by without a press, but the increases in shock produced by the lower rates the rat probably started with might prevent it from ever attaining such a performance.

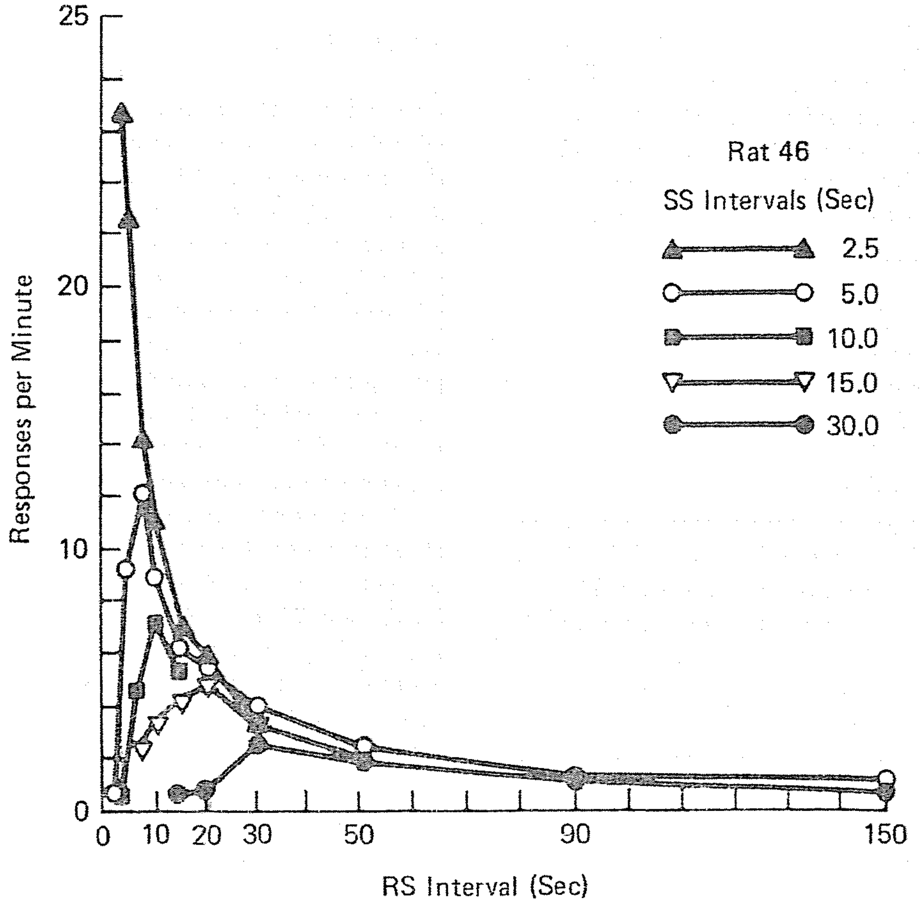

Figure 8-3 Rate of Rat 46's lever pressing as a function of RS interval, with SS interval as a parameter. The SS interval is the time between shocks if no responding occurs between shocks. The RS interval is the time by which each press postpones the next possible shock (from Sidman, 1953, Figure 1A).

Avoidance behavior may be persistent after a long history of avoidance; as we'll see below, it can be slow to extinguish. But the consequence of an effective avoidance response is that nothing happens: The aversive event is successfully avoided. So **despite the persistence of avoidance behavior once it is adequately in place, it is often hard to get it started.** This may explain why safety measures and other preventive procedures are not often shaped by natural contingencies. Someone who has never been in a serious automobile accident may be less likely to fasten a seat belt than someone who has, and someone who has never had a bad experience with fire may be less likely to install a smoke detector than someone who has.

One significant problem in medicine is the compliance of patients with regimens such as taking prescribed medications. Many patients stop taking their medications once their symptoms have disappeared even though further doses may have continued benefits. And with preventive medication, such as vaccination or the maintenance of vitamin or mineral supplements that can prevent dietary deficiencies, taking a dose is followed by nothing happening right from the start. This problem exists over a wide range of preventive measures, from safe sex to purifying drinking water to using sterile surgical equipment. We should not be surprised that such measures are sometimes difficult to shape up and maintain.

The Nature of the Reinforcer in Avoidance

Another issue is specifying what reinforces avoidance responding. When a successful avoidance response occurs, its important consequence is that nothing happens. How can the absence of an event affect behavior? According to one view, organism avoid because they are escaping from some feature of a situation that had previously included aversive stimuli. Earlier procedures in which a warning stimulus preceded shock and the organism prevented shock by responding in the presence of that stimulus made this view plausible. Avoidance was most easily acquired when the avoidance response not only prevented shock but also turned off the warning stimulus. It was assumed that the aversiveness of the warning stimulus was established through its consistent relation to shock, and therefore that this immediate consequence, escape from the warning stimulus, was the effective reinforcer. In fact, one purpose of Sidman's avoidance schedule was to demonstrate avoidance without a warning stimulus.

The view that the warning stimulus acquired its own aversive properties through its consistent relation to the aversive stimulus and that the termination of the warning stimulus therefore reinforced the avoidance response was called a *two-process theory of avoidance* (e.g., Kamin, 1956). The main issue was the nature of the reinforcer. Some theorists were willing to accept the fact of avoidance without appealing to some reinforcing event occurring at the moment of the avoidance response. Others felt that it was necessary to specify such an event. When the event couldn't be located in some environmental warning stimulus that terminated with a response, as in Sidman avoidance, these theorists instead put it inside the organism. They argued, for example, that the organism's state just after a response, when shock wasn't imminent, was reinforcing relative to its state long after a response, when shock might occur at any moment. The latter state acquired aversive properties through its relation to shock and, like a warning stimulus, was terminated by a response. This change of state presumably occurred immediately with each avoidance response, so it was assumed that an account based on other temporally remote events was unnecessary.

Debates on the status of such avoidance theories have a complex history (e.g., Anger, 1963; Herrnstein & Hineline, 1966; Hineline, 1977), with some even changing sides (e.g., Schoenfeld, 1950, 1969). Gradually the issues shifted to experimental questions about the conditions under which avoidance responding can be maintained: Must the organism be able to reduce the total number of shocks in a session or is it sufficient for the organism to be able to postpone individual shocks even though the same number of shocks is eventually delivered? It turns out that either condition can maintain avoidance responding, so these questions were soon replaced by ones about the establishing operations that make negative reinforcers effective (Hineline, 1970, 1981).

An establishing operation that makes positive reinforcers more effective is *deprivation*. The analogous operation for negative reinforcers is *presentation* (it would be called *satiation* were the stimulus food instead of shock); it is the presentation of aversive stimuli that makes their removal reinforcing. Even more so than with positive reinforcement, these establishing effects must be distinguished from discriminative, eliciting and other effects of stimuli. Furthermore, the relativity of reinforcement holds for negative as well as positive reinforcement. An avoidance procedure involves response-contingent transitions from one situation to another, and its effectiveness is determined by the situation following the avoidance response as well as the one preceding it.

An example is provided by a case in which not every shock is avoidable. Can avoidance responding be maintained when responding reduces the likelihood of shock but, unlike Sidman avoidance, doesn't reliably postpone each one? Herrnstein and Hineline (1966) arranged an avoidance schedule in which a rat was shocked with some probability at the end of every 2 s. Shock occurred with one probability if no lever press occurred and with a different probability if a press did occur. For example, in one condition a starting shock probability of 0.3 was reduced to 0.2 by a press: If the rat pressed at least once every 2 s, it reduced the shock rate from 9 to 6 shocks/min. Lever pressing was maintained by this procedure.

Now consider a case in which each lever press raises the shock probability from 0.1 to 0.2: by pressing at least once every 2 s the rat increases the shock rate from 3 to 6 shocks/min. Even though the consequence of pressing is still a shock probability of 0.2, the rat stops pressing. Relative to a starting shock probability of 0.3, the transition to a shock probability of 0.2 reinforces responding; relative to one of 0.1, it punishes responding.

Shocks were delivered probabilistically in this procedure, so even with shock probabilities that maintained avoidance responding (e.g., 0.3 reduced to 0.1 by lever pressing) some presses were immediately followed by shock and some periods without presses passed without shock. Thus, no consistent temporal relation existed between individual responses and individual shocks. Nevertheless, the likelihood of shock increased more rapidly on the average over time after no pressing than over time after a press. Thus, those who sought an immediate consequence of responding could argue that the reduction in average aversiveness produced by a lever press was sufficient to reinforce avoidance responding. And so the debate went.

These positions illustrate the difference between molecular and molar analyses of behavior. The former deals with behavior in terms of moment-to-moment sequences of events in a given setting; the latter deals with properties that can be measured only over extended time periods. For example, a molecular approach to avoidance examines the individual time intervals separating individual responses and individual shocks, whereas a molar view examines the more general relation between rate of responding and rate of shock over large samples of responses and shocks. But these categories do not have clear boundaries. Molar rates are determined by sampling discrete events over extended time periods, and distributions of molecular interresponse times or other detailed features of behavior also call for measures sampled over time.

The issues have not been and perhaps cannot be resolved. The relative importance of molecular or molar properties of behavior can be different in different settings. It is also reasonable to assume that evolution has equipped organisms with the capacity to respond differentially to many properties of the situations in which they find themselves. Situations can be created in which a rat postpones shocks within trials even though it doesn't reduce the overall shock rate, or in which it reduces the overall shock rate even though responding shortens the time to the next shock (Hineline, 1981). There is no reason to assume that an organism whose responding is determined by the molecular properties of one situation (e.g., the consistent temporal relations between responding and shocks created by the RS and SS intervals of Sidman avoidance) will be incapable of responding according to the molar properties of another (e.g., the consistent overall relations between response rate and shock rate created by a probabilistic avoidance schedule). If so, it is not a matter of choosing one or the other approach but rather of deciding which is more appropriate to the analysis of a given situation (Hineline, 2001).

Extinction after Negative Reinforcement

As with positive reinforcement and punishment, the effects of negative reinforcement are temporary. And as with these other operations, the effects of terminating contingencies between responses and aversive stimuli must be distinguished from the effects of simply terminating the aversive stimuli. In shock escape, turning off the shock eliminates responding simply because there is no occasion for escape in the absence of the aversive stimulus. But in avoidance, turning off the shock source is sometimes considered an extinction procedure. If avoidance responding is so well maintained that shocks are rare, the absence of shocks will make little difference and responding will continue for a long time. In fact, one widely acknowledged property of avoidance responding is its persistence after aversive stimuli are discontinued. For that reason, some have seen avoidance as relevant to some cases of persistence of human behavior, as in compulsions.

Consider the alternatives. With food reinforcement, we can arrange extinction either by turning off the feeder or by breaking the connection between responses and the feeder. Both have the same effect: Food is no longer delivered. That isn't so with shock escape or avoidance. Shock continues if responses can no longer remove or prevent it. In Sidman avoidance, for example, all shocks would be determined by the SS clock; responses would no longer operate the RS clock. This procedure would discontinue the response-shock contingency, but it would also increase the number of shocks if responding had been keeping shock rate low. Thus, by itself this procedure could not separate the effects of changing shock rate from those of changing the contingency.

Discontinuing the aversive stimulus has been the more common extinction procedure in avoidance, but presenting the aversive stimulus while discontinuing the consequences of responding more closely parallels extinction after positive reinforcement. The time course of extinction depends on which procedure is used and on how the procedure affects the rate at which aversive stimuli occur (cf. Hineline, 1977, pp. 377-381). In any case, paralleling our accounts of extinction after positive reinforcement and of recovery from punishment, extinction after negative reinforcement shows that the effects of negative reinforcement are temporary.

THE LANGUAGE OF AVERSIVE CONTROL

The presentation or removal of stimuli can reinforce or punish behavior. Reinforcement is most effective if the reinforced response is compatible with the responding occasioned by the reinforcer. Inversely, punishment is most effective if the punished response is incompatible with or at least independent of the responding occasioned by the punisher. Thus, it may be easy to reinforce jumping with shock removal (escape) but hard to punish it with shock presentation.

Stimuli that can reinforce by their presentation can also punish by their removal, and vice versa. We called shocks *punishers*, *negative reinforcers* and *aversive stimuli*. We introduced each usage in a different context, but this was fitting because context determines the behavioral functions of any stimulus. Thus, we introduced aversive stimuli in connection with elicitation, punishers in the discussion of consequences that reduced responding, and negative reinforcers in the treatment of consequences that increased responding. It would be convenient if we could always assume that each term identified different aspects of a single category of events.

This assumption is probably correct for many stimuli much of the time. If we know a stimulus is effective as a punisher, we can reasonably expect it to be effective as a negative reinforcer. Such consistencies justify calling it aversive and are to be expected because these categories have their origins in relations among the probabilities of different response classes. But these very probabilities should remind us of the relativity of reinforcers and punishers. We must not take too much for granted. The fact that we may easily reinforce jumping with shock removal whereas we may not so effectively punish it with shock presentation shows that the symmetry of reinforcement and punishment has limits.

Failures of symmetry between reinforcement and punishment have perhaps encouraged attempts to reduce either one to a special case of the other. Some instances of punishment have been described in the language of *passive avoidance*: by not responding, the organism was passively avoiding the stimulus arranged as a punisher for responding. But then we might as well say that not responding (e.g., not stepping down from a platform onto an electrified grid) is a response that can be reinforced. If the language works in this case, why not in any case of punishment?

The question of what counts as behavior is implicit in such arguments. Whenever responding is punished, we could say that not-responding is reinforced; whenever responding is reinforced, we could say that not-responding is punished. Once we extend our vocabulary that way, the difference between reinforcement and punishment seems to vanish. Yet we usually can tell the difference between cases of reinforcement and cases of punishment. In fact, the differences are often of serious concern. For example, we must not be indifferent to whether a parent reinforces a child's cooperative behavior with praise or punishes the failure to cooperate with beatings.

It is easier to speak in terms of discrete responses than in terms of their absence. When we can, therefore, we'll choose direct descriptions in terms of recordable responses such as lever presses or key pecks over indirect ones in terms of what is not happening. An organism exhibits more or less behavior at different times, and we needn't assume that all failures to act are in themselves actions. Just as we don't have to achieve absolute zero to acknowledge that temperature is a dimension that varies in quantity, we don't have to produce a totally nonbehaving organism to acknowledge that an organism's behavior is a dimension that can change in quantity. The behavior called not-responding (or other behavior) is a class that allows the totality of behavior to be constant, so that when summed the probabilities always add up to one, but that gives us one degree of freedom too many. If we punish a response, we should know what happens to it before we go looking for an account in terms of other behavior, and we should know what else the organism is doing before we start speaking of not-responding as behavior.

ADDENDUM 8A: SPECIES-SPECIFIC DEFENSE REACTIONS

An advantage of avoidance over escape procedures is that the reinforced response occurs in the absence of the aversive stimulus. Thus, other responses generated by the aversive stimulus don't compete continuously with the avoidance response. But just as it is easier to turn some responses than others into escape responses, it may be easier to turn some responses than others into avoidance responses. For example, pigeons avoid shock more readily if they can do so by moving from one side of a chamber to another than if they can do so only by pecking (e.g., MacPhail, 1968). In avoidance procedures, it is more difficult to make the case that such differences depend on competition between reinforced responses and other responses generated by the aversive stimulus. It has been argued instead that the differences arise because organisms are variously equipped with defense responses that are species-specific. If so, success with avoidance procedures may depend upon whether the experimenter chooses responses that the organism is already prepared to emit in aversive situations (Bolles, 1970; Seligman, 1970).

Bolles summarizes the argument as follows:

> What keeps animals alive in the wild is that they have very effective innate defensive reactions which occur when they encounter any kind of new or sudden stimulus.... These responses are always near threshold so that the animal will take flight, freeze, or threaten whenever any novel stimulus event occurs.... The mouse does not scamper away from the owl because it has learned to escape the painful claws of the enemy; it scampers away from anything happening in its environment, and it does so merely because it is a mouse. The gazelle does not flee from an approaching lion because it has been bitten by lions; it runs away from any large object that approaches it, and it does so because this is one of its species-specific defense reactions. Neither the mouse nor the gazelle can afford to learn to avoid; survival is too urgent, the opportunity to learn is too limited, and the parameters of the situation make the necessary learning impossible. (Bolles, 1970, p. 33)

It is generally accepted that species-specific behavior often limits what can be learned. Consider, for example, the transition from elicited to emitted responding in a rat's acquisition of signaled avoidance (a deletion procedure). In this procedure, a warning stimulus such as a buzzer precedes shock. If the rat responds during the buzzer and before shock, the shock is omitted; if it responds after the shock starts, the shock is removed.

In such experiments, a frequent response choice is some form of locomotion, such as jumping a hurdle or running from one side of the chamber to the other. Furthermore, the avoidance response is typically the same as the escape response. With rats, such locomotor responses are likely to be elicited by aversive stimuli even in the absence of a response-shock contingency. Once they've been produced by shock, they may continue when shock is absent. Thus, the rat's first few avoidance responses may occur mainly because of their earlier elicitation by shock (cf. Chapter 4 and Azrin, Hutchinson, & Hake, 1967). After avoidance responding begins, it is an experimental question whether it continues because it prevents shock or because it is an instance of species-specific behavior that is easily generated by aversive situations.

Some differences may also depend on species-specific determinants of what is aversive. For example, demonstrations of escape from or avoidance of the sound of running water by beavers raise the intriguing possibility that the aversiveness of such sounds contributes to the building and maintenance of their dams and lodges (Hartman, 1975).

Chapter 9
Operants: The Selection of Behavior

Operant, *a class of responses, can be traced to the Latin* opus, *work, which is also a source of* operation *and* copy. Class *can be traced to the Latin* classus, *a division of Roman citizens eligible for military draft and perhaps thereby a summons or call. In* classify, *it is linked to the suffix,* -fy, *a form of the Latin* facere, *to do.* Shape *is derived from an Old English form,* gesceap, *a creation, perhaps from an earlier root verb related to* shave *or* scrape.

We can change behavior by presenting stimuli and by arranging consequences for responding. One way to discuss the effects of these operations is in terms of the relative positions of responses in a behavior hierarchy: The organism's behavior consists of a repertoire of responses, each with a different probability. We can make some parts of the repertoire more likely than others by reinforcing some and not reinforcing others. But if we restrict our attention only to these responses, we miss some of the most interesting features of the changes in behavior we call learning; we miss how an organism comes to respond in new ways. How can we add new responses to an organism's repertoire? We'll find one answer in a procedure called *shaping*, which gives us a way to generate new responses. Shaping will also help us to see how reinforcement can create classes of responses. These classes are our units of behavior, which we call *operants*.

The rat's lever press and the pigeon's key peck have often been our examples of responses. But if we just place a rat in front of a lever or a pigeon in front of a key, we might not get presses or pecks. We might be lucky. With some organisms, responses might occur right away. But others might go so long without responding that we lose patience. Reinforcement can't work on a response that is never emitted. Fortunately, we have an alternative. Instead of waiting for the response, we can generate it by successively reinforcing others that more and more closely approximate it.

KEY TERMS: Operant, Descriptive Operant, Functional Operant; Shaping, Successive Approximations; Differentiation and Induction; Function, Topography.

SHAPING: DIFFERENTIAL REINFORCEMENT OF SUCCESSIVE APPROXIMATIONS

Suppose I have a pigeon new to the laboratory. I want to start it in an experiment, but I can't do that until it has begun to peck the key. I've already given the pigeon some time to get used to the chamber. When I operate the feeder the chamber lights turn off and a light comes on inside the feeder opening. Pretty soon this pigeon comes to the lit feeder to eat. Once it promptly starts to eat whenever the feeder operates, I operate the feeder only when the pigeon turns toward the key. After I reinforce two or three movements toward the key, I then reinforce not just any movement toward the key but only those including forward motion of the beak. By this time, the pigeon is spending most of its time in front of the key, so I can shift my attention from its turning toward the key to its forward beak movements. These more closely approximate key pecking than turns toward the key, and once I've guaranteed that they will continue by reinforcing them, I no longer have to reinforce turning toward the key. By this time, the pigeon's beak movements are full-fledged pecks and soon one strikes the key. At this point, I can withdraw, because the apparatus is arranged so that further pecks operate the feeder automatically.

An experienced experimenter can shape a pigeon's key peck with just 10 or 15 reinforcer deliveries. Some aspects of skill in shaping can be stated explicitly. For example, reinforcing movements is more likely to shape responding efficiently than reinforcing postures. Other aspects can't be formulated so readily. For example, shaping usually compromises between extremes of frequent and infrequent delivery of reinforcers. Frequent delivery leads to quicker satiation and may overly strengthen some early responses that later won't be part of the response to be shaped. On the other hand, infrequent delivery may reduce responding in general, and once the organism becomes inactive the progress in shaping up to that point may be lost. The experimenter must work within the limits imposed by these extremes, but no explicit rules exist for judging just where these limits lie (Eckerman *et al.*, 1980; Platt, 1973).

A shaper who can work close to the limits of extinction, thereby not reinforcing lots of responses now that will have to be extinguished later when the distribution of responses is closer to the target response, will probably finish shaping more quickly than one who is generous with reinforcers but risks satiation before reaching the target. Evolutionary contingencies work that way too, in that more extreme changes in the environment will ordinarily produced more rapid evolution, as long as they are not so extreme as to wipe out the population: "Since we know that natural selection is better at improvements than creation, selection can occur efficiently on one of the extreme states, which many if not all members of the population could occupy" (Kirschner & Gerhart, 2005, p. 74).

Some features of shaping are fairly specific to the particular response and organism being studied, whereas others are relevant to shaping a variety of responses in a variety of organisms. For example, an experimenter who has worked often with pigeons knows that reinforcing a small beak movement aimed directly at the key will more quickly produce key pecking than reinforcing a large sidewise beak movement that finishes in front of the key. On the other hand, whatever the response and the organism, an opportunity to reinforce a response shouldn't be missed if it more closely approximates the response to be shaped than any other response that has been reinforced before. But no set of rules really substitutes for what you learn about shaping by actually doing it. It is crucial to be sensitive to the moment-to-moment interactions of the organism's behavior with the reinforcers you deliver. In the course of those interactions, your behavior of shaping will itself be shaped.

The different properties of different responses that might be shaped make shaping an art. This art may be applicable to many skills: gymnastics; lovemaking; playing a musical instrument; seduction; penmanship; and setting someone up as a victim of a con game. As these examples suggest, shaping can be put to either good use or bad, and many use it without even knowing they are doing so. As with reinforcement and punishment, when shaping is put to good use, it might as well be done effectively; when it is put to bad use, the best defense against it is knowing how it works.

Shaping is based upon *differential reinforcement*: At successive stages, some responses are reinforced and others aren't. In addition, the criteria for differential reinforcement change as responding changes, in *successive approximations* to the response to be shaped. The property of behavior that makes shaping effective is that behavior is variable. No two responses are the same, and reinforcement of one response produces a spectrum of responses, each differing from the reinforced response along such dimensions as topography (form), force, magnitude and direction. Of these responses, some are closer to the response to be shaped than others and may be selected to be reinforced next. Reinforcing these responses produces still others, some of which may come even closer to the response to be shaped. Thus, reinforcement can be used to change the spectrum of responses until the one to be shaped occurs.

This aspect of shaping is sometimes supplemented by other effects of reinforcers. Some reinforcers increase activity. For example, food delivery makes a food-deprived pigeon active. It is therefore difficult to use food to reinforce its holding of a posture (Blough, 1958b). Thus, a response more closely approximating the one to be shaped occasionally occurs simply because the delivery of some reinforcers makes an organism more active.

There is a paradox to shaping. Reinforcement is said to raise the probability of the reinforced response. But no response is ever repeated exactly. How then can we appeal to reinforcement as the basis for the shaped response when reinforcers are delivered after responses that only approximate that response? In fact, if individual responses are never repeated how can we speak of reinforcement at all? We acknowledged this problem in Chapter 1 but didn't resolve it. Later we'll see that we can't deal with just single responses; we must deal with *classes* of responses.

Natural and Artificial Selection in Shaping

Shaping is a variety of selection that provides an ontogenic parallel to the phylogenic selection that occurs in biological evolution (Donahoe, Burgos, & Palmer, 1993). It is most obvious when used by a human trainer, as in the teaching of skills to a seeing-eye dog (cf. Pryor, Haag, & O'Reilly, 1969; Squier, 1993). In such cases, it is an example of *artificial* selection, just as the breeding of cattle is artificial selection in the phylogenic case. But shaping can also occur as a result of *natural* contingencies. "Evolution is then not a matter of 'waiting' for a new mutation, but *the preferential increase or decrease of alternative forms in response to changes in conditions*" (Carroll, 2006, p. 163).

For example, male cowbirds in different parts of the United States sing different dialects of birdsong (as is usually the case among songbirds, the female cowbird doesn't sing). A female is most likely to respond with mating-pattern postures to songs that sound most like the ones she heard in her youth, which were in the dialect of local males. When a foreign male is introduced, he begins singing in his own dialect. But he sings with variations, and the more time he spends in her presence, the more his song takes the form of the local dialect. His acquired dialect is a product of natural ontogenic selection: her differential reactions were reinforcers, and they shaped his song (e.g., King & West, 1985).

Shaping typically involves quantitative changes along one or more dimensions of an organism's behavior, but sometimes it appears to produce qualitative changes. Consider the following example involving the shaping of a rat's high-force lever presses (Catania & Harnad, 1988, p. 476). This example is based on a classroom demonstration developed by Skinner; I later constructed a similar apparatus and used it in my own classes over a period of many years.

In this apparatus the rat produces food by pressing on a counterweighted lever that protrudes into the chamber high enough from the floor that the rat must stand on its hind legs to reach it. With the weight set at a modest level the rat presses easily by resting one or both forepaws on the lever. Over successive reinforcers, we gradually raise the weight until it is near the rat's own weight. But once the rat's weight is exceeded, pressing is effective only if a new topography emerges. Pushing down on the lever with both hind legs on the floor doesn't work any more; instead, the rat's paws come up off the floor. While hanging on the lever the rat must now lift its hind paws to the chamber wall, where a wire mesh allows it a firm grip. Even with its own weight exceeded, the rat can now depress the lever by pulling between its forelegs and hind legs. Whether the rat makes the change to the new topography depends jointly on its behavior, its anatomy and the apparatus. For example, it will be less likely to produce the new topography involving both forelegs and hind legs if it got up to high-force pressing by jumping on the lever rather than by pressing it.

This example illustrates two kinds of ontogenic selection, one gradual and the other relatively abrupt. The first occurred with the counterweight below the rat's own weight, and the second occurred when that weight was exceeded. There are analogies in phylogenic selection, as when the gradual quantitative changes produced by selection relative to a population mean are contrasted with the more abrupt changes produced by catastrophic environmental events. In phylogenic evolution, the latter changes are sometimes called *saltations* (cf. Chapter 3).

One dimension of behavior that may be shaped is variability itself (Neuringer, 2002, 2004). When behavior is more variable, it provides more variants that may be reinforced and therefore may make shaping quicker and easier. Here too we can identify phylogenic parallels. The *genotype* is the genetic makeup of an organism, whereas the *phenotype* is the organism as it has developed within its environment. For example, two species of fish that live in highly similar lake environments may have developed highly similar phenotypes, but the genotype of one might be far more variable than that of the other. Research in conservation biology has shown that the range of genotypic variations available to the former species make it more likely than the latter species to survive drastic changes in its lake environment (e.g., Reed, 2007; Vellend, 2006).

One more dimension of shaping is illustrated by the rat and the counterweighted lever. Ordinarily in shaping, you watch current behavior and deliver reinforcers based on what you see. But with the weighted lever, you don't have to watch and wait. Once you've set the required force, you won't see the forces of the rat's successive presses. Instead, you will know by whether or not food was delivered whether a particular press exceeded the threshold force and produced food or failed to do so and was not. In shaping procedures that set thresholds your judgments are not which responses to reinforce but rather how quickly you can move the threshold up. If you move up too quickly and responding stops, you might have to drop the threshold down pretty far to bring back responding before you can start moving it up again. Here again, rules cannot substitute for practice with the procedure.

Shaping as Signal Detection

As we have seen, shaping is a skill you might acquire through instruction or through the nonverbal contingencies of actual practice. The momentary distribution of responses includes some closer to my shaping target than others. If I reinforce a response from the end of the distribution closest to my target, the distribution shifts in that direction; if I reinforce one from the other end, it shifts away. The distribution is always changing, so my task is to judge where within it each response comes from. If it comes from the end near my target, I should reinforce it; if it comes from the other end, I should let it go without reinforcement. If it comes from in between, whether I reinforce or not should depend on the time since the last reinforced response and on how many reinforcers I've already delivered; if I go too much one way or the other responding will stop, either because of extinction or because of satiation.

What I do while shaping has many of the characteristics of *signal detection* (Green & Swets, 1966;

Nevin, 1969; Swets, Dawes, & Monahan, 2000). Signal detection theory assumes that an observer responds or doesn't respond to a stimulus that consists of either a signal in noise or noise alone. Responding to the signal is a *hit*, but responding to the noise alone is a *false alarm*; not responding to the signal is a *miss*, but not responding to the noise alone is a *correct rejection*.

For our purposes, let's say that a response from the top half of the distribution should be reinforced, but one from the bottom should not. In other words, we should reinforce only responses that are above average and not those that are below average. Interpreted in this way, shaping can be analyzed as an instance of signal detection. Signal detection theory is usually applied to research on sensory systems and related topics, but it is appropriate here and we'll encounter it again later.

The possible outcomes are shown in Table 9-1. If the response is above average and I reinforce it, that's a hit; it moves the distribution of response probabilities toward my target. If I don't reinforce it, that's a miss. If the response is below average and I reinforce it, that's a false alarm; that reinforcer is wasted and moves the distribution away from my target. If I don't reinforce it, that's a correct rejection.

Once we know the proportions of hits and of false alarms for above and below average responses we can use them to calculate two statistics. One, *d'* (*d*-prime), is an index of the shaper's overall sensitivity to the current distribution; it depends on hits and correct rejections (true responses) relative to total responses. It grows larger with increases in hits and correct rejections and smaller with increases in misses and false alarms. A second statistic, *bias*, is an index of whether one type of response is more likely than the other. For example, if a false alarm matters more than a miss, you'll be biased toward not reinforcing rather than toward reinforcing; this might make you more likely to lose the behavior to extinction. But if you believe the miss is more costly and a false alarm will not have serious consequences, you'll be biased toward reinforcing rather than not reinforcing; this might make you more likely to lose behavior to satiation.

Judgments like these are the skills of someone who has become good at shaping, but they aren't easily expressed as rules. The skills come from being sensitive to the ways in which response distributions change from moment to moment over the course of shaping. As we'll see when we get to verbal behavior (especially Chapter 23), the learner who depends mainly on instructions may be less likely to master such judgments than one who is less dependent on them. We often make assumptions about new situations and those assumptions often take the form of instructions that we've generated for ourselves (cf. Rosenfarb *et al.*, 1992). But our words can get in the way and make our behavior insensitive to some of the contingencies that would otherwise shape and maintain it. That is probably why teachers are sometimes torn between telling students how to do something and letting the students discover it for themselves.

Shaping with Pets and with People

Animals used in combat and espionage and police work have a long history. On the eastern front in World War II, Russian dogs trained to run underneath German tanks did so with explosives strapped to their backs. As with the training of circus animals, these and other kinds of work with animals were often based not on behavior analytic principles but on practices that had evolved over time. The evolved or "intuitive" skills of existing trainers have gradually been integrated with the methods of behavior analysts to create effective new techniques. For example, dolphins were trained to locate underwater mines during both Gulf Wars, United States airports have long been patrolled by dogs trained to sniff for drugs and explosives (Johnston, 2000), and pouched rats have been taught to locate undetonated landmines (Poling *et al.*, 2011).

Table 9-1 Signal-detection contingencies in shaping: The entries show the consequences of reinforcing responses that are above or below average relative to a shaping target.

Reinforce the Response?	*Above Average Response*	*Below Average Response*
Yes	Hit (True Positive)	False Alarm (False Positive)
No	Miss (False Negative)	Correct Rejection (True Negative)

Skinner's discovery of shaping was a by-product of World War II. In the days before guided missiles Skinner had undertaken a project to train pigeons to guide aircraft toward enemy targets (Skinner, 1960). He and his coworkers discovered shaping virtually by accident when using hand switches rather than an automated system to deliver reinforcers (Peterson, 2004; Skinner, 1951). The technique was adopted by Breland and Breland (1951) for animal training and was later used to shape up animal performances for movies and television advertisements. Not long after, chimpanzees working on operant tasks became the first American astronauts (Brady, 2007).

The applications have since so expanded that they can only be touched on here. Clicker training has become an increasingly favored method for training pets (Pryor, 1999, 2009). If a small clicker or cricket is sounded just before a pet is fed or played with, the clicker itself will soon become an effective reinforcer on its own (see the discussion of conditional reinforcers in Chapter 14). The portability and small size of a clicker means that shaping need not be restricted to a narrow range of locations. Shaping has also been used with animals in zoos and in marine parks (e.g., Squier, 1993; Sutherland, 2006). Some who have seen how shaping works with pets have recognized that it may be relevant to interactions with children and spouses and other family members too (e.g., Sutherland, 2008). In any case, the stage had long been prepared in many ways for other human applications (e.g., Ayllon & Michael, 1959; Ferster & DeMyer, 1961; Lindsley, 1956).

Of course what can be shaped has limits that may depend on anatomies and other physiological constraints (don't bother trying to shape flying in either rats or elephants). But those are not the only constraints. For example, it is simply impractical to do much shaping with big snakes such as pythons. An animal that eats only large one-swallow meals spaced days and perhaps even weeks apart provides few opportunities for delivering reinforcers. Nevertheless, shaping is now all around us. We will later see that it can work not only on what we do but also on what we say.

DIFFERENTIATION AND INDUCTION

If we watched a rat's lever presses we might see the rat press the lever with either paw or with both paws or by sitting on it or perhaps even by biting it. Each is a different response; even if two presses were made with the same paw they wouldn't be identical. Nevertheless, we call all of them lever presses. On the other hand, if the rat made similar movements at the other end of its chamber, distant from the lever, we wouldn't call those responses lever presses no matter how closely they resembled the earlier ones that did operate the lever.

We can't discuss behavior just in terms of individual responses. Responses are instances of behavior, so each can occur only once; responses can have common properties but they can't be identical in all respects. Later responses will resemble the earlier reinforced response more or less closely, but they can't match it exactly. On the other hand, we can't group all responses together without distinction, because we'd be left with nothing to speak about but behavior in general. We have to settle for an intermediate level of analysis, at which we speak of neither individual responses nor behavior in general, but of *classes of responses* defined by common properties (Skinner, 1935a).

Response Classes

In experiments on the rat's lever press, the lever is attached to a switch that closes whenever the rat moves the lever with enough force through a sufficient distance. The common property of all lever presses is this consequence: Each response that closes the switch qualifies. Defining response classes in terms of common environmental effects is the basis for both recording responses and arranging their consequences. For example, an experimenter can record presses by counting closures of the switch and arrange to reinforce all such responses with food.

But this class produced by the experimenter has behavioral significance only if it is affected by the contingencies imposed on it. We must ask a fundamental question: Do consequences modify the likelihood of responses in this class? If so, it can be called an *operant* class; an operant class is a class of responses affected by the way in which it operates on or works on the environment. Lever presses and key pecks are convenient examples, but operant classes may include far more complex cases (see Chapter 10).

An operant is a response class modifiable by the consequences of the responses in it. This definition of a response class depends on behavioral properties of responding and not on physiological properties such as the distinction between somatic and autonomic responses. The behavioral properties of operant classes are based on the operation called *differential reinforcement*, the reinforcement only of responses that fall within a specified class. This operation makes subsequent responding conform more and more closely to the defining properties of the class. **The essential feature of an operant is the correspondence between a class of responses defined by its consequences and the spectrum of responses generated by these consequences.**

Some Examples of Differential Reinforcement

Imagine a food-deprived rat in a chamber with a horizontal slot 30 cm long on one wall (30 cm is about 12 in). Photocells record where the rat pokes its nose through the slot. We label successive 2-cm segments along the slot as positions 1 through 15, reading from left to right. Food can be delivered into a cup in the wall opposite the slot. With food deliveries accompanied by a distinctive noise, the rat quickly comes to the cup and eats whenever food is delivered (cf. Antonitis, 1951).

At the start, the rat spends little time near the slot. Occasionally it sniffs at it and puts its nose in it as it moves about, but these responses are relatively infrequent and aren't systematically related to the positions along the slot. A frequency distribution of the rat's responses as a function of position is shown in A of Figure 9-1.

Now let's reinforce the poking of the rat's nose through the slot, but only if the pokes occur at positions 9 through 12. The initial effect of reinforcement, in B, isn't restricted to positions correlated with reinforcement; responding increases at all positions. This phenomenon, the spread of the effect of reinforcement to other responses not included in the reinforced class, is called *induction* (an occasional synonym is *response generalization*). In the example, reinforcing responses at positions 9 through 12 affected responding not only there but also at other positions across the entire slot.

As we continue differential reinforcement, reinforcing responses at positions 9 through 12 but not elsewhere, responding gradually increases at positions correlated with reinforcement and decreases elsewhere, as shown in C, D and E. Eventually, most responses are within the boundaries that determine whether a response is reinforced, as in E, and a point may be reached at which, even though some unreinforced responses continue, the distribution of responses across positions doesn't change much further with continued differential reinforcement.

In the example, the distribution of emitted responses came to conform closely to the boundaries of the class of reinforced responses. This process is called *differentiation*; responding produced in this way is said to be *differentiated*. Differential reinforcement created a response class defined by response position. Yet if the distribution of responses in E represents the maximum differentiation possible, what can be said about the responses that continue at positions 6, 7 and 8, or 13, 14 and 15? They are outside the boundaries of the response class correlated with reinforcement. According to a strict interpretation of the defining properties of operants, they can't be counted as members of the operant class. They can be talked about in terms of induction: These responses are so close to the class of reinforced responses that the effects of reinforcement have spread to them. When responding stays within the boundaries of the reinforced class we speak of differentiation; when it falls outside we speak of induction. Nevertheless, the same procedure generates responding outside as well as inside, and all of the responding is part of one continuous distribution.

There is no difficulty if we acknowledge that the example involves two different classes. The first (1), represented by dashed vertical lines in Figure 9-1, is the basis for reinforcement; the lines show how consequences (food) depend upon response position, i.e., the conditional probability of food given a response is 1.0 for responses at positions 9 through 12 and zero elsewhere. This distribution defines a class of responses in terms of the consequences of responses within that class.

The second class of responses (2) is given by the actual performance produced by reinforcement. At any time this class is represented by the current response distribution (e.g., early during differential reinforcement, as at C, or late, as at E). The two classes needn't correspond exactly.

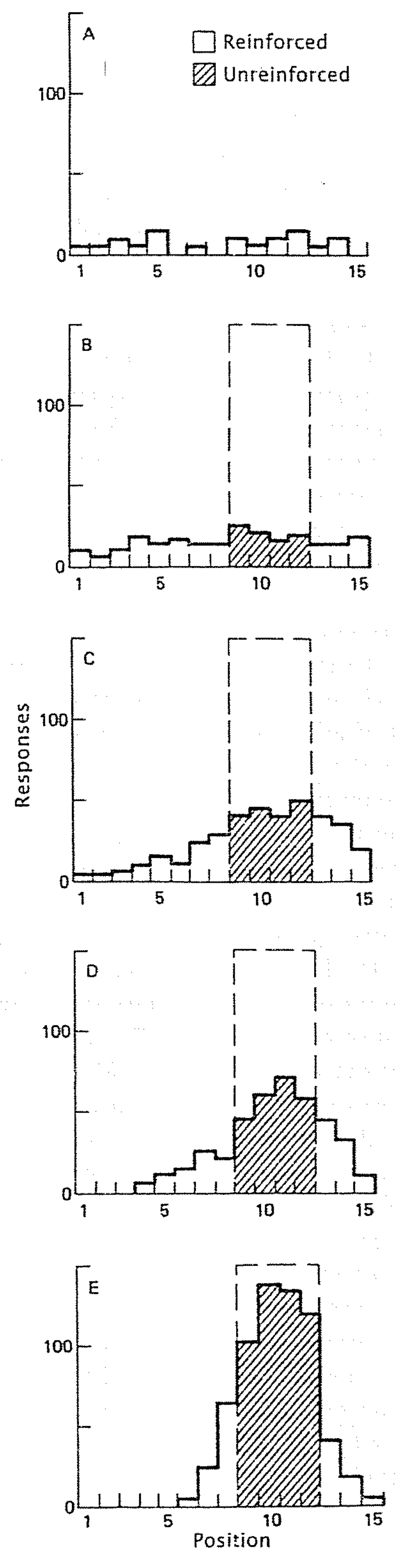

Figure 9-1 Hypothetical differential reinforcement of response location. A rat pokes its nose into a horizontal slot in the wall; photocells register this response at any of 15 positions reading from left to right. A shows the distribution of positions when no responses are reinforced. In B through E, responses at positions 9 through 12 (bounded by dashed vertical lines) are reinforced with food; filled areas show reinforced responding. In B, the effects of reinforcement spread across the length of the slot; this spread is called *induction*. In C through E, responding becomes restricted to positions correlated with reinforcement; this concentration of the effects of reinforcement is called *differentiation*. With continued differential reinforcement, the distribution becomes stable (E) and corresponds fairly closely to the response class (dashed verticals) correlated with reinforcement.

In fact, a fundamental dimension of any class of reinforced responses is the *degree of correspondence* between (1) the behavior that is reinforced and (2) the behavior generated by reinforcement. The behavior that is reinforced is sometimes called (1) the *descriptive* or *nominal* class; the behavior generated by reinforcement is sometimes called (2) the *functional* class. Figure 9-2 illustrates these classes: the x-axis is some parameter of responding, the dashed area shows how the reinforcer depends upon that parameter, and the solid curve shows the distribution produced by the contingency.

Now consider another hypothetical example, illustrated in Figure 9-3. Again we use photocells to record the positions at which the rat pokes its nose through a slot, but this time the slot is vertical rather than horizontal. The 15 positions are numbered consecutively reading from bottom to top. At the start, the rat occasionally pokes its nose in the slot as it sniffs about the chamber, but these responses mostly occur at the lower positions, as in A of Figure 9-1. Perhaps a response will eventually occur at 15, the top position, if we're patient enough. But maybe not. Shaping is a better option.

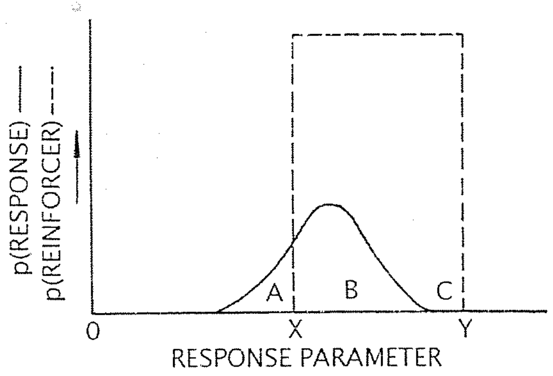

Figure 9-2 The dependence of a reinforcer on some parameter of responding (dashed line) and the distribution of responding produced by that contingency (solid curve). The two classes overlap in this example (B), but some of the distribution falls outside of the reinforced class (A), and not all regions of that class are represented in the distribution (C). (Adapted from Catania, 1973a, Figure 1)

First we arrange reinforcement for responses at position 7 or higher, as in B. Responding increases but most of it remains at the lower positions. Later we raise the criterion for differential reinforcement to 9, as in C. By this time, responding at the lower positions has decreased, the distribution has shifted toward higher positions, and responding has occurred for the first time at 13. Responding becomes more concentrated in the region correlated with reinforcement in D, when we raise the shaping criterion to 11. Finally, in E, we raise the criterion to 13. Had we imposed this criterion on the starting performance, in A, no reinforceable responses might have occurred. With differential reinforcement of successive approximations, however, the distribution has shifted to higher levels, with maximum responding at 12. Nevertheless, this maximum remains below the boundary separating reinforced from unreinforced responses, and more responses are unreinforced than reinforced. The basis for this outcome is simple: The rat can reach some positions more easily than others. But what are the implications of such an outcome for defining response classes in terms of their consequences?

In this instance, the class of responses defined by consequences (responses at positions 13 and higher, which were reinforced) differed from the class generated by those consequences (responding up to 15, but with the maximum at 12). Responding was obviously modified by its consequences, but some of it was outside the boundaries of the reinforcement criterion. We can't speak of either class alone; operants must be defined in terms of the relation between (1) the environment (the consequences it arranges for responses) and (2) behavior (the responding produced by these consequences). For convenience we'll sometimes speak of operants solely in terms of classes defined by consequences or solely in terms of the distributions generated by these consequences, but more strictly they should be defined in terms of the correspondences between the two classes.

Reinforcement inevitably involves differentiation. Responses such as lever presses and key pecks must occur at particular locations and be of sufficient force or they won't produce the consequences arranged for them. We could repeat the examples of Figures 9-1 and 9-3 by substituting such dimensions as force of a press or a peck. These examples illustrate the selection of behavior by its consequences. Within the lifetime of the individual organism, the selection of populations of responses is analogous to the selection of populations of organisms over evolutionary time. Both types of selection involve classes created by contingencies (cf. Chapters 3 and 30).

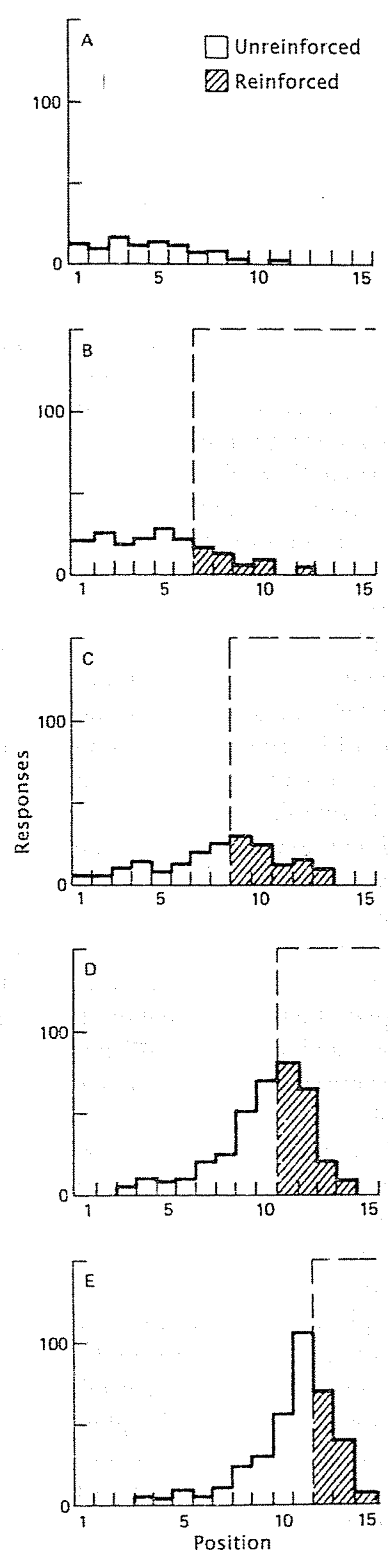

Figure 9-3 Hypothetical example of shaping (differential reinforcement of successive approximations to a response). The situation is like that in Figure 9-1, except that the slot is vertical with the 15 positions numbered from the bottom to the top. Again, unfilled areas represent unreinforced responding and filled areas represent reinforced responding. Before reinforcement (A), more responding occurs at lower than at higher positions. Next (B), responses at 7 or higher are reinforced; the dashed vertical line shows the lower boundary for reinforced responses. After this differential reinforcement has some effect, the criterion is moved up to 9 (C). Later it is moved to 11 (D), and finally to 13 (E), when shaping produces some responses at 15, the top of the slot, that weren't seen before reinforcement. At this point, the response distribution has its maximum at 12, just below the minimum required for reinforcement.

Function Versus Topography

We've seen how classes of behavior can be created through differential reinforcement. It might seem at first that these classes are defined by their topographical properties (what they look like). But that is not so. Consider even a superficially simple response like the rat's lever press. The rat might press the lever with its left paw or with its right paw or with both paws. It might also push the lever down with its chin or jump on it or sit on it. These responses look very different, but they all count as lever presses; they're all members of the same operant class. Despite their differences, they're members of that class because they share a common function: They all produce the same consequences.

In other words, **operants are defined by their functions and not by their forms.** But why should you care? It turns out that this distinction, which originated with questions about lever presses or key pecks, is crucial to significant human applications. One example is the problem of self-injurious behavior in children with developmental disabilities (e.g., Iwata *et al.*, 1994; e.g., Iwata *et al.*, 1990; Lovaas & Simmons, 1969; Wahler, 1975). Such behavior can include head banging, biting one's own flesh until drawing blood, poking a finger into the space between the eyeball and the eye socket, and many other topographies. It is dangerous behavior, and the children who engage in it can do themselves permanent damage.

Consider three male children who indulge in self-injurious behavior. They're housed in a treatment center, and each one spends time in head banging and in biting himself, so we can't distinguish among the children on the basis of the topographies of their behavior. Yet we find that the first child increases this behavior mostly when people are around but not paying attention to him; the second does so mostly when others ask him to complete some task (e.g., reading or simple arithmetic); and the third engages in this behavior independently of its social context. The self-injurious behavior of these three looks similar, but its functions are quite different. For the first child, the self-injurious behavior gets attention; for the second it avoids compliance with demands; and for the third the behavior doesn't seem to depend on any obvious environmental contingencies.

This analysis recommends very different treatment programs for each child. The first one must be taught other and more effective ways of engaging the attention of others and must be brought into situations where he'll be less deprived of attention. Tasks must be selected for the second child that are appropriate to his competence, and his success at those tasks must be reinforced (his behavior suggests that it has been too often punished). The source of the third child's self-injurious behavior is unknown, and we must consider the possibility that it has some organic source (perhaps the child was born to a drug-addicted mother and suffered prenatal damage to the developing brain).

The point is that it is more important to define behavior classes by their consequences than by their topographies. Even though the self-injurious behavior looks alike across the three children, the attention produced by that of the first child distinguishes it from the avoidance achieved by that of the second child; some aspect of self-stimulation may be involved in the behavior of the third child, but unfortunately we don't know enough about such cases.

Let's now concentrate on the first child. Suppose we try to extinguish the self-injurious behavior by ignoring it. First of all we might have trouble doing so because we can't tolerate the damage the child may do to himself. We persevere nevertheless and discover that the self-injurious behavior doesn't decrease. One possibility is that we haven't adequately identified the relevant class of behavior. Topography has again misled us. If the function of this behavior is to produce attention, it is probably part of a much larger class of behavior that includes shouting and acting up, hitting or otherwise abusing the caregivers in the treatment center, and any number of other responses that might function to get attention. This tells us how important attention is to this child. It also reminds us that we can't define response classes by what they look like. In this case, we must define a treatment program that uses attention to reinforce more effective and appropriate behavior. Both the child and his caregivers will benefit if the program is successful.

The self-injurious behavior was one class of behavior embedded in the larger class of attention-getting behavior (cf. Lalli *et al.*, 1995). The larger class was held together by the common consequences of its members, just as the various topographies of lever pressing (left or right paw, both paws, sitting on it) were held together by the common consequence of producing food. Common consequences are the glue that holds classes of behavior together (cf. Malone, 1990, p. 296). Furthermore, when a class of responses seems insensitive to its consequences, as when the first child's self-injurious behavior seemed not to extinguish, we must entertain the possibility that we've improperly defined the class, and that it is part of a larger class the other members of which continue to have the consequences it once shared with them (e.g, the attention-getting of the current example; cf. Chapter 13 on higher-order classes of behavior).

The issue of function versus form is relevant at many levels. In psychiatry, for example, the "Diagnostic and Statistical Manual of Mental Disorders" (American Psychiatric Association, 2000), which is periodically revised to deal with new diagnostic categories, is the resource for identifying psychiatric categories that enter into available treatment options and medical reimbursement. Throughout its history it has emphasized form over function. At great cost to patients and to our society, many behavioral disorders continue to be defined in terms of what they look like rather than in terms of their function and many continue to be characterized in terms of disease models rather than in terms of the contingencies that may have engendered them.

Chapter 10
The Structure of Operants

Classify *links class to the suffix,* -fy, *a form of the Latin* facere, *to do. Forms of* facere *appear in* fact, modify, difficult *and* effect. *In* specify, *it is linked to* species, *a name for another type of class that is selected. By way of Latin, these come from the Indo-European* spek-, *to see or observe, and are closely related to* expect, introspect, skeptic, telescope *and* spy. *In* office, *from* opi *plus* -ficere, *and thus the doing of work, it is also linked to* operant. Work *itself has Greek origins, and is related to* organism *through the Greek* organon *(tool).*

Chain *is derived from the Latin* catena, *perhaps from an Indo-European root,* kat-, *meaning twist or twine.* Chunk, *once a cut of meat but now usually meaning a large amount, is related to* chuck, *a piece of wood or meat, and the probably earlier* chock, *a piece of wood; their earlier derivations are obscure.* Mediate, *for dividing into two equal parts, comes from the Latin* medius, *in the middle.*

When I was a post-doctoral fellow, I was asked to prepare the pigeon demonstrations for B. F. Skinner's undergraduate course at Harvard, *Natural Sciences 114*. Skinner had typically spoken of the pigeon's wing flapping as emotional and ethologically determined behavior, not easily amenable to operant control. Those of us working around the lab at the time liked to identify places where we could show that Skinner was wrong about something. When a pigeon happened to flap its wings just after I had placed it in the demonstration chamber, here was my chance. I decided to see whether reinforcers could modify wing-flapping, so I operated the feeder. The pigeon stopped flapping and ate.

I was also teaching my first course, an undergraduate course in Comparative Psychology, and had recently prepared lectures on animal locomotion. I had learned about two modes of avian flight: energetic flapping that gets birds off the ground and up to air speed; and sustained flight with wings outstretched as airfoils, with movement mostly restricted to the ends of the wings where the tips of the feathers provide forward thrust. They are functionally like back-and-forth propeller blades and work by flipping their orientation as the wings move up and down. The breast muscles that drive the energetic takeoff flight, low in myoglobin, fatigue rapidly. The postural muscles of sustained flight are richer in myoglobin and less subject to fatigue. I knew that resting the bird often was crucial. In brief sessions over the next few days, always stopping before flapping began to fatigue, I intermittently reinforced wing-flapping and gradually shaped longer sequences of flaps.

KEY TERMS: Differential Reinforcement; Chaining, Chunking, Temporally Integrated Units; Mediating Behavior; Differential Reinforcement of Low Rate (DRL), High Rate (DRH), or Other Behavior (DRO).

On the day of the demonstration, I first put other birds through their paces. Then I placed this last bird in the chamber and demonstrated the wing-flapping, complete with pauses after reinforcement and high rate runs (cf. fixed-ratio schedules in Chapter 15). I quit before fatigue set in and briefly outlined the significance of the bird's muscle physiology, facetiously closing with the comment that now that we could reinforce wing-flapping our next project would be to teach the bird to fly. Skinner, unflappable, picked up his lecture from there. The only sign of his disapproval was his quiet comment after class that I really should have told him my plans ahead of time. He never argued with data, so he had no problem with the reinforcement of wing-flapping (Catania, 2002, pp. 333–334).

The point, of course, is that what can be shaped has limits, and successful shaping sometimes depends on knowing more than just when to deliver a reinforcer. Responding can vary not only in location or force but also in topography or form, direction, and so on. Differential reinforcement can be based on any response dimension, so any dimension or combination of dimensions might provide the defining properties of an operant class. We'll now explore the differentiation of operant classes along dimensions that have entered into significant theoretical issues in the history of the psychology of learning. Our examples include the temporal organization of behavior, its sequential patterning, and its novelty and variability.

DIFFERENTIAL REINFORCEMENT OF TEMPORAL ORGANIZATION

Temporal properties of responding include latency, duration and rhythm, any of which may provide the basis for differential reinforcement. For example, a procedure that differentially reinforces long latencies might reinforce a pigeon's peck on a lit key only if some minimum time has elapsed since the light turned on. With criterion latencies shorter than 10 s, the pigeon's mean latency usually exceeds criterion, so more pecks are reinforced than unreinforced, but with a longer criterion fewer latencies are long enough to qualify for reinforcement, so fewer pecks are reinforced. The tendency for responding to exceed a short latency criterion and to fall below a long one is a common feature of the differential reinforcement of temporal properties of behavior in human as well as nonhuman performances (Catania, 1970).

Differentiation of temporal properties of behavior is relevant to reinforcement schedules, treated in detail in Chapter 15. Substantial quantities of behavior can be maintained even when only occasional responses produce a reinforcer. In reinforcement schedules, responses can be made eligible for reinforcement based on the number of responses emitted, the time elapsed since some event, or some combination of these conditions. A property of special interest is the separation of responses in time.

For example, you'll successfully start a lawn mower with a flooded engine only if you wait long enough after your last attempt to start it; conversely, you'll successfully inflate a bicycle tire using a pump with a leaky connection to the valve only if you pump rapidly enough. Analogously, a pigeon's key peck might be reinforced only if it is preceded by 10 s of no pecking, or only if pecking has been so rapid that at least 20 pecks were emitted during the last 5 s. In the first case, pecking decreases; in the second, it increases.

If we concentrate just on the pecks, we might be tempted to say that the first case doesn't qualify as reinforcement, because pecking decreased. But the unit of responding reinforced in this case isn't just a peck; it is a sequence consisting of a pause plus a peck. As this combination becomes more likely (thereby demonstrating the effect of the reinforcers on the response class producing them), the rate of pecking necessarily decreases.

The procedure in the first example, with responses reinforced only if preceded by a minimum time without a response, is called a *differential reinforcement of low rate* or *DRL* schedule (Ferster & Skinner, 1957); it is sometimes also called a schedule of *interresponse time* or *IRT* reinforcement (Malott & Cumming, 1964), because reinforcement is based on the spacing in time of individual responses rather than on the average rate generated by many responses over an extended time. In general, the longer the interresponse time required for reinforcement, the lower the rate of responding. In DRL performance, the rate of responding decreases because the likelihood of responses preceded by long pauses (long IRTs) increases. Thus, DRL contingencies differentially reinforce a complex operant consisting of an IRT plus a response, in that order.

In the second example, with reinforcement depending on 20 or more pecks within 5 s or less, responding might increase, but only because the high rate of responding itself is differentially reinforced. Such a schedule is called a *differential reinforcement of high rate* or *DRH* schedule. The DRH schedule has received less research attention than the DRL schedule, mostly because it is harder to work with DRH than with DRL schedules. Consider a pigeon whose pecking has been raised to more than 4 pecks/s by a DRH schedule requiring 20 pecks/5 s. (In fact, skillful experimenters can routinely produce rates in excess of 10 pecks/s in pigeons by gradually raising the criteria.) As long as the pigeon maintains this high rate, pecks produce reinforcers frequently. But if for any reason the pigeon slows down, pecking meets the rate criterion less often and thus produces fewer reinforcers. This in turn reduces the pigeon's rate of pecking and so leads to a further decrease in reinforcers. This vicious circle will probably end in a rate of pecking so low that the reinforcement criterion is never met, so the pigeon's pecking ceases completely. To reinstate responding, the experimenter has to gradually bring the pigeon back up to the high rate through shaping.

In a DRL schedule, however, an initial effect of reinforcers following pecks is that pecks occur more often. This higher rate of pecking means that pecks occur closer in time, so fewer IRTs are long enough to meet the reinforcement criterion. Thus, reinforcers decrease and the rate of pecking decreases in turn. But this decrease in rate simply makes it more likely that IRTs will again be long enough to meet the criterion. Thus, pecking oscillates between increased rates accompanied by decreased reinforcement and decreased rates accompanied by increased reinforcement. Pecking can be maintained over long periods by such contingencies.

The significance of these examples is that we must be cautious about taking response rate as a fundamental measure of the effects of reinforcement. Like latency and other measures (cf. reflex strength in Chapter 4), response rate, once regarded as a direct measure of response strength, was recognized as a property of behavior that could be differentiated like other properties such as force and topography (Nevin, 1992). Although DRL responding occurs at a low rate, it is easily produced and maintained; and although DRH responding occurs at a high rate, it is hard to produce and fragile once it has been produced.

Other classes of differential-reinforcement schedules set other temporal criteria. The *differential reinforcement of paced responding* sets both upper and lower limits on the IRTs that can precede reinforced responses and tends to maintain a fairly constant response rate. The *differential reinforcement of other behavior* (usually abbreviated *DRO*, and occasionally also called *differential reinforcement of zero behavior*), delivers a reinforcer if a specified time elapses without a response. This is in fact the technical name for a procedure often used to reinforce alternative responses during extinction in applied settings (cf. Chapter 5).

Consider a DRO schedule of 10 s arranged for a pigeon's pecks. This schedule reinforces the class of responses defined by 10 s without pecking. But if a reinforcer is delivered after every 10 s without pecking, then each peck must delay the reinforcer by at least 10 s. If pecks occur less often because they delay reinforcers, it would be appropriate to say they're negatively punished. In this case, as in distinguishing positive reinforcement of not-responding and negative punishment of responding, speaking in terms of specified responses may seem preferable to speaking in terms of their absence. But the DRO vocabulary is well established. Though potentially troublesome it defines a procedure and should not be taken as attributing functional properties to not-responding as a response class.

RESPONSE SEQUENCES: CHAINS VERSUS CHUNKS

Once we break down a behavior sequence into components, we can treat the sequence as a succession of different operants, each defined by the reinforcing consequence of producing an opportunity to engage in the next one until the sequence ends with a reinforcer. Such sequences are called response *chains*. An example was provided in Chapter 6 by Sidman's detailed analysis of a rat's lever pressing reinforced by food. Rising up to the lever produced contact with the lever, which set the occasion for pressing the lever, which produced a seen food pellet, which set the occasion for moving to the food cup, and so on. Any segment of the sequence serves the dual function of reinforcing the last response and producing the conditions that occasion the next one. A discriminative stimulus that serves such a reinforcing function is sometimes called a *conditional reinforcer* (cf. Chapter 14).

Chained sequences are inevitable products of some environments. If I'm on my way up to my campus office, I must open the door of the building before I can enter; when I get to the elevator, I must press the elevator button unless the elevator car is there already; I can enter only when its door opens; once inside I press the button for my floor; I leave only when the elevator door opens on my floor; and so on. Each new stimulus allows the next response of the chain, and some of those responses cannot occur without them: I can't pass through a door I haven't opened yet. I can alter the sequence in some ways, as by taking the stairs instead, but I can't much change the order of the components of the sequence. That order is determined by the environment.

Some behavior sequences can be reduced to smaller units in this way, and the analysis into such components can be confirmed experimentally by checking how independent the components are from one another (Skinner, 1934a). For example, if lever presses no longer produce food, lever pressing may decrease, but by delivering food independently of lever pressing we may discover that food continues to occasion movement to the food cup. This demonstrates that the integrity of one component of the chain isn't affected by altering the reinforcement contingencies for another.

But not every sequence can be broken down this way. For example, one experiment (Straub *et al.*, 1979) arranged four keys lit green, white, red and blue in a pigeon chamber, with the location of each color varying from trial to trial. If the pigeon pecked green, white, red and blue in that order, the final peck on blue produced food; the trial was cancelled if the keys were pecked out of order. The pigeon learned to peck white after green and red after white and blue after red even though the color locations changed from trial to trial and even though successive pecks in the sequence didn't produce stimulus changes. Furthermore, the pigeon's pecks were likely to conform to the reinforced sequence even with one of the colors absent (e.g., after pecking white, pecking blue rather than green if red was missing). The current response in the sequence must depend on the pigeon's past behavior. The same issue arises whenever an organism learns an arbitrary sequence in which responses are not accompanied by stimulus changes (e.g., Boren & Devine, 1968); can the organism's own behavior provide discriminative stimuli that may occasion its later behavior?

Let's get back to the trip to my office. If my office door has a combination lock, I punch in my five-digit code and enter once the door is unlocked. Do I have to wait for some consequence of entering the first digit before entering the next, and then again for the third one, and so on? Is it appropriate to call this sequence a chain? Of course the lock won't open if I enter the digits in the wrong order, but where are the discriminative stimuli and the consequences that hold the sequence together? The question isn't whether some behavior sequences are held together as chains. Rather, it is whether chains are the only way in which to put organized behavior sequences together.

In the history of the psychology of learning, the positions taken on the nature of behavior sequences were often symptomatic of serious divisions among researchers of different orientations, some holding that sequential behavior could always be interpreted in terms of an ordering of independent units (variously called associations, stimulus-response bonds or conditioned reflexes), and others holding that sequential behavior could not be adequately interpreted in such terms. As we'll see, **sequential behavior of both sorts is possible**; the crucial question in any particular case is determining the nature of the sequence.

Perhaps the most telling argument was made by Lashley (1951), who summarized earlier conceptions as follows:

> the only strictly physiological theory that has been explicitly formulated to account for temporal integration is that which postulates chains of reflexes, in which the performance of each element of the series provides excitation of the next.... The...kinesthetic impulses from each movement serve as a unique stimulus for the next in the series. (Lashley, 1951, p. 114)

Lashley then described cases, including illustrations from language and music, that made such accounts implausible. For example, in answering the argument that each movement serves as a unique stimulus for the next, Lashley considered the complex sequence of movements required to pronounce the sounds of the word "right" in proper order. The order isn't given by the sounds themselves, because they can occur in a variety of orders and combinations (e.g., in the opposite order, as in "tire"). If the sound sequence cannot be based on direct connections it must depend on some larger organization. Lashley extended the case from sequences of sounds within words to sequences of words within sentences:

> The word "right," for example, is a noun, adjective, adverb, and verb, and has four spellings and at least ten meanings. In such a sentence as "The millwright on my right thinks it right that some conventional rite should symbolize the right of every man to write as he pleases," word arrangement is obviously not due to any direct associations of the word "right" itself with other words, but to meanings which are determined by some broader relations. (Lashley, 1951, pp. 115–116)

Lashley used music to point out that the sheer rapidity of some sequences constrains how they might be generated:

> The finger strokes of a musician may reach sixteen per second in passages which call for a definite and changing order of successive finger movements. The succession of movements is too quick even for visual reaction time. In rapid sight reading it is impossible to read the individual notes of an arpeggio. The notes must be seen in groups.... Sensory control of movement seems to be ruled out in such acts. (Lashley, 1951, p. 123)

Some sequential patterns cannot be reduced to a succession of stimulus-response or S-R units. When a skilled typist rapidly types *the*, each letter can't be a discriminative stimulus for the next stroke, first because the typist will be executing that next stroke even before last typed letter exists yet on the page; and second because as discriminative stimuli these letters can't be unique if they can be followed by hitting the space bar or varied other keys depending on what the typist is typing (e.g., *the* or *these* or *then* or *thermometer*).

Lashley provided no particular name for such sequences. So far here I've called them *temporally extended units*, but a more concise term, *chunk*, has some precedent in the cognitive literature (see Chapters 26 and 28). Sequences clearly built up as successive links in which each response produces a discriminative stimulus that occasions the next response will be called *chains*. Those consisting of integrated components the links of which cannot be broken down in that way will be called *chunks*. Sometimes chunks may be embedded within chains, as when opening a combination lock is one integrated link within an extended chain.

In the face of arguments such as Lashley's, some researchers felt a choice had been forced between assuming that all sequential behavior depended upon stimulus-response sequences or assuming that all depended on temporally extended units of behavior not reducible to such sequences. The issue may instead be regarded as an experimental one. Some sequences clearly can be put together so that each response produces stimulus conditions that set the occasion for the next (chaining), whereas others must be integrated so that responses appear in the proper order without each depending on the consequences of the last (chunking). **For any given sequence, the issue is deciding whether it is made up of chains or chunks or some combination.**

Mediating Behavior

When we observe sequences, another issue is the function of the responses we see in the sequence. Are they determined by the contingencies or are they simply accidental fillers or do they fall somewhere in between? When responses make sequential contingencies more likely to be met even though they are not required by the contingencies, they are sometimes referred to as instances of mediated generalization (Keller & Schoenfeld, 1950). If a DRL schedule requires a long pause followed by a lever press, a rat that moves around exploring the chamber before pressing may be more likely to wait long enough than one that stays right by the lever (Laties *et al.*, 1965), and both the exploratory behavior and the lever press may be affected by the reinforcer that follows.

Suppose a pigeon's task is to match the color of some sample stimulus after a delay. A center sample key might come up green or red. Either way, the pigeon pecks that sample and it goes dark. Several seconds later two side keys light up, one red and one green. If the pigeon pecks the one that matches the sample, its peck will be reinforced. What if the pigeon leans left after a green sample and if later leaning left always pecks green again, but leans right after a red sample and if later leaning right always pecks red again. The pigeon's matches depended only indirectly on the colors; instead, they were supported or mediated by which way the pigeon leaned.

In research settings mediating behavior is typically unplanned and its function is determined only by experimental analysis. But natural environments provide many opportunities for establishing such sequences, especially where situations require waiting. If mediating behavior during a wait extends the time and therefore changes the likelihood that some response at the end of the wait will be reinforced, that behavior may be properly regarded as a product of the contingencies.

VARIABLE BEHAVIOR

The close correspondence between the class of responses with consequences and the class of responses generated by these consequences is the criterion for speaking of an operant class. As we've seen, these classes may be defined along single dimensions such as force or location or may have more complex properties, as in differential-reinforcement schedules and response sequences. In each case, our major interest is with the dimensions along which responding conforms to the class of responses that is reinforced.

The structure of behavior is such that we can't always define such dimensions independently of reinforcement contingencies. For example, consider the reinforcement of novel responses in porpoises (Pryor, Haag, & O'Reilly, 1969). The novel performances were shaped by reinforcing, in each session, some class of responses not reinforced in any previous session. For example, if a porpoise's backflips were reinforced in one session, slapping the water with its tail might be reinforced in the next and beaching itself at the side of the pool in the one after. After several sessions, the porpoise began to emit responses in each new session, such as leaping up from the water with a corkscrew spin, that the experimenters had never seen before. Response novelty had been differentiated, but how else is this operant class to be specified except by describing the criteria for reinforcement? Novel behavior must be emitted before it can be incorporated into other behavior, but the fact that we have difficulty measuring it doesn't rule out novelty or other complex dimensions of behavior as properties that can define operant classes.

Another dimension that cannot be defined independently of reinforcement contingencies is the variability of responding, which has also been a target for differential reinforcement (Neuringer, 1986, 2002). One experiment (Schwartz, 1982a), apparently intended to be critical of the concept of reinforcement itself, had been explicitly designed to show that reinforcers produce stereotypy rather than variability. It required the pigeon to peck on two keys to drive a light from one to another corner of a four-by-four array. Pigeons came to do this in a stereotyped way, but the working of the array constrained possible sequences of left and right pecks. An analogous procedure that allowed all possible sequences (Page & Neuringer, 1985) showed that reinforcers could in fact maintain highly variable behavior, and these results led to a research program that demonstrated that variable responding was a dimension of behavior that could be reinforced (Neuringer, 2004). A curious footnote to this history is that the original experiment, from which flawed conclusions were drawn, may have been undertaken on the basis of a political rather than a scientific agenda (cf. Catania, 1987; Hooper, 2002; Schwartz, 1986).

But differential reinforcement with respect to novelty or variability raises questions. Reinforcers are produced by individual responses, and yet properties such as novelty and variability cannot be properties of individual responses. They can only be properties of responses in the context of other responses that occurred earlier. A given response might be variable in the context of one sequence of past responses and stereotyped in the context of another. Thus, the fact that novelty and variability can be differentially reinforced means that organisms are sensitive to populations of responses and consequences over extended periods of time, and not merely to individual response-stimulus sequences (cf. Chapter 8 on molar and molecular analyses).

We'll return to the issue of behavior classes later, especially in connection with verbal behavior, when we'll treat the finding that some grammatical structures are more easily learned than others as similar in kind to the finding that rats learn some maze sequences more easily than others. In each case, the problem is to identify the dimensions along which responding may come to conform to the class of responses that has consequences. Like the novelty and variability explored here, these dimensions can sometimes be specified only by verbal description. In Chapter 13, we will return to these dimensions in a review of the various sources of novel behavior.

ADDENDUM 10A: THE SHAPING OF PHYSIOLOGICAL RESPONSES

Early in the psychology of learning, when operant behavior was called *instrumental*, it was assumed that only responses of the skeletal musculature could enter into classes modifiable by their consequences. Other classes of responses called *autonomic*, those of glands and smooth muscles, hadn't been shown to be modifiable in this way. Such responses were typically elicited, and Pavlov's conditioning procedures (see Chapters 2 and 17) had shown how new stimuli could come to elicit them. At that time, evidence that these responses could be modified through reinforcement or punishment was negligible. Furthermore, it was always possible to argue that such changes in autonomic responses were mediated by other kinds of behavior (as when exercise raises heart rate).

Despite controversy (e.g., Crider, Schwartz, & Shnidman, 1969, on pupillary contraction; Dworkin & Miller, 1986), some autonomic responses seem modifiable by their consequences. Consider the salivary response. Salivation occurs spontaneously as well as when elicited by a stimulus such as food in the mouth (Zener & McCurdy, 1939; cf. Chapter 4). Consequences can be arranged for spontaneous or emitted salivation, measured in drops of saliva. The consequences of emitted salivation mustn't be food, though, because the effect of food as a reinforcer would be hard to distinguish from its effect as an elicitor of salivation. Thus water, which doesn't elicit salivation, was used to reinforce salivation in water-deprived dogs (Miller & Carmona, 1967). Salivation increased when it produced water (reinforcement) and decreased when it prevented the delivery of water (punishment). Barring some procedural artifact, these autonomic salivary responses seem to have been modified by their consequences; if so, emitted salivation could appropriately be called an operant class, and these salivary responses were members of that class.

The possibility of mediation by other behavior makes this a challenging research area (cf. Harris & Turkkan, 1981, on shaping of blood pressure elevation). I can modify my salivation by working my tongue and other mouth parts; I can modify my heart rate or other physiological measures by heavy breathing or by holding my breath; and so on. Even if such potential mediating responses are monitored over the course of these procedures, there is always the chance that some other candidate for mediation has been missed. Examples from such sources as biofeedback and control of physiological processes in Eastern meditation may tempt us to assume that such processes are modifiable by reinforcement contingencies, but perhaps it is wise to remain skeptical until appropriate data are in.

ADDENDUM 10B: MAZE LEARNING

A rat's negotiation of a maze can be regarded as a single complex response. Consider the simplest maze, the T-maze (C in Figure 5-2). With the left goalbox empty and the right goalbox containing food, a sequence that includes running from startbox to choicepoint, turning right, and then running to the right goalbox will be reinforced. Sequences that include turning left won't be reinforced or (if we allow the rat to retrace its way from left to right goalbox) at least won't lead to the reinforcer as quickly. As the rat comes to choose the right path more often we may say that this T-maze performance has become differentiated.

But maze learning can involve more complex sequences. Consider the sequence of responses in the *double-alternation problem* (e.g., Hunter, 1928). At one end of a central runway, a choicepoint offers a left turn (L) and a right turn (R), but both paths lead back to the other end of the runway. Without being removed from the apparatus the organism comes back up the central runway and again chooses left or right at the choicepoint. If it makes a series of left and right turns in the appropriate order, food is presented to it as it returns to the central runway.

In such apparatuses, the question was whether organisms can learn double-alternation sequences, either LLRR or RRLL. The sequence was learned with varying degrees of success by different species and different organisms of the same species, but a more general issue was whether an organism's behavior at one time can determine its behavior later. An organism can't complete a double-alternation sequence successfully unless its behavior at the choicepoint somehow takes into account which turns it made the preceding times around.

A difference between this and other examples of differentiation is in how the response class is specified. Linear position or force or duration are single dimensions along which responses can be located, but sequences of responses in double-alternation learning can't be ordered unambiguously along any one dimension. For example, LLLR contains only a single transition from L to R and LRLR contains as many Ls as Rs, but which one is more closely related to the reinforced sequence, LLRR?

As organisms learn response sequences in these tasks, differential reinforcement acts upon the behavior they bring to the experiment. This behavior is often systematic. For example, a rat learning a maze might at first always turn left, then later always right, then still later alternate between left and right, and so on until mastering the entire maze. In a Lashley jumping stand with light and dark stimuli in which choices of dark lead to food and in which light and dark irregularly change sides, a rat at first might mostly choose light, then mostly left, then mostly right, until finally the rat begins consistently choosing only dark. Such systematic patterns were noted by Krechevsky (1932), who called them hypotheses:

> It is either implicitly assumed or explicitly stated in most descriptions of learning that in the beginning of the process the animal is a "chance" animal. His responses are without purpose, without form, and without meaning. (Krechevsky, 1932, p. 157)

Based upon his data, Krechevsky contrasted this with a different view:

> The animal, in executing a series of movements which we call "perfect," "errorless," "learned," "integrated," is not doing something which has arisen from a series of "imperfect," "unintegrated," "chance" responses. He is now merely running through a different set of integrated responses, which series of integrated responses were *preceded by other just as integrated responses*.... When the human individual *behaves* in the very same such and such way we must also say that he has an "hypothesis." However, we are primarily interested not in defending our terminology but in describing certain behavior. (Krechevsky, 1932, pp. 528–529)

Krechevsky here provides a useful reminder that complex organisms do not enter new environments behaving randomly. What has been selected up to that point is a product of nonrandom environments. Again we can find parallels in biology. "Despite the randomness of mutation..., phenotypic variation cannot be random because it involves modification of what already exists" (Kirschner & Gerhart, 2005, p. 220).

Chapter 11
Discriminated Operants: Stimulus Control

The Latin habere, *to have, and* capere, *to take or seize, are traceable to closely related Indo-European roots.* Habere *is an ancestor of* behavior, habit *and* inhibit. Capere *led to* concept *and* perception, *words relevant to stimulus classes; thus, these words and* behavior *are distant relatives. They rejoin in* capability

There are no obvious ties among differentiation *and* induction, *applied to response classes, and* discrimination *and* generalization, *applied to stimulus classes.* Differentiation, *from the Latin* dis-, *apart, plus* ferre, *to carry, is related through the Indo-European* bher-, *to carry or bear, to* birth, transfer, preference *and* metaphor *(but not to* interfere*).* Induction, *from the Latin* in-, *in, plus* ducere, *to lead, is related through the Indo-European* deuk, *to drag or to lead, to* duke, adduction, educate *and* conduct. Discrimination, *from the Latin* dis-, *apart, plus* crimen, *judgment, is related by way of the Indo-European* skeri-, *to cut or separate, to* crime, describe *and* criterion. *And* generalization, *from the Latin* genere, *to produce or cause, is related through the Indo-European* gen-, *to give birth or beget, to* ingenious, kind *and* nature.

When you're a Jet and you're walking toward some gang members down the street, it really matters whether you can discriminate between your own gang and your rivals, the Sharks. The two gangs dress differently and talk differently and act differently, so it is easy to tell them apart. But discriminating isn't just easy; it is inevitable. You couldn't keep from discriminating if you tried. Continuing down the street is likely to be reinforced if you find yourself among other Jets, but it will probably be punished if you end up in the middle of a group of Sharks. This is a straightforward three-term contingency: given the discriminative stimulus Jets, walking down the street is reinforced, but given Sharks it is punished. So too for Sharks who see either other Sharks or only rival Jets down the

KEY TERMS: Discriminated Operant, Stimulus Control, Three-Term Contingency; Discrimination and Generalization; Attention; Fading

street. Under such circumstances, can we imagine that any members of either gang would fail to discriminate on the basis of the stimuli correlated with gang membership?

When we discuss discrimination in the technical vocabulary of behavior analysis, we do so in the context of specifying the conditions under which organisms come to behave one way in the presence of some stimuli and a different way in the presence of others. When we discuss discrimination in everyday talk, however, we usually do so in the context of concerns about how people deal with and treat each other. We ask whether some individuals might be discriminated against on the basis of race or ethnicity or gender or sexual orientation, or whether discriminatory practices within certain educational or professional or social or religious institutions might confer advantages or disadvantages upon some individuals but not others. Sometimes the term has more favorable implications, as when we speak of someone who has discriminating tastes. These issues may seem worlds apart from what makes a pigeon peck a green key but not a red one. Here, however, I'll be arguing for a common thread that ties all of these issues together.

What are my credentials? Well, first of all, I'm an expert in discrimination. By that I mean that I completed and received a grade of A in an undergraduate course, Psychology 3, called *Discrimination*, and I've continued studying and later teaching on that topic ever since. My course, part of an Experimental Psychology sequence, concentrated mainly on vision and hearing and other sensory systems, but its content led directly to the topic of discrimination learning in the field that would eventually come to be called behavior analysis.

More important, I was fortunate in the diversity represented in the environments in which I've lived and worked. I grew up in the Washington Heights area of Manhattan in New York City in an apartment building at the intersection of several ethnic neighborhoods. My classmates in the local elementary school, P.S. 132, though racially homogeneous, included Hispanics, Irish, Scandinavians, Jews, Greeks and Italians, no doubt among a few others. Later, during high school and college, I worked summers in the city hospital system, first at Harlem Hospital and later at the then newly opened Bronx Municipal Hospital Center. I was the only white summer employee at Harlem Hospital, where I worked in the Outpatient Department and then as a ward clerk, substituting for staff members on vacation. In those days, most of the doctors and senior staff were white, but the rest of the staff and virtually all of the patients were negro (or colored; back then, the terms *black* and *African-American* were not yet readily available as alternative designations).

Much later, after my father died, my mother lived alone into her nineties in an apartment building on Main Street in Flushing, New York. Those who lived there and on the neighboring streets represented such a range of racial and ethnic groups that the locals referred to it as the United Nations. When I visited and took the elevator to her fifth floor apartment, it was not unusual for the riders to include people from three or four different continents. Yet despite their vast diversities they all helped each other and looked out for each other. Where differences lead to conflict in so many parts of the world, what made this community work? A treatment of places where diversity worked devoted an entire chapter to Flushing and other areas of Queens (Meyer & Brysac, 2012). We'll come back to this later.

After working in other positions at other places, including as stint as a psychopharmacologist at the Smith, Kline and French Laboratories in Philadelphia and nine years as a faculty member at the University Heights campus of New York University, I arrived at the University of Maryland, Baltimore County, where I've taught since 1973. The UMBC student body reflects the diversity of its regional demographics. An important academic fringe benefit has been the opportunity to learn from international students and from international visitors to my laboratory; their countries of origin included Norway, Brazil, Japan, Spain, Italy, Wales, India, and Russia. I discovered that differences in language or culture can sometimes impact even research directions and experimental styles (e.g., Motokawa, 1989).

For some years UMBC has had as its President Freeman Hrabowski, an African-American. Through most of my teaching I had perhaps too casually introduced the behavior analytic topic of discrimination in my courses by briefly mentioning that I recognized that the term had social significance but that I would be using it mainly in its technical sense. The proportion of minority students in my classes had been growing, however, and it became clear at some point that I should address the issue directly and in more detail. With Freeman Hrabowski's explicit encouragement, I eventually did so. Some of this chapter is a product of that teaching, and its content has been shaped by many discussions with my students and my colleagues.

DISCRIMINATION IN THE VERNACULAR

Despite the apparent range of usages of the term *discrimination*, all have an underlying commonality. Whether we deal with pigeons responding differently to green keys and red ones or teachers responding differently to the mathematical aspirations of the male and female students in their classes or law enforcement officers responding differently to members of the public depending on what they look like or how they talk or what they wear, the issue here is how such discriminated responding might be brought about by contingencies and what can be done if those contingencies produce problematic behavior.

In itself, discriminated responding is neither good nor bad. Depending on what is discriminated we might want to examine it when people try to scout out artistic or musical or athletic talent or when they try to identify people at risk of certain illnesses so as to take appropriate preventive measures, or instead we might want to examine what goes on when discrimination takes the form of bigotry or racial profiling or restricted access to public facilities and resources. This approach differs fundamentally from what we find in the social sciences at large. For example, experiments showing that the hands of a white person change skin resistance more upon seeing a black face than upon seeing a white one are concerned with a by-product of discrimination, but they tell us little about how that by-product was created.

Some kinds of prejudice are based on verbal histories, but that is an issue we must defer for a while. Our treatment of verbal behavior in Chapters 23 will give us some tools for extending our analysis to cases involving substantial verbal components. But many cases of discrimination depend on actual experiences and not just word of mouth. If a member of some racial or ethnic group is mistreated in some way by members of another group, that person will come to discriminate on the basis of group membership as surely as a rat will come to discriminate between a stimulus in the presence of which its lever presses are punished by shock and one in the presence of which they are not. For this reason, it is useful to distinguish between *discrimination*, which is based upon experienced contingencies, and *prejudice* (literally, pre-judgment), which is based not on experienced contingencies but on what one has been told or has indirectly learned in other ways.

However much we may wish it otherwise, contingencies maintain discriminations involving the neighborhoods in which people travel and the individuals with whom they interact. To the Sharks and Jets scenario we could add ones involving Christian Fundamentalists and Atheists, or Israelis and Palestinians, or Sunnis and Shiites and Kurds, or Hutus and Tutsis, or Koreans and Chinese and Japanese, or the Catholics and the Protestants of Northern Ireland, or Serbians and Croatians, or Indians and Pakistanis, or Armenians and Turks—our list could go on and on.

Closer to home we can take pride in how far we have come since slavery, the intelligence typing of European immigrant groups before and during World War I, the displacement of Native Americans from their homelands, the internment of citizens of Japanese ancestry during World War II, and segregated schools and restaurants and toilet facilities recently enough still to be in living memory, to mention just a few historical stains. Discrimination still exists nevertheless. Have you never been discriminated against on the basis of your height or your weight or your gender or your race or your accent or your national origins or your religion or your sexual orientation or your politics or your nerdiness or, if not one of these, something else? Here, though, is a tougher question: Can you honestly say that you yourself have never discriminated against others based on such dimensions?

If existing social and political environments make such discriminations sometimes inevitable, are there things we know about behavior that can help us to cope with them constructively? Conditional discrimination is a topic we'll cover in more detail in Chapter 12. A pigeon has learned to peck green and not red when a chamber light is on but red and not green when the chamber light is off. Whether it pecks green or red is conditional on whether the chamber light is on or off. What can this tell us about our human discriminations? For one thing, though it may be inevitable that some will behave discriminatively based on their experiences with others in particular environments, such as different neighborhoods within a city, generalization of this behavior to other environments doesn't necessarily follow. In some places or situations the failure to discriminate can be costly; the consequences might even make the difference between life or death. In other places or situations, however, different judgments along different dimensions are required.

In a university, for example, what should matter most is not what a student looks like or where the student is from but rather what competencies the student demonstrates. As an extended conditional stimulus, the college should be a place not only where some of the discriminations learned elsewhere are not tolerated but also where they will be replaced by new ones based not on what others look like or where they are from but rather on what they have achieved. Once those kinds of discriminations become more important in those few environments than the ones from other environments that they've replaced, perhaps they will eventually become unconditional rather than merely conditional discriminations. To the extent that some problematic discriminations arise from actual contingencies, it is unlikely that we can eradicate them, either by legislation or by other means. But we can at least nurture environments where those discriminations are replaced by conditional discriminations with which they are incompatible. There is precedent for hoping that enclaves of tolerance can produce the seeds of change (Appiah, 2010; Meyer & Brysac, 2012). We are beginning to learn how contingencies work, so identifying the contingencies that enter into these societal problems may help us to create practical solutions.

THE NATURE OF DISCRIMINATED OPERANTS

In Chapters 9 and 10 we dealt with reinforcement based on response dimensions, but differential reinforcement can also be based on the stimuli in the presence of which responses occur. For example, a rat's lever presses in light are different from its presses in darkness, so reinforcement can be arranged for presses in the presence but not the absence of light. When responding is reinforced only in the presence of some stimulus, we say that reinforcement is *correlated with* that stimulus. For convenience we may occasionally abbreviate *in the presence of a stimulus* to *during a stimulus*. Responding controlled in this way by stimuli is said to be *discriminated*. A response class created by such differential reinforcement with respect to stimulus properties is called a *discriminated operant*.

Discriminated operants are pervasive. When driving, you proceed through an intersection if the traffic light is green but not if it is red. When speaking to someone, what you say is affected by what the other person said, the other's posture and facial expression, the setting of the conversation, and so on. Many earlier examples of reinforcement included discriminative control of responding. In discussing the rat's learning of a maze, we emphasized the increase in the choice of correct turns, but a rat that failed to discriminate among different sections could hardly master the maze as a whole.

In fact, there is probably no such thing as an operant class without discriminative stimuli. A pigeon's key peck can't be emitted in the absence of a key and a rat's lever press can't be emitted in the absence of a lever. Features that remain relatively constant throughout an experiment, such as the chamber itself along with the devices in it, are sometimes called *contextual* stimuli. We'll usually be most interested in stimuli that change within experimental sessions, but a laboratory pigeon's or a rat's environment exists in a broader context, including its living quarters and other features besides the experimental chamber (cf. Donahoe & Palmer, 1994).

Discriminative stimuli correspond to those colloquially called signals or cues. They don't elicit responses. Rather, they set the occasion on which responses have consequences and are said to *occasion* responses (cf. affordance: J. J. Gibson, 1979). Figure 11-1 shows an example (Herrick, Myers, & Korotkin, 1959). Rats pressed a lever in the alternating presence or absence of light. With the light on, presses were occasionally reinforced with food; with it off, they weren't. The notation for the stimulus correlated with reinforcement is S^D for *discriminative stimulus* or S+ for positive stimulus;

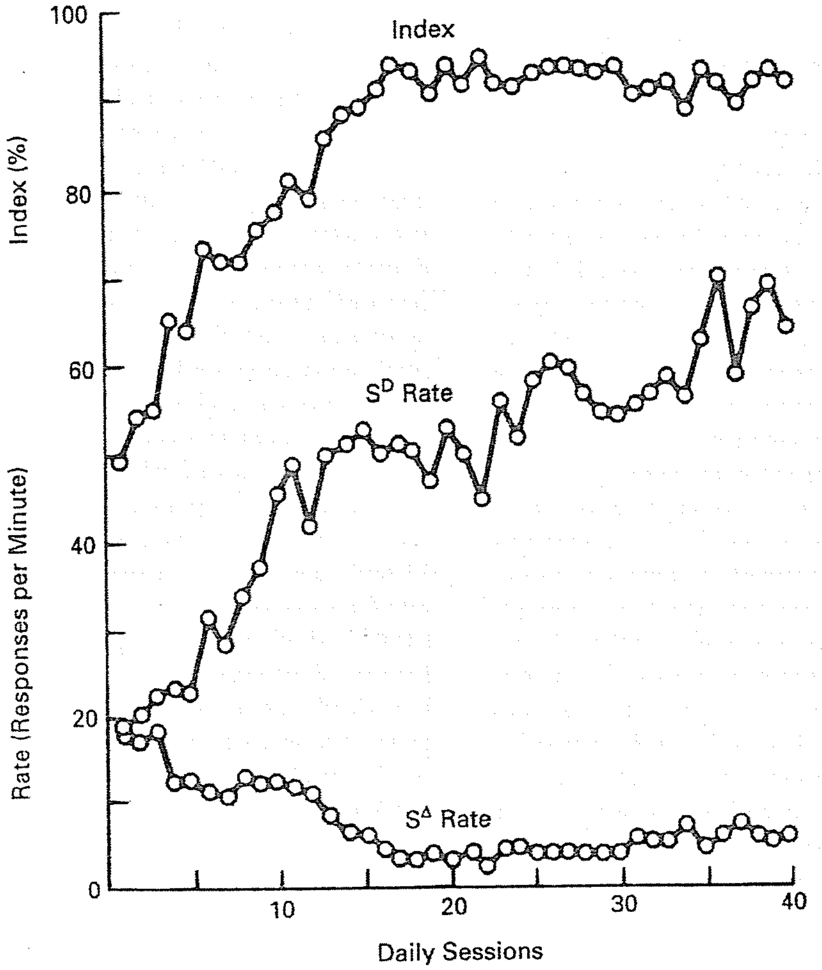

Figure 11-1 Rats' lever pressing in the presence and absence of light. The light (S^D) was correlated with reinforcement and its absence (S^Δ) with extinction. The discrimination index is the percentage of total presses during light: 100 x S^D rate / (S^D + S^Δ rate). Data are medians from eight rats. (From Herrick et al., 1959, Figure 2)

for the one correlated with nonreinforcement or extinction it is S^{Δ}, also for discriminative stimulus (Δ is delta, the Greek d), or S– for negative stimulus (strictly, S^0 is more appropriate to denote the absence of a stimulus, but S– is a more common usage).

In this procedure, light and dark alternated irregularly; when lit, the lamp remained on for 5 to 30 s. Lever presses were reinforced according to a *variable-interval* or *VI* schedule: On the average, only one lever press per 30 s was reinforced while the light was on. The important features of this schedule are: (1) it maintains a moderate and relatively constant response rate; and (2) the varying times between reinforcers make time since the last reinforcer an unreliable predictor of when the next response will be reinforced. Stimulus changes and reinforcer deliveries vary unsystematically, so if the rat presses more during light than dark we can be confident that the rat's responding is not based on the temporal regularities of those other events.

Over sessions, lever pressing increased during light and decreased in its absence. Figure 11-1 also shows a discrimination index, responding in light as a percentage of total responding, which increased over sessions. Equivalent ways to describe this outcome are to say that lever pressing during the light was a discriminated operant, or that the light functioned as a discriminative stimulus for lever pressing, or that lever pressing was under the stimulus control of the light.

We can illustrate a discriminated operant with a hypothetical example comparable to that of Figure 9-1. Again, a rat is in a chamber with a slot in one wall, but this time it can't poke its nose into the slot. Instead, the slot is covered by a translucent plastic strip that can be lit from behind by any one of a row of 15 lamps; these lamps provide a stimulus dimension rather than a response dimension. A lever is centered just below the slot and just above a feeder. Now we light the lamps one at a time in irregular order and record the lever pressing that occurs during each. If we don't reinforce presses, pressing is infrequent and bears no systematic relation to which light is lit. The data might look like A in Figure 9-1, with the difference that the x-axis now represents stimulus instead of response position.

Now we reinforce presses only when positions 9, 10, 11 or 12 are lit; with any others lit, we don't reinforce presses. The initial effect of reinforcement is much like B in Figure 9-1: Responding increases across all positions. In other words, the effect spreads. The spread of the effect of reinforcement from one stimulus correlated with reinforcement to nearby stimuli is called *generalization*. Here, reinforcing presses when any light from 9 to 12 is lit affects pressing during light at any of the other positions.

With continued differential reinforcement with respect to stimulus location, pressing gradually increases during stimuli 9 through 12 and decreases during others, as in C, D and E of Figure 9-1. Eventually most pressing occurs when positions correlated with reinforcement are lit, as in E. Though some presses still occur with other positions lit, at this point the distribution of responses doesn't change much with continued differential reinforcement.

In this example, the stimuli that occasioned responding came to conform closely to the class of stimuli correlated with reinforcement. Differential reinforcement created a class defined by the stimuli in the presence of which responses occurred. But what about responses during stimuli outside the boundaries (e.g., positions 6, 7 and 8, or 13, 14 and 15)? Strictly, we shouldn't count these as members of the discriminated operant; we should speak of them in terms of generalization. But differential reinforcement generated the responding both inside and outside those boundaries, so all are parts of a single distribution.

We resolve this problem as we did in Chapter 9. One class is correlated with a reinforcement contingency; the other is the class during which responding occurs. **We are interested not in either class alone but rather in their correspondence.** This parallels our discussion of differentiation and induction, as is appropriate: A stimulus during which a response occurs is another property of that response, like its force or duration or topography. Why then the different terms for differential reinforcement with respect to response properties (differentiation and induction) and with respect to stimulus propmerties (discrimination and generalization)?

The answer probably lies in their separate experimental histories. When we study differential reinforcement with respect to response properties, we record responses in different classes, but aside from arranging contingencies there isn't much else we can do about them. If we see a rat about to poke its nose into position 7 along the slot, we can't stop the rat from emitting that response even if it has responded there a lot more often than elsewhere. But if we are working with stimulus properties, we can choose among many possible orders and relative frequencies of our stimuli. For example, instead of presenting lights equally often at each position we could present lights at some positions but not others, so that the rat gets many opportunities to press during some stimuli but never gets a chance to press during others. This is why we say that stimuli in discrimination procedures set the occasion for responses: When a class of responses is defined by the presence of a stimulus, responses in that class can't occur when the stimulus is absent.

This methodological distinction runs through these two research lines, but even it has limitations. For example, we could as easily discuss the differential reinforcement of long interresponse times (the DRL schedule: Chapter 9) in terms of the differentiation of a complex operant consisting of a pause plus a peck as in terms of a discrimination based on time elapsed since the last peck, with the duration of the pause treated as a stimulus property. In this instance, the vocabularies are interchangeable. **Differential reinforcement underlies discrimination and generalization just as it underlies differentiation and induction. Both involve correspondences between the dimensions upon which differential reinforcement is based and the dimensions of the resulting behavior.** But with unambiguous cases we'll adhere to the established distinction between these vocabularies, because they have an extensive and widely accepted historical foundation.

SIGNAL DETECTION: BREAST SELF-EXAMINATION

A stimulus may be below *threshold* or too weak to be detected. Two stimuli may be so close together they cannot be told apart; their separation may not exceed their *difference threshold*. But a threshold is not some absolute value. It may vary from moment to moment and it is best regarded as a distribution: the probability of detection as a function of stimulus value. The mathematics of these probability distributions provided the basis for signal detection theory (Green & Swets, 1966; Swets, Dawes, & Monahan, 2000), which we applied to shaping in Chapter 9.

Sensory capacities may have substantial effects in learning procedures, and overlooking them can create problems with experimental designs and the interpretation of research findings. In experiments with rats that use auditory stimuli, for example, it helps to know that they can hear high frequency sounds that we cannot and that they will probably have difficulty with low frequency sounds that are easy for us to hear. What might seem slow learning might be telling us more about how the rat's hearing differs from ours than about whether it is a good learner. In experiments with pigeons that use visual stimuli in different locations, it helps to know that a pigeon is essentially near-sighted with respect to stimuli in front of its beak but far-sighted with respect to stimuli on one or the other side of its head. In experiments with squirrel monkeys using colors as visual stimuli, it helps to know that some subspecies are color-blind whereas others are not, so that a simple color discrimination that is easy for one may be impossible for the other. Issues like these become even more crucial when a child's behavior problems addressed in terms of environmental contingencies turn out to depend on instead on some sensory deficit such as hearing loss or uncorrected vision.

Table 9-1 presented signal-detection contingencies, and it is easy enough to translate from that shaping scenario to one for a pigeon pecks reinforced during green but not during red (Nevin, 1969). In that case, pecks on green are hits and failures to peck are misses, whereas pecks on red are false alarms and failures to peck are correct rejections. But let's instead consider a more important human discrimination, breast self-examination, where the response is the palpation of different areas of the breast and the relevant signal is a potentially malignant node or lump. If there is such a lump, a hit is obviously better than a miss; if not, a correct rejection is obviously better than a false alarm. But both misses and false alarms can be costly, though in different ways, and those relative costs will bias judgments in one direction or the other. Misses may seem inherently more important than false alarms, but consider what it might be like if someone free of cancer was mistakenly told that a biopsy had tested positive.

Breast self-examination has been studied extensively as a discriminative skill that can be taught (Pennypacker & Iwata, 1990). Early studies emphasized the detectability of the lump as a function of various properties such as its size or its depth within simulated tissue (Bloom et al., 1982). As the work progressed, however, search patterns and methods of palpation emerged as more significant concerns. Recommendations from the American Cancer Society at the time advocated a spiral pattern of search, but analyses of detection success and area covered demonstrated the superiority of a grid pattern. Based on data such as those presented in Figure 11-2, that search pattern was eventually adopted as the primary recommendation of the American Cancer Society. It has undoubtedly saved lives. It is probably also worth noting that men too are at risk of breast cancer, though at a lower rate than women. But their tumors are typically detected at later stages of their cancers, perhaps because they have not been taught to do their own examinations.

Breast self-examination is relevant here not only because it is an example of how behavior analysts have extended research with nonhuman organisms to areas of significant human concern (Pennypacker is one of the co-authors of the pigeon research presented below in Figure 11-5). It also reminds us that despite our privileged access to our own bodies we need to be taught much of what we know about ourselves. We'll examine this point further in Chapter 22 when we consider how we learn to talk about private events such as our thoughts and our feelings.

Perhaps even more important, this example emphasizes discrimination learning as behavior rather than as a passive outcome of exposure to stimuli. Effective breast self-examination depends on patterns of search and palpation. It is obvious that we can feel things only when we touch them, but we also see things when we look at them and hear things when we listen to them. In any modality, shifting attention from one thing to another is behavior. Put simply, as this example clearly illustrates, **discriminating is something we do.** It is too easy to focus mainly on the stimuli rather than on the behavior occasioned by them.

ATTENDING TO PROPERTIES OF STIMULI

Stimuli have varied properties and there are no guarantees that an organism will respond to just those properties we select. In differentiation, a rat might reliably press with its left paw even though this property isn't critical to whether the press is reinforced. Similarly, in discrimination a rat might respond on the basis of the intensity of a visual stimulus even though differential reinforcement is based only on its shape; we might then say that the rat is attending to intensity. (Stimulus properties to which an organism is likely to respond are sometimes called *salient*, but salience isn't a property of a stimulus; it is actually a property of the organism's behavior with respect to that stimulus.)

Organisms typically respond to some stimulus properties and not others, so we can hardly talk about discriminated operants without also talking about *attention*. Furthermore, if attending to one stimulus property or another is something organisms do, we must treat attending as behavior. And if attending is behavior, then it can have consequences. For example, if a reinforcement contingency is correlated with the brightness of a visual stimulus but not with its size, it will make a difference whether the organism attends to brightness or to size, and if attending to brightness occurs more often because of its consequences, it is appropriate to talk about that attending as an operant.

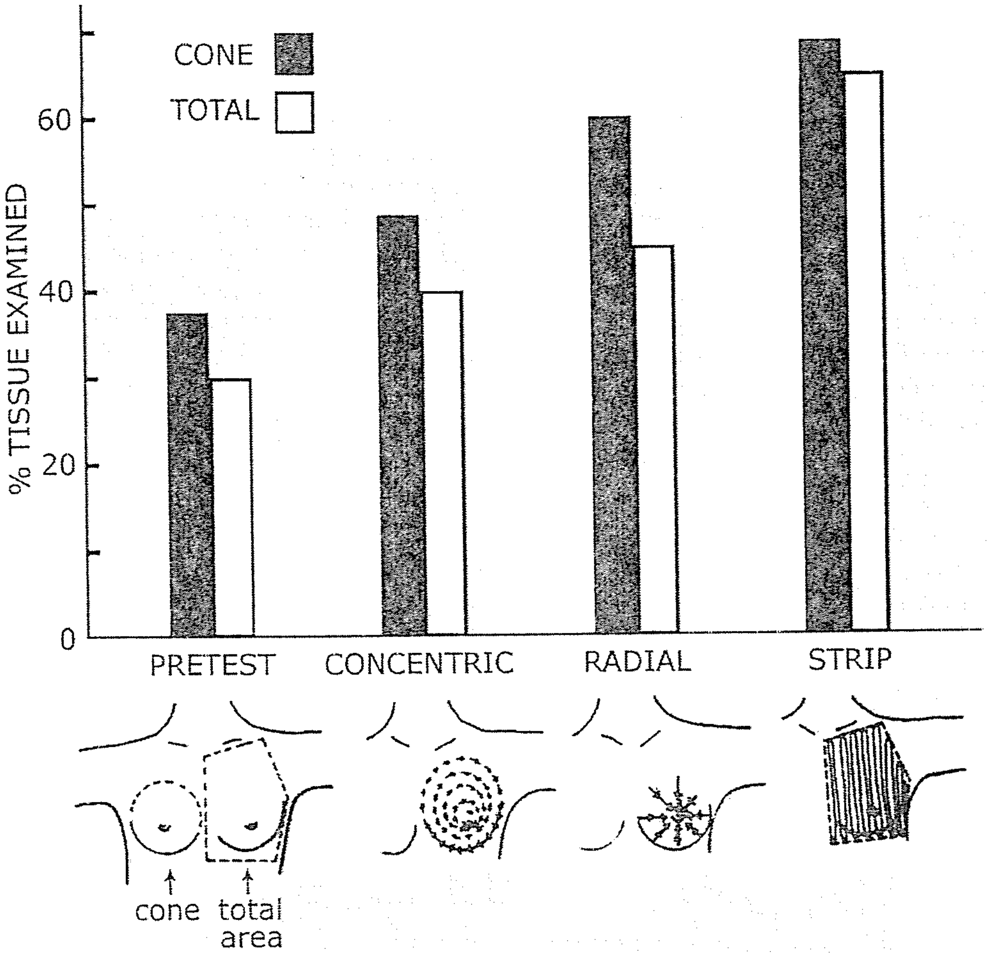

Figure 11-2 Percent tissue examined in pretest breast self-examination (left) and after training with the search patterns illustrated by the sketches beneath the graph. Training helps in all cases, but more tissue is palpated with strip or grid search (right) than with either concentric or radial search. The data show percentages relative to the areas outlined on the left: cone or total recommended area of examination. (Adapted from Figures 13.10 and 13.11 in Pennypacker and Iwata, 1990).

Imagine I'm working on a jigsaw puzzle and I'm searching for a piece. I might scan the unplaced pieces for shape or for image, but it's hard for me to do so for both. Mostly I attend to one feature at a time. I might move and turn pieces or try their fit, but the difference between searching for a shape or searching for an image doesn't reveal itself in the movements of my eyes or my fingers. **Attending is not defined by movement; it is defined in terms of its consequences.** For example, if I'm working on a puzzle in which the pieces have been cut with very similar shapes but the images vary a lot from one part of the puzzle to another, I'll attend more to image than form, but with a puzzle with very differently cut pieces and very subtle differences in image, I'll attend more to form than to image.

Now let's look at some data on attention with pigeons whose key pecks are occasionally reinforced with food. One of two stimulus combinations is presented on the key: a triangle on a red background or a circle on a green background. After 3 min of triangle-on-red, the next peck during this stimulus is reinforced; after 3 min of circle-on-green, the stimulus turns off without reinforcement. The arrangement during triangle-on-red is called a *fixed-interval* or *FI* schedule of reinforcement; the arrangement during circle-on-green is extinction. An FI schedule usually maintains responding that increases in rate as time passes in the interval instead of the relatively constant rate maintained by a variable-interval schedule (VI; see Chapter 15 for more details). With FI reinforcement no peck is reinforced until the interval ends, so we don't have to worry that the pigeon is responding because reinforcers are being delivered rather than because of the stimuli on the keys.

Figure 11-3 shows data for two pigeons (Reynolds, 1961a). The left graphs show rates of pecking during each stimulus combination after 18 hours of training. Both pigeons emitted more than 40 pecks/min during triangle-on-red but pecked much more slowly during circle-on-green. In a test without reinforcement, the elements of each combination were presented separately. Almost all of Pigeon 105's pecking occurred during triangle; red, the color correlated with reinforcement, occasioned little more responding than either circle or green, the elements previously correlated with extinction. On the other hand, almost all of Pigeon 107's pecking occurred during red; even though triangle had been correlated with reinforcement during training, by itself it occasioned even less pecking than circle or green. Pigeon 105 was attending to form and not color and Pigeon 107 was attending to color and not form. Form and color had been similarly correlated with reinforcement during training, but we saw what their discriminative functions were only by separately examining the effects of each element.

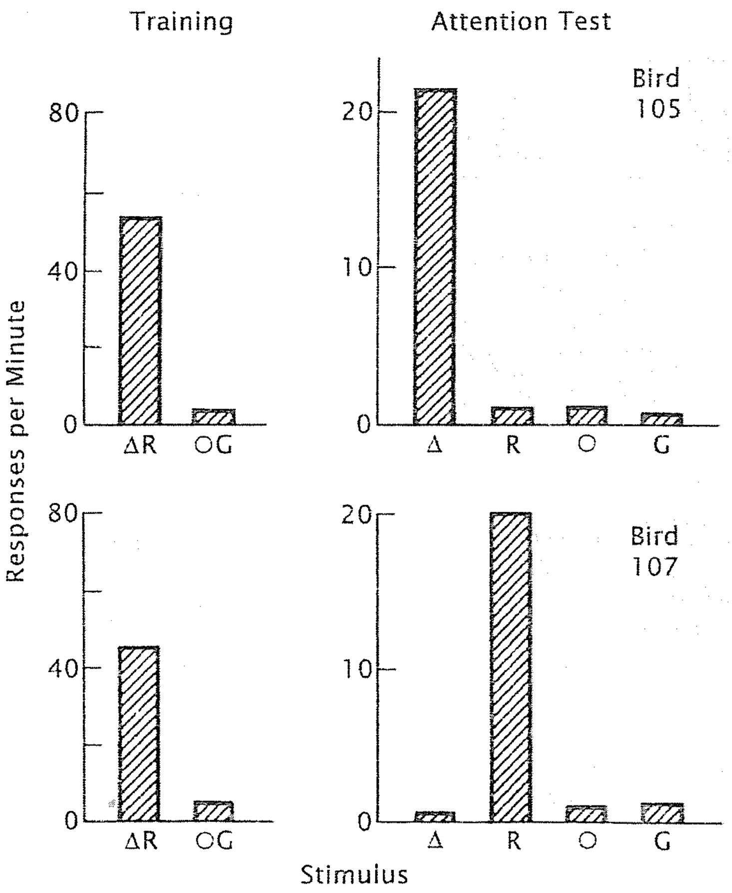

Figure 11-3 Key pecking of two pigeons when reinforcement is correlated with triangle-on-red and extinction is correlated with circle-on-green (left: training), and during extinction tests with forms and colors presented separately (right: attention tests). Pigeon 105 was attending to form but not color; Pigeon 107 was attending to color but not form. (From Reynolds, 1961a, Figure 1)

The pigeons' responding here isn't mere generalization. The responding of Pigeon 105 generalized from triangle-on-red to triangle without red but not to red without triangle. We speak of attention not as a response to some stimulus but rather as attention to some stimulus dimension: Paying careful attention means listening to everything that is said and not just to some parts. We say Pigeon 105 attended to form and not color because discriminated responding occurred with changes in form but not with changes in color. When an organism discriminates along one stimulus dimension but not another, we say the organism is attending to the first dimension but not the second.

Once an organism has attended to some stimulus properties in one situation, it is likely to attend to those properties in new situations. We can change the likelihood that an organism will attend to stimulus properties simply by changing the way in which reinforcement is correlated with them (Johnson & Cumming, 1968). In the training phase of Figure 11-3, form and color were correlated equally well with reinforcement. For this reason, the experiment was particularly suitable for demonstrating some properties of attention. But if our major interest was form discrimination in the pigeon, we'd make color irrelevant and correlate changes in reinforcement only with changes in form.

Attention is crucial in development. In *joint attention*, a parent's routine interactions with a child will ordinarily shape the child's looking so that instead of looking at the parent's face when the parent speaks the child begins to look where the parent is looking or pointing. As we will see in Chapter 22, their joint attention toward events in the environment may be a vital prerequisite for other important skills, such as learning the names of things pointed to (e.g., Dube et al., 2004; Tomasello & Farrar, 1986).

Generalization Gradients

If responding is reinforced during some stimulus and a property of that stimulus is varied, responding may depend on how much the stimulus has changed. For example, if a pigeon's key pecks are reinforced when the key is yellow during training, the pigeon will ordinarily peck at lower rates as the light is changed to orange and red and violet during a test. This demonstrates generalization: the effect of reinforcement during yellow spreads to the other colors.

Figure 11-4 illustrates generalization of pigeons' key pecks to tones of different frequency after pecks were reinforced only during a tone of 1000 cycles/s (Jenkins & Harrison, 1960). With no discrimination training, tone was always present and pecks were reinforced according to a variable-interval or VI schedule. In presence versus absence training, tone was sometimes present and sometimes absent and pecks were reinforced according to the VI schedule only during tone. After training, reinforcement was discontinued and tones of other frequencies were presented for the first time, eight times each in mixed order, along with no tone and the original training tone. (Another type of gradient, the postdiscrimination gradient, is discussed in Addendum 11A.)

Without discrimination training (top), neither tone frequency nor its presence or absence had much effect on pecking. The gradient was relatively flat, so these pigeons weren't attending to tone; changes in its frequency made no difference to them. With discrimination training

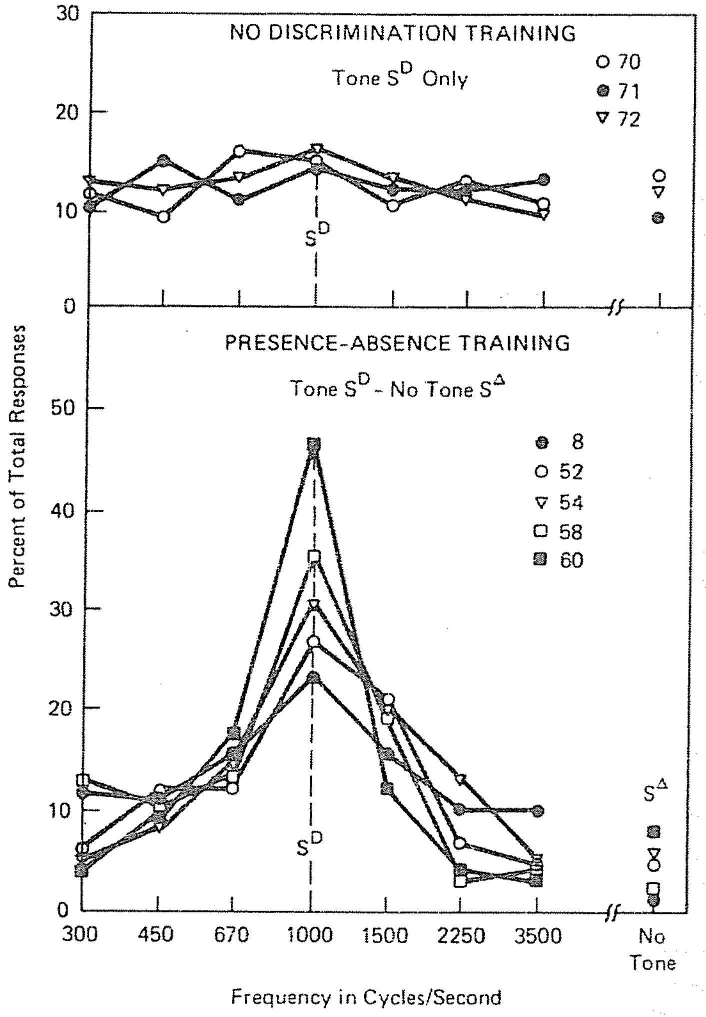

Figure 11-4 Generalization gradients as a function of tone frequency after reinforcement of key pecks during a 1000-cycle/s tone (top: 3 pigeons) or after reinforcement during tone and extinction during its absence (bottom: 5 pigeons). Without discrimination training, gradients were relatively flat; after presence versus absence training, they peaked at the reinforcement stimulus (S^D). (From Jenkins & Harrison, 1960, Figures 1 and 2)

(bottom), the original frequency produced more pecking than any other during the test; in general, the closer to the original frequency, the more the pecking. Low rates of pecking occurred when the tone was absent. No surprise: Absence of tone was correlated with extinction. For these pigeons response rate varied with frequency even though discriminated responding depended only on the presence or absence of tone and not on its frequency: These pigeons were attending to the tone.

Such experiments must be conducted with special care. The form of generalization gradients is affected by other variables, such as the reinforcement schedule during training (e.g., Hearst, Koresko, & Poppen, 1964). Furthermore, any major procedural change can disrupt performance, and the transition from reinforcement during training to extinction during generalization testing is undoubtedly a major change. We want to keep behavior going while we present varied stimuli in our extinction tests. But when contingencies change in natural environments (e.g., some new food with toxic effects has displaced an earlier food), organisms that are choosy and quickly stop responding after a change may have advantages over those that are persistent. The phylogenic effects of such natural contingencies may influence the outcomes of our tests (cf. Catania, Ono, & de Souza, 2000).

Inhibitory Gradients

If the effects of reinforcement spread, how about the effects of extinction? There is a methodological difficulty. To determine whether a stimulus reduces responding, there must be some responding to start with. A procedure is needed that separates the dimension along which the extinction gradient is determined from the one correlated with reinforcement. Such a procedure is illustrated in Figure 11-5 (Honig et al., 1963). With one group of pigeons, the reinforcement stimulus was a vertical line on the key and the extinction stimulus was a lit key with no line; with a second group, these stimuli were reversed. For both groups, VI schedules were used during training and gradients along the dimension of line orientation were obtained during nonreinforcement.

With vertical correlated with reinforcement and its absence correlated with extinction (filled

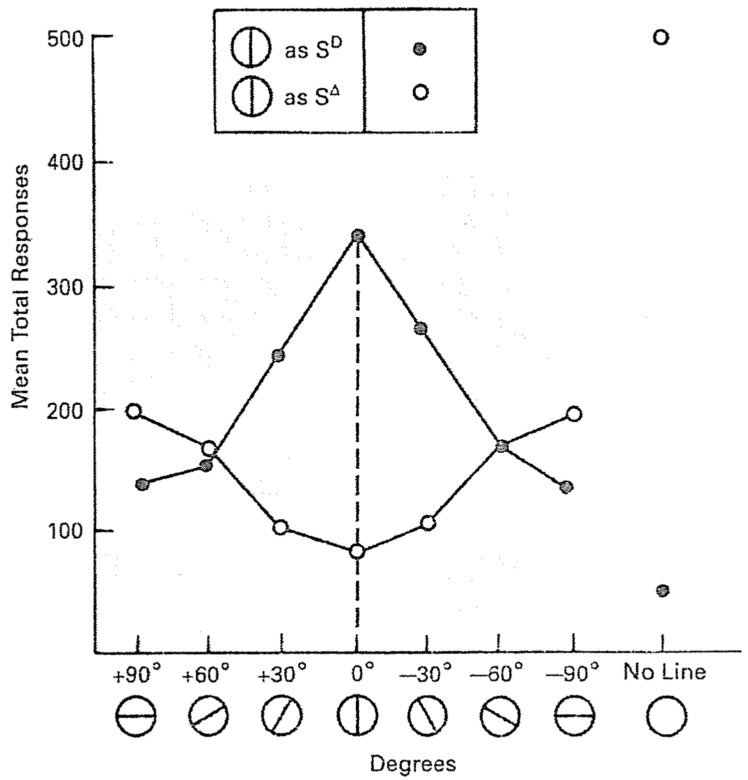

Figure 11-5 Excitatory and inhibitory stimulus-control gradients after VI reinforcement of pigeons' key pecks during one stimulus (S^D) and extinction during another ($S^Δ$). For one group, a vertical line was correlated with reinforcement and its absence with extinction (filled circles); for another, the stimuli were reversed (unfilled circles). Data obtained during nonreinforcement show responding during various line orientations and during absence of the line. (From Honig et al., 1963, Figure 1)

circles), pecking decreased with greater deviations from vertical; this is a reinforcement gradient like that of Figure 11-4 (bottom). But with vertical correlated with extinction and its absence correlated with reinforcement (unfilled circles), pecking increased with greater deviations from vertical. In this group, changing the line orientation of the line affected its distance from the vertical extinction stimulus but not its distance from the reinforcement stimulus, the absence of a line. This is an inhibitory gradient. As is typically the case, this gradient was shallower than the corresponding excitatory gradients, probably because organisms are more likely to attend to stimulus properties correlated with reinforcement than to those correlated with extinction (Dinsmoor, 1995).

Feature-Positive Discriminations

Suppose the alternating stimuli of a reinforcement-extinction discrimination consist of an array of squares with an asterisk substituted for one of the squares or the same array with no asterisk, as illustrated in Figure 11-6. The asterisk is the distinctive feature of these stimuli, and a pigeon will acquire this discrimination far more quickly if the asterisk is correlated with reinforcement (feature positive) than if it is correlated with extinction (feature negative). The pigeon is much more likely to look at the asterisk in the former case than in the latter (cf. Jenkins & Sainsbury, 1970), and as we'll see later, human behavior can be similarly captured by such contingencies.

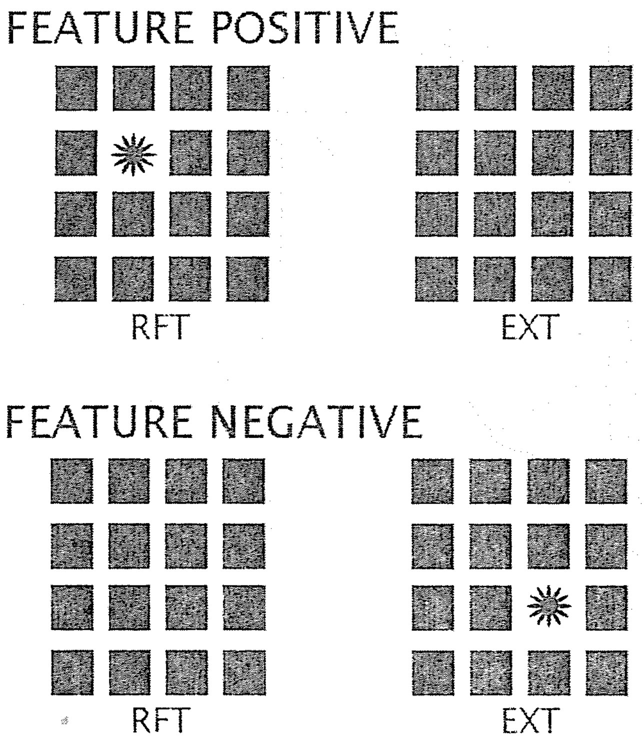

Figure 11-6 In a feature-positive discrimination, responding is reinforced only if a distinctive stimulus is present (top). In a feature-negative discrimination, responding is reinforced only if that stimulus is absent (bottom). The asterisk is the distinctive stimulus and appears in different places in the area of squares over successive trials. With pigeons, feature-positive discriminations are easily learned whereas feature-negative discriminations are learned with difficulty.

The distinctive stimulus appears in different places in the array over successive trials. As the pigeon masters a feature-positive discrimination, it begins to look at the distinctive stimulus and to peck the key by pecking that stimulus wherever in the array it is located. Because it is correlated with reinforcement, the distinctive stimulus becomes a conditional reinforcer, so looking at it is reinforced. With the feature-negative discrimination, however, the distinctive stimulus is correlated only with extinction, so looking at it is not reinforced. If anything, the pigeon looks at the squares, which are correlated with reinforcement during reinforcement trials. The pigeon that fails to look at the asterisk is less likely to learn to discriminate between its presence and its absence, and so learns this discrimination more slowly. Differential attention is important in both nonverbal and verbal behavior, and we'll explore its implications further in other chapters.

Whatever the stimulus, we can always raise questions about the actual features to which an organism attends. For example, suppose a pigeon looks only at the upper edge of a key with a vertical line projected on it. If we rotate the line from vertical, its top is gone from the upper edge of the key, which now looks more like it did when no line was there. In this case, the line is the stimulus in name only: It is a *nominal stimulus*. The *functional stimulus* is what appears at the upper edge of the key. We can evaluate such possibilities experimentally, for example by removing parts of the line to see whether its removal produces the same sorts of changes in the pigeon's behavior as rotating it (Touchette, 1969). The point may seem trivial with regard to a pigeon's pecks but it can be exceedingly important in application, as when a teacher tries to find out whether a child is attending to the words in a storybook or is faking reading by attending mainly to the pictures.

FADING: STIMULUS CONTROL BY SUCCESSIVE APPROXIMATIONS

As usual, there are biological parallels to our behavior analytic categories. Just as a response may be differentially reinforced in the presence of one stimulus but not another, organisms may be selected by evolutionary contingencies in one environment but not in others: "... because species often occupy multiple habitats, a mutation may be favored in some locales and selected against or ignored in others" (Carroll, 2006, p. 162). Fine discriminations can develop with the co-evolution of sensory and motor capacities, as when the sensitivity of a primate mother's hearing is selected based on her responsiveness to the calls of her offspring while the calls are in turn selected by their distinctiveness from calls by the offspring of other mothers (e.g., Rendall, Owren, & Rodman, 1998; Rendall, Rodman, & Emond, 1996). Here both phylogenic and ontogenic selection for discriminative capacities may be intimately intertwined.

Just as the response properties that define an operant can be changed gradually by shaping, the stimulus properties that define discriminated operants can be changed gradually by analogous procedures called fading. Training doesn't have to start with stimuli that are difficult to discriminate. Instead, it can start with stimuli that are easy to discriminate and then move gradually to more difficult ones. For example, it is usually harder to teach a pigeon to discriminate between vertical and horizontal lines than between red and green. Once a discrimination between red and green exists, one between vertical and horizontal can be successively approximated by superimposing vertical on red and horizontal on green and then gradually fading out the colors (Terrace, 1963b). In this chapter I've arranged some crude fading by dropping out *variable-interval* in favor of its abbreviation, *VI*, after both had sometimes appeared together.

Creating stimulus control through fading is often effective (e.g., Sidman & Stoddard, 1967), but as with shaping no rules exist for how rapidly stimuli should be faded in or out in different situations. For example, if we superimpose vertical on red and horizontal on green and then partially fade out the colors, we might remove the colors completely and find that the pigeon has learned to discriminate between vertical and horizontal. On the other hand, fading might be unsuccessful; if a pigeon attends only to colors even when they became very faint, discriminated responding might disappear every time we dim them below some minimum level.

Shaping requires that some behavior is available to be shaped, and fading requires that some discriminative responding is available to be shifted to a new stimulus dimension. For example, consider *errorless discrimination learning* (Terrace, 1963a). Soon after a pigeon's pecks on a red key had been shaped with food reinforcers, reinforcement during 3-min periods of red was continued according to a VI schedule. Periods of red alternated with another stimulus during which pecks weren't reinforced. At first, this other stimulus was a dark key lasting 5 s. Its duration was gradually lengthened and it was changed from dark to dim and then to brighter and brighter green, until to a human eye its brightness matched that of the red key. By the end of these conditions, 3 min of red and reinforcement alternated with 3 min of green and extinction. Each pigeon treated this way pecked the extinction key fewer than 10 times in the entire course of training; pecking occurred almost without exception on red rather than green. Pigeons introduced to this procedure later after initial training or introduced abruptly to green at full duration and intensity pecked the green extinction key hundreds and even thousands of times during equivalent training periods.

The gradual fading in of green was effective in part because turning the red key dark early in training stopped the pigeon's pecking for a few seconds (any abrupt stimulus change might have had such an effect). Pecking wasn't likely to start again before the 5-s extinction stimulus ended. Thus, a difference in responding to the two stimuli occurred at the outset. The gradual changes in the duration and intensity of the extinction stimulus built upon that difference. After such a history, we might convert the extinction stimulus to reinforcement and the pigeon, never pecking that stimulus, might never discover that we had done so. The sense in which such a performance is errorless isn't obvious. We must be cautious about the language of errors; the term error implies a judgment about the value of responding and may be inappropriate to a behavioral analysis. But fading isn't solely of theoretical interest; its possible applications to education give it practical significance.

THE VOCABULARY OF DIFFERENTIAL REINFORCEMENT

Both differentiation and discrimination involve differential reinforcement. The major difference is whether differential reinforcement is imposed on properties of responding or on properties of the stimuli during which responding occurs. The main implication of the difference is procedural: In studies of differentiation the experimenter must wait for the organism's responses, whereas in studies of discrimination the experimenter controls the order and duration of stimulus presentations. The vocabulary of differentiation and discrimination is summarized in Table 11-1.

Differential reinforcement can be based on simple dimensions of stimuli such as intensity or location. The experimental question is whether responding conforms to the differential consequences, in that more responding occurs during the stimuli correlated with reinforcement than during those correlated with nonreinforcement. The implications are profound. What we've learned from training simple discriminations in pigeons or rats has been used to develop methods for training primates to discriminate among different drugs that have been administered to them (e.g., Schuster & Balster, 1977), for studying the sensory capacities of nonverbal organisms (e.g., Blough, 1958a), and for teaching young parents to identify their children's illnesses (e.g., Delgado & Lutzger, 1988) to mention just a few of a vast number of successful applications.

Table 11-1 The Vocabulary of Differential Reinforcement

Differential Reinforcement (Procedure)	*Concentration of Effects of Reinforcement (Outcome)*	*Spread of Effects of Reinforcement (Outcome)*	*Differential Reinforcements by Approximations*	*Class*
With Respect to Response Properties	Differentiation	Induction	Shaping	Operant
With Respect to Stimulus Properties	Discrimination	Generalization	Fading	Discriminated Operant

As illustrated by these applications, differential reinforcement can be arranged for complex properties of stimuli that aren't easily quantified. For example, children learning to read must be able to name the letters of the alphabet. But the properties important for distinguishing among some letters are different from those important in distinguishing among others (e.g., straight line versus curve is important in distinguishing *U* and *V* but not *V* and *N*), and different distinctions are important for lowercase than for uppercase letters (e.g., no pair of uppercase letters has the up-down or left-right reversals that must be mastered to read *b*, *p*, *d* and *q*). The way in which a child learns to distinguish among letters of the alphabet depends on relations among such stimulus properties as symmetry, curvature and closure. Those properties essential to discriminating among different letters are called *critical features* (e.g., E. J. Gibson, 1965).

It isn't enough, however, to enumerate critical features. For some letters, the uppercase and lowercase forms differ more from each other than they differ from other letters (e.g., *e*, *E* and *F*, or *h*, *n* and *N*). Given the multitude of forms, what then defines the class of stimuli that occasions our saying *A* or *B* or *C*? This question is about the stimulus structure of the letters of the alphabet. The problem becomes even more complicated when different contexts are considered. For example, *O* could be either a letter or zero, and *I* could be either a letter or a roman numeral. The concept of an *X* or a *Y* or a *Z* is defined by the class of stimuli to which we respond with the corresponding letter name, but such classes are based on behavior, not on common physical properties (cf. Chapter 22).

Discriminated operants are behavior classes defined by the stimuli that occasion responding. We often identify such classes in our everyday vocabulary, as when we speak of stopping at a red traffic light or answering a telephone. The red light can vary in brightness and size and the telephone ring can vary in loudness and timbre, but our behavior is reasonably independent of variations along such dimensions, so we speak in terms of these classes rather than in terms of specific instances.

Natural Concepts and Probabilistic Stimulus Classes

We can speak of concepts as generalization within a class of stimuli and discrimination between classes (Keller & Schoenfeld, 1950). Thus, our concept of red must involve generalization among all stimuli we call red and discrimination between those and the ones we call by other names. Thus, concepts are to classes of stimuli what operants are to classes of responses. Responding on the basis of some single property of stimuli is sometimes called *abstraction*, and the language of concepts is sometimes restricted to responding that is based on some combination of properties. But these are ambiguous distinctions. For example, being-to-the-left-of can be treated either as a single relational property or as a combination of properties necessarily including both a reference point and a stimulus to the left of that reference point.

Time is another dimension with ambiguous properties. What sort of stimulus is it? If we differentially reinforce a rat's holding down of a lever, how should we speak of what it does? Are the holding times differentiated, or should we say the rat is discriminating how long it has held the lever down. Yet we can run experiments on temporal discrimination. Figure 11-7 provides an example. After a squirrel monkey heard a 1-s noise, its presses on a left lever were reinforced

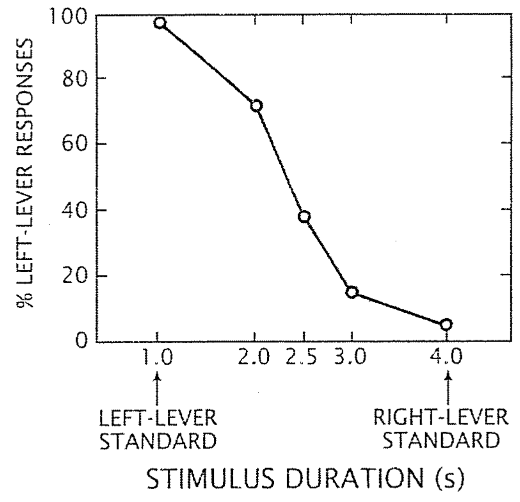

Figure 11-7 Temporal discrimination data for a squirrel monkey. Left-lever presses were reinforced after 1-s noises and right-lever presses were reinforced after 4-s noises. The orderly data over durations are similar to psychophysical data obtained along visual and auditory dimensions of stimuli. (Adapted from Catania, 1970, Figure 1-1)

with food. After it heard a 4-s noise, its presses on a right lever were reinforced. No reinforcers were arranged after noises of 2, 2.5 or 3 s. This procedure has the advantage that the discriminative response does not have temporal properties. The monkey accurately discriminated 1 s from 4 s, with intermediate responding at intermediate durations; its data along this temporal dimension are similar to psychophysical functions obtained along various dimensions of standard sensory stimuli such as lights and sounds.

It is often difficult to define discriminative stimuli by physical dimensions. For example, the properties that define the letter *A* vary according to whether it is uppercase or lowercase and whether it appears as type or as script. The capacity to discriminate among such stimuli exists in animals as well as humans (e.g., Morgan et al., 1976). But the difficulty of defining stimuli in terms of measurable physical properties isn't limited to arbitrary classes created by humans, such as letters and numbers. It is shared with natural objects and events. On what basis do we generalize between huskies and chihuahuas by calling both dogs while we discriminate between huskies and wolves even though they look more alike than huskies and chihuahuas?

Pigeons have been taught to discriminate between pictures that contain a human form and those that don't. Such discriminations have been called *natural concepts*. In one study (Herrnstein, Loveland, & Cable, 1976), slides were presented on a screen next to a pigeon's key and its pecks were reinforced during some slides but not others. Some pigeons learned discriminations between pictures with and without trees; others learned discriminations between pictures with and without water; still others learned discriminations between pictures with and without a person. After training with one set of slides (e.g., slides with and without trees), the pigeons discriminated among slides from those classes that had never been presented before. The new slides were sometimes more accurately discriminated than those used in training. The implications were that

> we cannot begin to draw up a list of common elements. To recognize a tree, the pigeons did not require that it be green, leafy, vertical, woody, branching, and so on (overlooking the problem of common elements nested within terms like leafy, vertical, woody, and so on). Moreover, to be recognizable as a nontree, a picture did not have to omit greenness, woodiness, branchiness, verticality, and so on.... If not common elements, what?... Pigeons respond to clusters of features more or less isomorphic with the clusters we respond to ourselves. (Herrnstein et al., 1976, pp. 298–299)

Natural concepts are examples of *probabilistic stimulus classes*, classes in which each member contains some subset of features but none is common to all. The number of features in the subset may vary from one class member to another. Such classes, sometimes called *fuzzy sets*, don't have well-defined boundaries, though class members may have family resemblances (Rosch, 1973).

Other types of stimulus classes include polymorphous stimulus classes (Lea & Harrison, 1978; see Glossary) and equivalence classes (about which more in Chapter 12). In fact, the class of stimulus classes is itself a probabilistic stimulus class, in the sense that its definition changes as we expand the boundaries of relevant research.

Some probabilistic stimulus classes are defined by reference to a *prototype*. A prototype is a typical member of a probabilistic class; it is derived from a weighted average of all of the features of all members of the class. For example, birds are a probabilistic stimulus class; most fly, but ostriches and penguins don't. In the production of a prototypical bird, feathers must be weighted more heavily than webbed feet because more birds have feathers than have webbed feet. Thus, a robin is more prototypical than a duck because it shares more features with other birds than does a duck.

Defining stimulus classes is a general problem. Appeals to physical measurement do not resolve it, because the reading of meters or other instruments is also discriminative behavior. Behavior classes depend on the common contingencies that created them and not on their physical properties. Even the behavior of the scientist depends on discriminations learned in the laboratory. Stimulus control is as fundamental with respect to our own scientific behavior as it is with respect to the behavior of the organisms we study. Thus, any effective philosophy of science must take it into account.

ADDENDUM 11A: POST-DISCRIMINATION GRADIENTS AND PEAK SHIFTS

Stimulus-control gradients can also be obtained after discrimination between two or more stimuli along a dimension. Figure 11-8 compares such a postdiscrimination gradient with a generalization gradient (Hanson, 1959). For one group of pigeons (generalization), key pecks during a single wavelength on the key were reinforced according to a VI schedule, and then rates of pecking during this and other wavelengths were obtained during nonreinforcement. The gradient peaked at the reinforcement stimulus; in other words, rate decreased as distance from this stimulus increased.

For a second group (postdiscrimination), key pecks were reinforced according to a VI schedule during the same wavelength as the first group, but this wavelength alternated with another during which pecks weren't reinforced (extinction). As for the first group, rates of pecking during this and other wavelengths were obtained during nonreinforcement. In this case, the peak of the gradient was displaced from the reinforcement stimulus in a direction away from the extinction stimulus; this displacement is called a peak shift. Similar effects occur when the discrimination is based on a higher frequency of reinforcement during one stimulus than another (Guttman, 1959).

Other interactions among gradients can be observed when stimulus control procedures are arranged concurrently on two different keys. For example, when a generalization or post-discrimination gradient along the dimension of line orientation is arranged on one key, an inverted inhibitory gradient can sometimes be obtained at the same time during a constant red or green stimulus on the other key (Catania, Silverman, & Stubbs, 1974).

Effects of discrimination training on the shape of gradients raised questions about the origins of peaked generalization gradients. One suggestion was that gradients should be steeper or shallower depending on whether stimuli in that region of the gradient were easier or harder to discriminate. But comparisons of the steepnesses of generalization gradients around stimuli in different regions of the spectrum with the thresholds for the detection of a change in wavelength in those regions found no simple relation between generalization and discriminability (Guttman & Kalish, 1956).

Another suggestion was that the peaked gradients depend on discrimination learning that occurs before the organism is brought into the experimental situation. For example, a pigeon presumably learns to discriminate among grains it eats long before it sees yellow projected on a pigeon key; its discriminations of color might be sharpest in the yellow region of the spectrum simply because yellow predominates as a color in its food. It is hard to control the color discriminations acquired in natural environments, but we can create environments in which color discriminations aren't possible. In an environment lit only by monochromatic light, a very narrow band of wavelengths such as the yellow emitted by a sodium vapor lamp, objects have no color; to a human in such an environment, everything appears in shades of gray. An organism reared in such an environment has no opportunity to learn color discriminations.

A suitable organism for monochromatic rearing is the duckling, which is capable of walking and pecking soon after it hatches; thus, reinforcement procedures can start early. Generalization gradients across wavelengths from ducklings reared monochromatically have been sometimes flat (Peterson, 1962; cf. Figure 11-4, top) and sometimes peaked (Rudolph, Honig, & Gerry, 1969; cf. Figure 11-4, bottom), suggesting that ducklings sometimes attend to color even without color experience. Furthermore, when ducklings reared monochromatically were trained to discriminate between only two wavelengths, their postdiscrimination gradients were similar to those of normally reared duckings and included peak shifts even though the stimuli consisted mostly of wavelengths they'd never seen before (Terrace, 1975).

ADDENDUM 11B: PLACE LEARNING VERSUS RESPONSE LEARNING

The response of moving from one place to another has had a privileged place in the psychology of learning (e.g., Olton, 1979). Going to food isn't the same as producing food while remaining in place. One important difference is that the environment changes more drastically when we move to a new location than when we introduce a new stimulus into our present one. We humans in particular have substituted other responses for ordinary locomotion: To go places, we step on gas pedals, turn steering wheels, press elevator buttons and stand on escalators or moving walkways. Furthermore, movement produces continuous changes in the environment, whereas sequences of other responses may produce no change until the sequence is completed. Beyond these differences, then, does it matter whether an organism moves to a new place or produces new stimuli where it is?

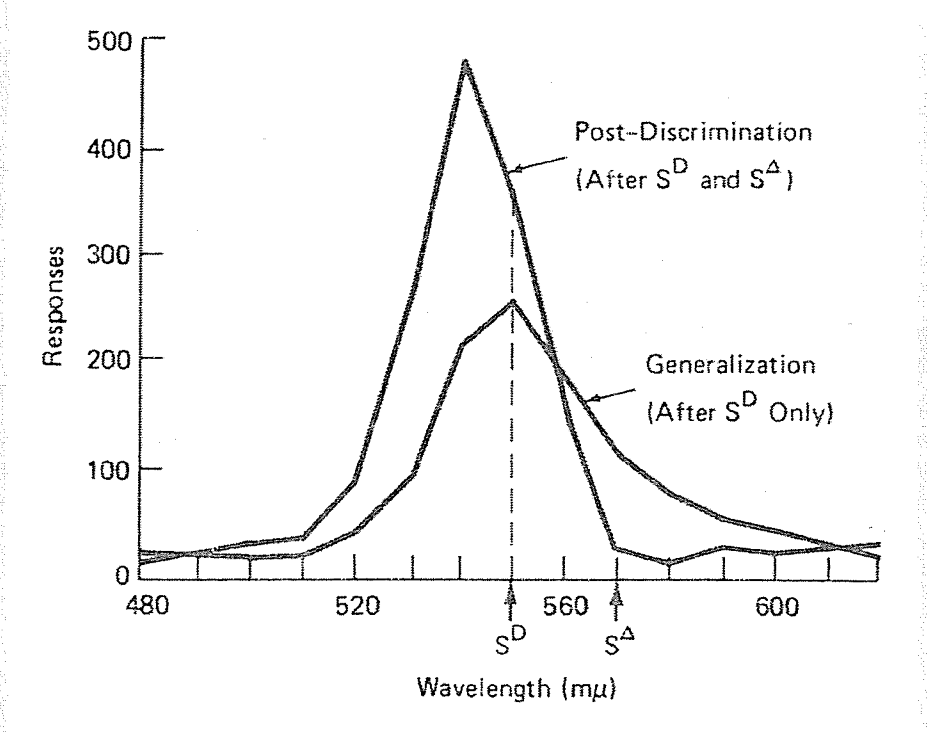

Figure 11-8 Gradients of stimulus control. The generalization gradient shows pigeons' key pecking after reinforcement at a wavelength of 550 millimicrons (S^D); the gradient peaks at the S^D. The postdiscrimination gradient shows key pecking after reinforcement at 550 millimicrons (S^D) and extinction at 570 (S^Δ); it shows a peak shift: Maximum responding has shifted from the S^D in a direction away from the S^Δ. The spectrum goes from violet and blue at short wavelengths to red at long wavelengths, so the S^D and S^Δ were in the green-yellow region. (From Hanson, 1959, Figure 1)

Such questions were implicit in the controversy over place learning versus response learning (Restle, 1957). We can differentiate the right turns of a rat at the choicepoint of a T-maze by reinforcing right but not left turns. We can then ask whether the rat's responding is based on response dimensions (movements to the right as opposed to those to the left) or on stimulus dimensions (movements toward a particular place, without regard to the direction from which the rat approaches). For example, suppose the right arm of a T-maze points toward the brighter windowed east wall of a laboratory. The rat might learn right turns or it might learn to run toward the window. We could test these alternatives by turning the T-maze around so the right arm points west and the rat now approaches the choicepoint from the north. If the rat still turns right but now away from the window, it shows *response learning*. If it turns left but still toward the window, it shows *place learning*: It moves toward the same place, even though turning in a different direction. Does the rat learn right versus left turns or east versus west turns?

The rat's performance depends on the stimuli available both inside and outside the maze. The typical maze used to be topped by wire mesh or some other cover that allowed the experimenter to watch the rat. If an experimenter can look in, a rat can look out. Although it is nearsighted, the typical rat can discriminate the general direction of lights and other gross features of a room. As long as stimuli are available outside the maze, they may become the basis for the rat's turn in a particular direction. But if those stimuli are eliminated by placing a roof over the maze, the direction that the maze faces in the room becomes irrelevant and the rat can't show anything but response learning. Whether the rat learns place or response depends on how the experimenter sets up the maze. By judiciously choosing conditions, an experimenter can make things come out either way.

In natural environments, food at a given location isn't necessarily replenished as in the goalbox of a laboratory maze. In foraging, an animal is more likely to move to a new location than to return to one where it has already consumed the available food. Again, the environmental features to which the organism attends can vary with constraints imposed by the experimenter (Collier & Rovee-Collier, 1981; Lea, 1979). For example, when a rat is given daily sessions with food located at the ends of each of several alleys and the food isn't replenished during the session, the rat learns not to repeat visits to alleys where it has already eaten (Olton & Samuelson, 1976). Spatial properties of environments are particularly important, but in appropriate circumstances a rat may learn other properties as well.

Chapter 12
Conditional Discrimination and Stimulus Classes

Matching *comes by way of Old English from the Indo-European* mag *or* mak-, *to knead or fashion; it is related to* make, among *and* mass. Sample *can be traced to the Indo-European root,* em-, *to take or distribute. Like* example, *it combines the Latin* ex-, *out, and* emere, *to buy or obtain; it is related to* exempt, prompt *and* consume. Comparison *is derived from the Latin* com- *with, plus* par, *equal; it is related to* part, pair *and perhaps also* repertory. Oddity *fittingly has a more singular etymology; it is derived from the Old Norse* oddi, *a point or a triangle.*

Equivalence *combines the Latin* aequi-, *equal, and valere, to be strong or have power. Relatives include* value, valid, available *and* wield.

- Relations as Stimulus Dimensions
  - *Matching to Sample and Oddity*
  - *Symbolic Behavior: Equivalence Classes*
- Higher Order Classes of Behavior
  - *Learning Set*
  - *Contingencies Operating on the Subclasses within Higher-Order Classes*
  - *Origins of Structure*
- Addendum A: Animal Cognition and Cognitive Maps

During many years of team-teaching the learning course with my colleague Eliot Shimoff we used animal demonstrations in the classroom. One was a demonstration of shaping in which a rat pressed a counterweighted lever, as described in Chapter 9. Another started out with a pigeon that pecked red but not green. I arranged the contingencies and the key colors with hand switches, alternating between red and green every 10 or 15 seconds or so. During red I reinforced about one in every four or five pecks on average and during green I didn't reinforce any. The deal was that Shimoff talked while I concentrated on the apparatus.

But we had a problem. As I worked with the pigeon Shimoff talked about stimulus control and his example was a pigeon that conformed to traffic signals, going on green and stopping on red. At one point he said that he could name the key color without even looking at the demonstration chamber, which was located somewhat behind him and to the side on the stage area of our classroom. He said all he had to do was listen, so while I had the pigeon pecking on a red key he said the key must be green and when I had it stopping during green he said the key must be red. The class seemed amused.

KEY TERMS: Conditional or Conditioned Discrimination; Matching-to-Sample and Oddity; Sample and Comparison; Arbitrary Matching; Equivalence Classes; Reflexivity, Symmetry, Transitivity; Higher-Order Classes.

We soon got into an argument about what we had agreed on for this demonstration and what instructions the teaching assistant had been given. The teaching assistant had been standing nearby but wanted nothing to do with us and wouldn't take sides. The only thing I could do was to offer to fix things by teaching the bird to reverse itself and start pecking green instead of red. When Shimoff suggested that doing the reversal in the available time would be impossible and that the class had been ruined, I replied that I expected to complete the reversal in less than ten minutes. He took me up on the offer and started a clock to see how long I would take.

It was a setup, of course, though most students claimed after class that they hadn't been fooled by our phony fight. The first thing I did was to stop reinforcing pecks on red, but the bird didn't slow down much at all while the key remained that color. Next I switched to green, and the pigeon, whose performance had been very well prepared by the teaching assistant, became practically immobile, facing the key but at some distance from it. I operated the feeder and the pigeon quickly came over to eat, so I immediately began shaping key pecks during green. It took just two or three reinforcers to get the bird to move its beak toward the key, and with only three or four more I had it pecking on green. But when I switched the key back to red the bird still pecked. It was now pecking both green and red, so only half the job was done. The bird now went on green, but I had yet to get it to stop on red.

I began reinforcing one of every three or four pecks on green, and after delivering a few reinforcers I switched back over to red, during which I reinforced no further pecks. That meant that green was itself becoming more potent as a reinforcer while red was becoming less so. In other words, only pecks on green now produced reinforcers, so by virtue of that contingency green was becoming a *conditional reinforcer.* And if it was a reinforcer, I should be able to use it. The pigeon was still pecking red a lot, but the next step was to wait for any slight pause in that pecking and to change the key from red to green during the pause. I would use the onset of green to reinforce pauses during red.

Within the first minute or so I caught some good pauses and let the bird produce two or three reinforcers during green before going back to red. As the pauses during red became longer I could now pick up head movements away from the red key and make the onset of green contingent upon them. Pretty soon the pigeon was backing away from the key as soon as it turned red and moving toward the key and pecking it as soon as it turned green. By alternating quickly between green and red I could even get the pigeon to bob its head in front of the key for a few seconds: It moved toward pecking the green key, but when I turned the key red before it got there it backed off and I could turn the key back to green again. Less than ten minutes had elapsed on Shimoff's clock and he now had a pigeon that went on green and stopped on red. We then debriefed the class and were especially pleased if our contrived argument had been convincing enough that some students had really begun to worry about us.

But things went strangely one semester. The pigeons for our demonstrations began in our research laboratory. Most of our research involved steady-state procedures, so they were around for a long time (our oldest was twenty-one when it died of natural causes). One bird had been devoted to the reversal demonstration for several semesters. Ahead of time each semester one of us or a teaching assistant shaped up the performance of pecking red but not green. Then we did the reversal in the classroom. After returning to the laboratory, we again shaped up pecking red but not green so the bird would be ready the next time around. This time the bird seemed just fine when we checked it out a day or so before our demonstration, but when we got it to the classroom it began to pause on red and peck on green well before Shimoff and I got to our argument. Though we then demonstrated a reversal, it lacked the classroom drama that we had planned for and come to expect.

What had happened? We finally realized that our pigeon had been through these procedures so often that it had learned a *conditional discrimination*, where the conditional stimulus was whether it was in the laboratory or the classroom. In the laboratory we always reinforced pecks on red but not green; in the classroom, after the first couple of minutes we always reinforced pecks on green but not red. The bird had finally reached the point where it didn't wait for us before starting its reversal. We solved the problem by doing reversals in the laboratory as well as in the classroom. And with that example let's move on to conditional discrimination as a primary topic of this chapter.

Like contingencies, discriminations may be effective under some conditions but not others. For example, your response to a green traffic light will depend on whether you're facing it or looking at it from the cross traffic. Such discriminations, in which the role of one stimulus depends on others that provide its context, are called *conditional discriminations*. Consider the case of attention in the pigeon, discussed in Chapter 11 (Figure 11-2). The stimuli were triangles or circles on red or green. We might have added a lamp above the key and reinforced pecks in the presence of triangles when it was lit and pecks in the presence of red when it wasn't. Under these circumstances, the pigeon will learn to peck triangles but not circles without regard to color when the lamp is on, and it will learn to peck red but not green without regard to form when it is off. In other words, whether the pigeon attends to form or color is conditional on whether the lamp is lit.

In Chapter 11 we considered more socially significant conditional discriminations, such as those that that might be learned on city streets or in academic institutions. Here we'll examine other ways in which they may function. We'll consider several procedures in which discriminative contingencies depend upon the context within which they're arranged. These procedures sometimes generate higher-order classes of behavior, in the sense that the classes are defined not by particular stimuli or responses but rather by relations that include those stimuli and responses as special cases. The procedures we'll review include matching-to-sample, oddity, arbitrary matching and learning set. They provide the context in which we'll explore the conditions under which one stimulus can become the functional equivalent of another (as when, in varied settings, uppercase *A* and lowercase *a* function as the same letter even though they look very different).

Note that matching and oddity are relational properties of stimuli. We often treat stimuli as if they were restricted to concrete objects or environmental events. But as we learn to respond in consistent ways to objects or events, we also learn to discriminate among features, sometimes called *abstract* or *relational*, that are independent of particular objects or events. The term *stimulus* often functions in this more general way, in the sense of some property of environmental events. For example, we might say that a chair had been placed to the right of a table. Although the chair and the table are concrete objects, being-to-the-right-of is not, and yet we discriminate this relation from being-to-the-left-of. In some discrimination experiments, therefore, relations among stimuli have been dimensions of special interest.

RELATIONS AS STIMULUS DIMENSIONS

Many conditional discriminations involve arbitrary relations between a conditional discriminative stimulus and the discriminations for which it sets the occasion. Some cases in which those relations aren't arbitrary are of special interest. For example, whether one stimulus in a set of stimuli matches one or more of the others or is an odd stimulus depends on, or is conditional on, the context of other stimuli within which it's presented. For example, if stimuli A and B are blue and stimulus C is yellow, then B is a matching stimulus in relation to A and C is an odd stimulus in relation to A or B. We'll now consider some properties of such conditional discriminations.

Matching-to-Sample and Oddity

Matching-to-sample is illustrated as we might arrange it in a three-key pigeon chamber in Figure 12-1 (cf. Ferster, 1960; Skinner, 1950). During an intertrial interval, all keys are dark. A trial begins when the center key is lit, presenting the *sample*. Typically a peck is then required on the center key. This peck, sometimes called an *observing response*, turns on new stimuli and also makes it likely that the pigeon has looked at the sample (cf. the discussion of observing responses in Chapters 11 and 16). The two stimuli produced by the center-key peck are called *comparisons*; one comparison matches the sample and the other doesn't. A peck on the matching key then produces a reinforcer followed by a new intertrial interval, but one on the nonmatching key is followed directly by the intertrial interval without a reinforcer. Sometimes nonmatching pecks also extend intertrial interval, which may then function as a mild punisher (cf. Holt & Shafer, 1973).

Both the sample stimulus (G or R in Figure 12-1) and the position of the matching comparison (left or right in Figure 12-1) usually change over trials. A common feature of matching-to-sample is a correction procedure, which repeats the same samples and comparisons on the next trial if a trial ends with a peck on a nonmatching key. This procedure prevents the development of responding

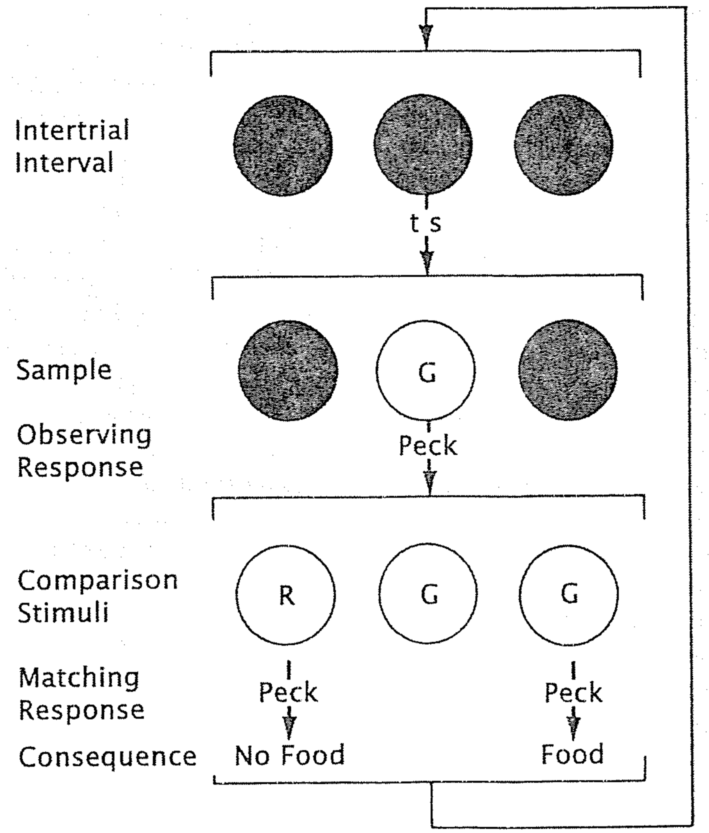

Figure 12-1 Diagram of a matching-to-sample trial in a three-key pigeon chamber. After an intertrial interval of *t* s, a sample stimulus (green: G) appears on the center key. A center-key peck turns on the two side keys. One comparison stimulus matches the sample; the other doesn't . A peck on the matching comparison stimulus produces food, after which the next intertrial interval starts; a peck on the nonmatching comparison starts the next intertrial interval without food. The sample stimulus and the left-right locations of the matching comparison vary from trial to trial.

restricted to only one key or color, but it also guarantees that errors are often closely followed by the reinforcers produced by correct responses.

Suppose a pigeon pecks only the comparison on the left. If the matching key alternates irregularly between left and right, these pecks will be reinforced on half the trials (those with the matching stimulus on the left). Reinforcement on half the trials will probably be enough to maintain pecking on the left key indefinitely. With a correction procedure, however, the pigeon has to switch to the right key sooner or later, because a trial with a matching stimulus on the right will repeat itself until a right-key peck occurs and is reinforced. Similarly, if the pigeon always pecks red comparisons, these pecks will be reinforced on half the trials unless a correction procedure forces the pigeon to switch occasionally to green.

Another procedural refinement is to enhance the pigeon's attention to the sample by modifying the contingencies arranged for pecks on the sample key. For example, if production of the comparisons depends on different patterns of responding for each sample (e.g., different response rates), the response patterns the pigeon produces can tell us how well it's discriminating between the samples. Such procedures sometimes produce faster acquisition of matching, but unfortunately enhanced attention to the sample sometimes also reduces attention to the comparisons (e.g., Urcuioli, 1985).

In any case, suppose now that a pigeon is responding accurately given red (R) and green (G) in a matching-to-sample procedure. How should we describe its performance? Has it just learned to peck left given the configurations RRG and GGR and right given GRR and RGG? Or has it learned matching in general, the identity relation? If we now present blue or yellow and the pigeon matches the new sample colors we'll feel more confident about speaking of generalized matching (in fact, matching in pigeons doesn't transfer easily to new colors, though the likelihood of such transfer depends on the details of training). Even if we saw matching with new colors, what if we failed to get matching with geometric figures? We might just say the pigeon has learned color but not form matching, noting that the human concept of matching seems not so limited by specific dimensions of the stimuli. Matching-to-sample has sometimes been called *identity matching*, but we reserve that term for cases in which matching generalizes to novel samples and comparisons, such as matching of forms after training with colors. With a limited stimulus set, what looks like matching on the basis of identity too often turns out to be matching based on specific stimulus configurations.

There are many variations on matching-to-sample. If reinforcement is arranged for pecks on the nonmatching rather than the matching comparison, the procedure becomes an instance of *oddity* responding, because this key is necessarily the odd one of the three keys. This version of oddity requires a minimum of two pecks per trial, one on the sample and the other on the nonmatching comparison. Other versions of oddity use no sample; with pigeons in a three-key chamber, for example, each trial consists of lighting all three keys, with the odd key any one of them. This procedure requires just one peck per trial.

The relations among stimuli can also involve *arbitrary matching*. For example, we could train the pigeon to peck green given a square as a sample and red given a circle as a sample. We might then ask whether the pigeon might peck a square given a green sample and a circle given a red sample. We often expect such reversibility when we deal with words and objects, as when a child who's learned to point to a picture of a car on seeing the word *car* can also point to the word on seeing the picture. This reversibility, a property of *symbolic behavior*, isn't to be taken for granted. For example, a child may show such reversibility without explicit training but a pigeon won't. Arbitrary matching can also be extended to cases in which the same matching response is trained with more than one comparison, as in pecking a green comparison given either a circle or an ellipse as a sample, and pecking a red one given either a square or a triangle; (cf. Zentall & Urcuioli, 1993, on many-to-one and one-to-many matching).

One recurrent problem in analyzing procedures like these is that the coping of a pigeon or

some other organism with matching or oddity or related conditional discriminations varies considerably depending on procedural specifics. In some experiments pigeons rarely show generalized matching or generalized oddity; in others they do well (e.g., Katz & Wright, 2006). One condition that makes a big difference is the number of stimuli used in the procedure. The matching-to-sample procedure in Figure 12-1 uses just two stimuli, red and green. Other procedures use large populations of slides, such as travel scenes, so that the stimuli used in one particular match are never repeated. A strong case can be made that the two-stimulus procedure is far more difficult than the many-stimulus procedure, in that the contingencies arranged for common stimuli across different trials in the two-stimulus procedure interfere with each other, whereas the large populations in procedures with many stimuli minimize such interference (Sidman, 1987; Wright, 2012; Wright, Katz, & Ma, 2012; Wright & Roediger, 2003).

Tasks learned later can interfere with performance on tasks learned earlier (retroaction), but the effects of tasks learned earlier on those learned later (proaction) are usually even more powerful. We'll consider retroactive and proactive interference later, especially in the context of human verbal learning and memory (Chapters 21, 26 and 27). We'll also examine other variables that influence the acquisition of relational discriminations such as matching and oddity, as in the differential-outcome or reinforcer-specific effects to be discussed in Chapter 14. Still another variable is suggested by a distinction analogous to that between response rate and resistance to change, reviewed in Chapter 6. The accuracy of conditional discriminations and the rate at which that discriminated responding occurs are separate measures that may be dissociated (Nevin *et al.*, 2003), so interpretations of such performances may sometimes be biased depending upon which measures are used.

The various tasks we've just considered here so far involve relations among different stimulus classes. Some are arbitrary, as when we designate pecks on green as correct given square as a sample and pecks on red as correct given circle as a sample; but others such as matching and oddity seem to involve more fundamental dimensions. Still another issue is the role of verbal behavior when we compare human and nonhuman performances on such tasks (see Chapter 23). How then should we deal with relations such as same and different and opposite, and what can we say about the contingencies that created them (cf. Hayes, 1994; Hayes, Barnes-Holmes, & Roche, 2001, on relational frames) ? Questions about stimulus structure seem to be inextricably interwoven with questions about contingencies and the structure of behavior (e.g., Fujita, 1983; Lamb & Riley, 1981; Wasserman, Kiedinger, & Bhatt, 1988). As long as interference and differential outcomes and verbal behavior and incommensurable procedures or data analyses cloud our interpretations of conditional discriminations, caution in broadly extending findings from this literature may yet be well-advised. Nevertheless, it is appropriate to consider some of its achievements.

Symbolic Behavior: Equivalence Classes

If a pigeon's pecks are reinforced in the presence of green but not red, we wouldn't even consider the possibility of the pigeon's greening in the presence of pecks. Such a reversal makes no behavioral sense. But this isn't so in matching. Both the sample stimulus and the comparison response are defined by the stimuli presented on the keys. We could therefore ask about a red response to a red stimulus, or about the reversibility of a vertical response to a diagonal stimulus, or about whether round responses to dim stimuli can be created by training round responses to large stimuli and then large responses to dim stimuli.

These cases illustrate properties of the relations called *reflexivity*, *symmetry* and *transitivity*. Reflexive properties are those that hold between a term and itself (e.g., A = A); symmetrical properties are those in which the order of terms is reversible (e.g., if A = B then B = A); and transitive properties are those in which the common terms in two ordered pairs determine a third ordered pair (e.g., if A = B and B = C then A = C). Equivalence relations are those that have all three properties, and the terms that enter into them (here, A, B and C) are said to be members of an *equivalence class* (Sidman, 1994). Typically, an equivalence class is never considered in isolation. For example, we might contrast a class identified as A_1, B_1 and C_1 with another class consisting of A_2, B_2 and C_2, or of α, β and γ, or of X, Y and Z.

For other relations, only some of these three properties hold. For example, the relation of *oppositeness* is symmetrical (if D is the opposite of E, E is the opposite of D), but it's neither reflexive (D is not the opposite of itself) nor transitive (if D is the opposite of E and E is the opposite of F, D isn't the opposite of F; instead, D is the same as F). And a magnitude relation such as *greater than* is transitive (if G is greater than H and H is greater than I, then G is greater than I), but it's neither reflexive nor symmetrical.

Figure 12-2 illustrates how standard and arbitrary matching procedures can be used to demonstrate reflexivity, symmetry and transitivity, where the first equivalence class consists of red, Δ and ◍ (light), and the second one consists of green, O and ● (dark). Each procedure includes two samples, shown as Sample A in the middle row of the three-key array on the left and as Sample B in that on the right. For convenience, the matching comparison is always shown in the left position of each three-key array, even though in practice the side positions of the comparisons vary from trial to trial. As in Figure 12-1, a peck on the sample produces the comparisons, and a peck on the matching comparison produces food whereas one on the nonmatching comparison doesn't.

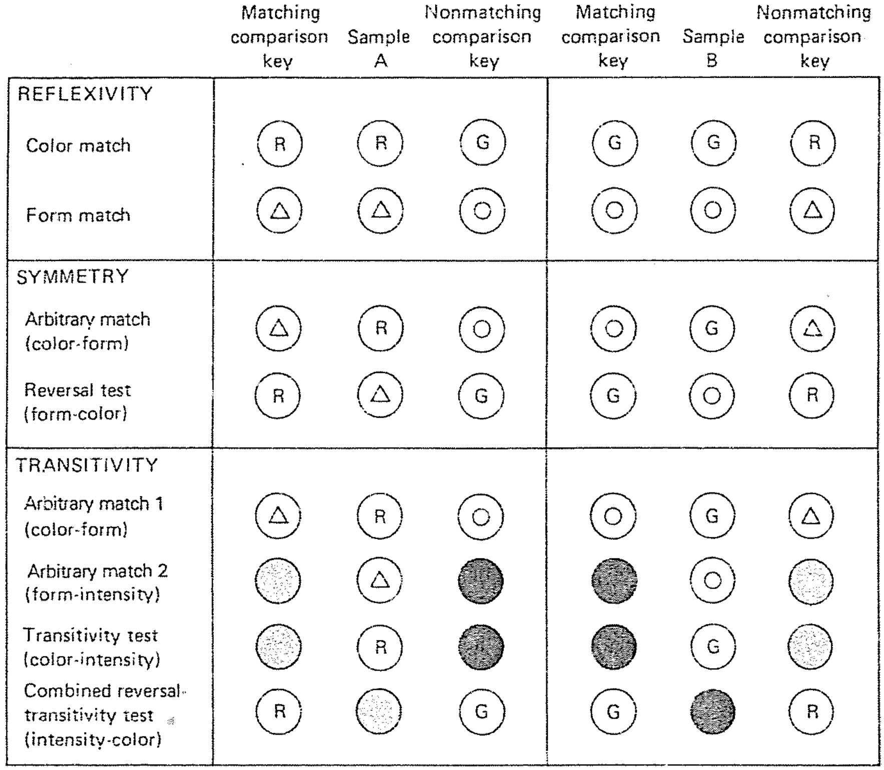

	Matching comparison key	Sample A	Nonmatching comparison key	Matching comparison key	Sample B	Nonmatching comparison key
REFLEXIVITY						
Color match	R	R	G	G	G	R
Form match	Δ	Δ	O	O	O	Δ
SYMMETRY						
Arbitrary match (color-form)	Δ	R	O	O	G	Δ
Reversal test (form-color)	R	Δ	G	G	O	R
TRANSITIVITY						
Arbitrary match 1 (color-form)	Δ	R	O	O	G	Δ
Arbitrary match 2 (form-intensity)	◍	Δ	●	●	O	◍
Transitivity test (color-intensity)	◍	R	●	●	G	◍
Combined reversal-transitivity test (intensity-color)	R	◍	G	G	●	R

Figure 12-2 Matching procedures for studying three properties of equivalence relations: reflexivity (identity matching), symmetry (reversal of arbitrary matching), and transitivity (transfer across ordered pairs of arbitrary matches). Each three-key array is shown in only one of its two possible arrangements, with the matching comparison on the left. Stimuli in equivalence class A, on the left, include red (R), triangles (Δ) and light (◍); those in class B, on the right, include green (G), circles (O) and dark (●).

The top two rows of Figure 12-2 (reflexivity) illustrate standard matching procedures both with color and with form. The next two rows (symmetry) illustrate symbolic matching with color samples and form comparisons and then a reversal test with form samples and color comparisons. The bottom rows (transitivity) show how the common stimuli in two symbolic matching procedures (color to form and form to intensity) can be combined in a test of transitivity (color to intensity). An alternative test, illustrated in the last row, combines the reversal and transitivity tests (intensity to color); it's sometimes called an equivalence test, and the relation between the stimuli of the equivalence test is sometimes called an equivalence relation. The relations of the reversal and equivalence tests were never explicitly taught. If appropriate matching occurs in these tests, the new relations demonstrated by that behavior are called *emergent* relations, in the sense that they've emerged without explicit training; they are examples of novel behavior produced by the arbitrary matching contingencies.

When we say that a pigeon that pecks in the presence of green can't green in the presence of pecks, we're saying that the relations between antecedent stimuli and the responses they occasion in a three-term contingency aren't symmetrical. This means that we can't reduce equivalence classes to classes of discriminative stimuli. Equivalence relations are symmetrical, but the relations among the terms of a three-term contingency are not.

We've already questioned whether the pigeon's standard matching performance depends on the identity relation in general or only on relations among specific stimuli. We might assume that the identity relation between sample and comparison stimuli would make standard matching easier than arbitrary matching. For pigeons, however, acquisition of matching depends more on the stimulus dimensions chosen for samples and comparisons than it does on whether the procedure is standard or arbitrary matching. For example, both standard and arbitrary matching develop more rapidly with red and green samples than either does with vertical and horizontal samples (e.g., Carter & Werner, 1978). Symmetry and transitivity have been demonstrated in a limited way with pigeons, but the effects have not ordinarily been robust (e.g., D'Amato *et al.*, 1985; Richards, 1988). Furthermore, there are technical as well as logical pitfalls to be avoided in the analysis of equivalence relations via transfer tests, such as whether the relevant conditional discriminations have been learned as picking out the matching comparison rather than as *exclusion*, or rejecting the nonmatching one (e.g., Carrigan & Sidman, 1992). Thus, the case for equivalence relations in the pigeon's matching behavior is at best not compelling (no doubt all of the component relations could be explicitly taught, of course, but that outcome would be mainly of procedural interest).

Equivalence relations are easily generated in humans. For example, retarded youths who'd already shown reflexivity (identity matching) were taught matches involving (A) spoken words, (B) pictures corresponding to those words and (C) the corresponding printed words (e.g., car, dog, boy: Sidman, Cresson, & Willson-Morris, 1974), where each combination of spoken word, picture and written word constituted a different class. First, given any one of 20 spoken picture words they learned to pick out the corresponding pictures from a comparison set (AB). Then, given the same 20 pictures they learned to pick out the corresponding printed words from a set (BC). For all 20 words, these two kinds of matches (spoken words to pictures and pictures to printed words) also generated four other relations without additional training: two new relations through symmetry (BA, given a picture saying its name; and CB, given a printed word picking out its picture), one through transitivity (AC, given a spoken word picking out the corresponding printed word); and one through transitivity plus symmetry (CA, given a printed word saying it). Forty relations had been taught (20 each in AB and AC) and another 80 emerged indirectly (in BA, CB, AC and CA). The reinforcement of arbitrary matching had created the beginnings of a reading repertory in these youths. The emergent relations justify calling their performances *symbolic matching* rather than *arbitrary matching*.

Each of the 20 equivalence classes consisted of a picture and the corresponding printed and spoken words. New classes could have been added with new pictures and corresponding words, and the number of equivalences could have been expanded by adding new relations (e.g., printed words and words written in longhand, or pictures and actual objects). There are no obvious limits to the number of classes that can be created or the number of stimuli that can be included within each class. The status of such classes in the behavior of nonhumans remains a matter of controversy (e.g., Dube *et al.*, 1993; Horne & Lowe, 1996; Schusterman & Kastak, 1993; Zentall & Urcuioli, 1993). But equivalence classes define symbolic behavior, so much of the interest in equivalence classes arises because of their possible relevance to the phenomena of language, which we'll consider in Chapters 20 through 24.

The members of an equivalence class are equivalent in the sense that we can exchange one for another in the context of arbitrary matching procedures. But this doesn't necessarily mean that they're functionally equivalent, in the sense that one will substitute for another in other contexts. In other words, *functional equivalence* isn't the same as membership in an equivalence class, and it must not be assumed that the logical properties of these classes are fully consistent with their behavioral ones (Saunders & Green, 1992). The functional equivalence of the members of an equivalence class must be tested experimentally. The question is whether a function acquired by one member of an equivalence class, such as a discriminative function, will transfer to other members of that class (cf. Sidman *et al.*, 1989).

Consider an example. A child has learned to obey a parent's words, *go* and *stop*, when crossing with the parent at a traffic intersection. In a separate setting, the child is taught that *go* and green traffic lights are equivalent and that *stop* and red traffic lights are equivalent (in other words, *go* and green become members of one equivalence class and *stop* and red become members of another). If the discriminative functions of the words *go* and *stop* transfer to the respective traffic lights, the child will obey the traffic lights without any additional instruction. An analogous type of transfer has been demonstrated experimentally with children: High and low rates of responding occasioned by one set of stimuli transferred to another set when the stimuli in the sets were made members of equivalence classes (Catania, Horne, & Lowe, 1989a; cf. de Rose *et al.*, 1988).

It is a different story, of course, when transfer occurs along other dimensions, such as judgmental words or actions and racial or ethnic or other identifications. We must look to the conditional discriminations that can be shaped within academic institutions and within communities that are tolerant of diversity to counter the effects of such social contingencies. We'll consider the role of words in our treatment of verbal behavior in Part IV.

HIGHER-ORDER CLASSES OF BEHAVIOR

In our examples of matching-to-sample and related procedures, we've typically described performances at more than one level of analysis. We described pecks to red comparisons given a red sample as conditional discriminations, but we also asked whether we could characterize this performance as identity matching. Our decision depended on whether we had created an operant class defined by particular stimuli in a particular context, or one that generalized across relations over a range of different stimuli. It makes a difference whether a pigeon that matches green to green and red to red also matches with other colors, or whether a pigeon that matches to sample across a range of colors also does so when we instead present shapes or textures or other dimensions. Similarly, it makes a difference whether a pigeon's arbitrary matching is an operant class restricted to a particular set of samples and comparisons or instead includes all of the reflexive, symmetrical and transitive relations that define equivalence classes.

These examples involve classes embedded within other classes. If we demonstrate identity matching, each specific match defines a class, but identity matching is then a higher-order class that includes all of the specific matches as its components. In matching-to-sample, the matching of green to green or red to red may exist as separate operants. They can be treated as instances of identity matching only if they can be shown to be components of a single higher-order operant defined by the identity relation between sample and comparison. One test for this higher-order operant is whether novel relations can be demonstrated, such as the matching of new colors (e.g., blue to blue or yellow to yellow); another is whether all of the matching subclasses hang together as a class if we change contingencies only for some subset of them. It is all too easy to get caught up in coordinations among classes, as with relational frames, without adequately taking into account the contingencies that created them or the different and not necessarily compatible contingencies that operate on the various levels of higher-order classes. What follows examines some higher-order classes, reviews some of their properties, and considers some implications that can be drawn from them.

Learning Set

The learning of a new discrimination can depend on what the organism has already learned (what an organism has learned in the laboratory is its *experimental history*). As illustrated by the phenomenon called *learning set* (Harlow, 1949), responding may depend on relations among stimulus properties independently of specific stimuli. Two different objects were presented to a food-deprived monkey in successive trials, with their positions varying from trial to trial. Food was always available under only one of them, so picking that one up was differentially reinforced. After a number of trials, the monkey mastered the discrimination, choosing the one with food under it almost every time. Then another pair of objects was presented in a new problem, again with picking up only one of them reinforced. Again a discrimination emerged. A new problem with another new pair then followed, and later still another, and so on.

Averages across eight monkeys are shown in Figure 12-3. Each set of points shows percent correct responses over the first 6 trials across blocks of successive problems. Over the first 8 problems (1-8), correct responses increased gradually over trials, but by trial 6 they hadn't even reached 80%. In the next block of problems (9-16), correct responses increased more rapidly over trials. Over successive blocks, correct responding rose more and more rapidly over trials until, by the last block (trials 289-344), responding reached nearly 100% accuracy on the second trial of each new problem. In other words, the more problems the monkey had mastered, the more rapidly it mastered each new one.

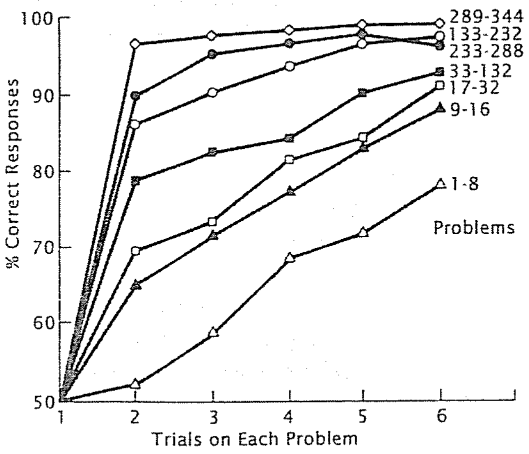

Figure 12-3 Learning set or learning-to-learn. Eight monkeys learned successive two-choice discrimination problems. Each set of points shows percent correct responses over the first 6 trials of each problem for blocks of successive problems (the y-axis starts at 50%, chance level). The rate of learning increased over successive problems. By the last block of problems (problems 289-344), responding was almost 100% accurate by trial 2 of each new problem. (Adapted from Harlow, 1949, Figure 2)

In this procedure, we can't describe the discriminated operant just in terms of a stimulus pair. When a monkey acquires successive discriminations so rapidly that it consistently picks the stimulus correlated with reinforcement after a single trial with a new pair, its performance depends on relations between stimuli and their correlated consequences within each problem and not on particular stimuli. The monkey had to learn many other features of the procedure besides which stimulus was correlated with food. It had to learn that food was correlated with just one of the two stimuli rather than with position or other dimensions of the setting, that this correlation didn't change within a problem, that reaching simultaneously for both stimuli wasn't reinforced, and so on. During early problems, learning took place slowly because the monkey was learning these many things; eventually, after the monkey had learned these other things, all it had to learn in a new problem was which stimulus was correlated with food.

At this point, we can define the discriminated operant as follows: If picking up one of the objects is reinforced on the first trial of a new problem, pick that one on all later trials; if it isn't, switch to the other object on all later trials. In learning set, this is the behavior that is reinforced and this is what the organism does. The correspondence between reinforcement contingencies and the behavior generated by these contingencies remains the criterion for this operant class. Learning set qualifies as a higher-order class because it is defined by these relations and not by the stimuli and responses of any particular problem.

One variation on learning set is *learned helplessness* (Maier, Seligman, & Solomon, 1969). For example, rats that receive inescapable and unavoidable electric shock in one situation are sometimes less likely to learn avoidance when it becomes possible in a new situation than rats never exposed to inescapable and unavoidable shock (cf. Maier, Albin, & Testa, 1973). Here again it is difficult to define the response class, except to note that the rats given inescapable and unavoidable shock apparently learned that their responses didn't have important consequences; these contingencies generated behavior that transferred to situations in which responding could have had consequences.

Learning set shows that we can't treat the introduction of an experimentally naive organism into a laboratory setting as a simple case (cf. Chapter 5 on acquisition). To study learning, some investigators have therefore turned to the repeated acquisition of simple discriminations or simple response sequences. Paradoxically, they have come to study learning, defined by changes in performance, in the context of steady-state procedures, defined by stability in performance. For example, with four groups of three levers each in a monkey's chamber, only one particular sequence of presses was reinforced within any session (e.g., left lever of group 1; right lever of group 2; middle lever of group 3; right lever of group 4), but the sequence changed from session to session (Boren & Devine, 1968). Once the monkey was consistently mastering a new problem each day, these repeated acquisitions provided a baseline for studying effects of different fading procedures, or of drugs on learning, or of creating the response sequence as a whole as opposed to building it up either from the beginning or from the end, and so on.

Contingencies Operating on the Subclasses within Higher-Order Classes

A significant property of higher-order classes was implicit in our treatment of self-injurious behavior in Chapter 9. In that example, self-injurious behavior maintained by attention was not easily reduced by extinction as long as it remained part of a larger class (attention-getting) other members of which continued to be effective. Whenever the contingencies maintaining a higher-order class differ from those maintaining one or more of its subclasses, we must ask which contingencies will prevail. For example, imagine a higher-order class consisting of a child's imitations of the actions modeled by a puppet. When the puppet claps, the child claps; when the puppet laughs, the child laughs; when the puppet jumps, the child jumps; and so on. Let's identify a dozen different imitations by letter, A through L. Once we've reinforced those, one way to tell whether we've created a higher-order class is to see whether the child imitates new actions that the puppet has never modeled before.

Now suppose we stop reinforcing G while we continue reinforcing all of the other imitations, A through F and H through L. Whether imitations of G decrease or instead persist long after they've been excluded from the reinforcement contingencies is an experimental question. If they decrease, we've demonstrated that we can pull G apart from the other members of the higher-order class. If they don't, we have to assume that G is still functioning as a member of the higher-order class called imitation, defined by the correspondences between what the puppet does and what the child does. As long as a higher-order class maintains its integrity, its subclasses are maintained along with the other members even if they are not consistently involved in the contingencies that maintain the others. When that happens, the lower-order subclasses will seem insensitive to the changed contingencies that have been arranged for them.

When we arrange new contingencies for subclasses of higher-order classes, we ordinarily expect those subclasses eventually to be differentiated from the others. But if the subclasses overlap in various ways with other classes that share in other contingencies, this may not happen. For example, imitating what someone has said shares class membership with other imitations but it may also participate in social contingencies that don't involve imitation. This may be enough to maintain it as a subclass when contingencies change for it but not for other kinds of imitations. These complexities are implicit in the following account of ontogenic selection, where the "rules, principles, strategies and the like" correspond to what we've been calling higher-order classes.

> In more mature human beings, much instrumental behavior and, more especially, a great part of verbal behavior is organized into higher-order routines and is, in many instances, better understood in terms of the operation of rules, principles, strategies and the like than in terms of successions of responses to particular stimuli.... If one who is attempting to describe and predict the behavior of an adult human learner fails to take account of these behavioral organizations, and attempts to construct an account in terms only of individual stimulus-response units, the principles of operation of rewards and punishments may appear to be quite different from those revealed in simpler experiments with animals or immature human learners. Actually, it may be that the principles of operation of these factors are the same in all cases and that the difference lies in the nature of the behavioral units whose probabilities are being modified.... (Estes, 1971, p. 23)

In matching-to-sample and in imitation, as in learning set, training with many specific instances may sometimes be a sufficient prerequisite for higher-order or generalized classes. For example, training with many symmetry problems may produce generalized symmetry, training with many transitivity problems may produce generalized transitivity. We'll encounter other examples of higher-order classes in other contexts, as when we deal with novel behavior (Chapter 13), social learning (Chapter 19) and verbal behavior (Chapters 22 and 23).

Origins of Structure

Most of the classes we've considered so far have been structured classes. The properties that defined them weren't arbitrary; instead they were derived from systematic relational properties of environmental events (e.g., matching, oddity, symmetry). Let's now examine some arbitrary classes, as created in an experiment by Vaughan (1988), who arbitrarily divided a group of photographic slides into two sets of 20 slides each. The slides were presented one at a time, and pigeons' pecks were reinforced given a slide from one set but not given a slide from the other. The pigeons learned to peck accordingly.

Once performances were fairly consistent, the correlation between slide sets and reinforcement was occasionally reversed. After several reversals, the pigeons began to switch their pecking from one slide set to the other after only some of the slides had been shown. In other words, the common contingencies arranged for the 20 slides in either set made them functionally equivalent, in the sense that once contingencies changed for just a few slides in that set, behavior changed appropriately for all of them. This functional equivalence emerged because the same consequences were arranged for pecks in the presence of all the stimuli within a set.

This procedure created two arbitrary discriminated operants, pecks to one slide set and pecks to the other, by arranging common contingencies for the members within each set. The correlation with reinforced pecks was the only thing distinguishing one class from the other. We probably think of classes differently when their constituents are arbitrary, as in Vaughan's study, than when the constituents have some natural coherence (e.g., when they're selected from a narrow range of colors or are all instances of some natural category). But **common contingencies select the members of operant classes.** A rat's left paw presses and right paw presses and both paw presses are arbitrary too, until they share in the common contingencies of depressing the lever. We should view Vaughan's experiment as one more example of using common contingencies to create a discriminated operant; it differs from others mainly in the particular classes that it created. Whatever structure existed in the two slide sets was imposed by the two sets of common contingencies (reinforcement for one, extinction for the other).

One expression of the problem of how structure emerges from undifferentiated beginnings is familiar to psychologists in the form of William James' metaphorical description of a newborn's response to the world as "one great blooming, buzzing confusion" (James, 1890, p. 488). James appealed to the child's earliest discriminations as the basis for the child's organization of what started out as disorganized. We know a lot more about discrimination than was available to James. He didn't have the benefit of Thorndike's research on the consequences of behavior or Pavlov's on the signaling functions of stimuli or Skinner's on operant classes and three-term contingencies.

To tackle James' problem, we can start by considering how discriminated operants are created by contingencies. Operants, defined in terms of their response properties and the stimuli during which they occur, are selected by their consequences. But common consequences must be distinguished from common contingencies. We know that a single reinforcer can sometimes maintain two or more different response classes. For example, when a pigeon's left-key pecks produce food according to one interval schedule while its right-key pecks produce the same food according to a different interval schedule, left-key pecks and right-key pecks are different operants even though they produce the same consequence, food.

Operant contingencies make all members of an operant functionally equivalent; the Vaughan experiment demonstrated a totally arbitrary discriminated operant based only on common contingencies. Accounts of complex discrimination too often appeal to the relative contributions of the stimulus and the organism without including the contribution of contingencies (e.g., Fetterman, 1996). When the class members have no common physical features, any approach that looks to stimulus properties to define how such a class was formed must fail (cf. Lakoff, 1987). We must look instead to what created the class, and the only common feature that is inherent across all of its members is the common contingencies they enter into.

We've just argued that common contingencies can create the sometimes arbitrary functional classes called operants. But even when class members share physical properties, nonarbitrary functional classes may arise not because of direct effects of those shared properties but rather because, by virtue of those shared properties, all class members are necessarily involved in common contingencies. For example, running one's hand over a sphere differs from running it over a cube; only in the latter case does one encounter an edge. These natural contingencies may therefore be the basis for discriminating between spheres and cubes. In other words, what seem to be nonarbitrary natural categories may well be created in just the same way as arbitrary ones, over a lifetime of experience with the common contingencies they engender. Doesn't anything significant involve contingencies of some sort, and aren't there plenty of opportunities for these contingencies to bring order, in the form of discriminated operants, out of the "great blooming, buzzing confusion"?

Undoubtedly some aspects of operant selection must be constrained by properties of the sensory and motor systems and neural organizations that have been selected in phylogeny and then shaped in ontogeny. Skinner spoke of such constraints as the "natural lines of fracture along which behavior and environment actually break" (Skinner, 1935a, p. 40); "We divide behavior into hard and fast classes and are then surprised to find that the organism disregards the boundaries we have set" (Skinner, 1953, p. 94). Such contingencies are so pervasive that we must never discount them when we try to identify the origins of behavior. Behavior structure is determined by contingencies, but contingencies are features of environmental structure.

ADDENDUM 12A: ANIMAL COGNITION AND COGNITIVE MAPS

The field of animal cognition is concerned with what animals know. It addresses this problem by identifying events and relations that can be discriminated by different species. Studies of animal cognition are of special interest when they involve discriminations of complex relational properties of the environment. Examples include judgments of visual symmetry (e.g., Delius & Nowak, 1982); discriminations of numerousness (e.g., Davis & Pérusse, 1988); visual search (e.g., Blough, 1989); discriminative control by reinforcement contingencies or stimuli correlated with those contingencies (e.g., Washburn, Hopkins, & Rumbaugh, 1991); the organization of behavior within sequentially discriminated sequences (Terrace & Chen, 1991); and responding under the stimulus control of the organism's own behavior (e.g., Shimp, Sabulsky, & Childers, 1989), to mention just a few. This section presents only a highly selective sample.

Consider a pigeon who watches a clock hand projected on the center key of three pigeon keys (Neiworth & Rilling, 1987). The hand starts at vertical and rotates at a constant rate from vertical through 90°; then it disappears. A little later it reappears farther along, at 135° or 180°. The timing of its reappearance is either consistent or inconsistent with it having moved at a constant rate of rotation while invisible. After a trial consistent with a constant rotation rate, left-key pecks are reinforced; after an inconsistent one, right-key pecks are reinforced. Pigeons learned to discriminate trials consistent with a constant rotation rate from those in which the constant rate was violated even though the rotating stimulus was absent some of the time; this discrimination also transferred to new locations of reappearance of the clock hand, so it could not have been based on trial durations or specific locations of the clock hand. It therefore demonstrates visual tracking in the absence of the visual stimulus; such tracking is sometimes called imagery (cf. Chapter 28). Colloquially, we might say that the pigeon knew where the stimulus was even while it was invisible (a useful skill, as when an edible insect passes behind an obstruction and a bird awaits its emergence from the other side).

Accounts of animal cognition are often framed in terms of the structure of relevant stimuli. For example, if an organism discriminates among stimuli on the basis of some critical feature, an animal cognitivist might say that the organism represents the stimuli to itself in terms of that feature (e.g., in the clock hand example, that the pigeon represents or imagines its constant motion, or in the case of search for prey, that a predator compares potential targets with a search image).

Whenever local environments have different properties it is advantageous for an organism to be able to find its way from one to another. An environment with a rich and stable supply of food is preferable to one with a poor and variable supply; an environment with easily accessible food is preferable to one in which the same food exists but is less accessible; an environment with safe areas for breeding and for rearing offspring is preferable to one that is more dangerous; and so on. (The argument holds for most animal groups. It would take us too far afield also to consider the phylogenic contingencies that operated in the evolution of plants, but it is appropriate to note that many plants disperse their seeds; animals are often involved in that dispersal, as when bees pollinate flowers.)

Once some kind of orientation has emerged, natural selection is likely to sharpen it over phylogenic time. It is therefore not surprising that many animal species readily find their way around in the world. Some of their navigation is learned and some is unlearned. Gallistel (1990) provides a detailed treatment ranging from foraging in ants to echolocation by bats and choice of routes by chimpanzees, and from orientations based on simple stimulus dimensions such as gradients of odor or light to varieties that are functionally equivalent to celestial navigation. Organisms that leave their nesting areas must be able to return; the more accurately they can do so, the more widely they can forage. Organisms that store food over the winter must be able to locate the food later; the more variable their sites and the more sites they can keep track of, the less likely they'll be to lose what they've stored to their competitors (e.g., Balda, Kamil, & Grim, 1986). Organisms that evade predators must be able to locate escape routes; those that allow themselves to be pursued into blind alleys don't survive.

We considered some issues relevant to spatial orientation in discussing place learning versus response learning. The two kinds of behavior were distinguished mainly by whether stimuli outside the maze were available to the rat. If they were, the rat oriented itself within the larger stimulus complex; it learned places. If not, it mastered only specific turns within the maze; it learned responses. Additional complexities were introduced in other maze experiments (cf. Olton, 1979). For example, a rat sometimes chooses the shortest available route through a maze when another route preferred earlier is blocked, or sometimes takes appropriate shortcuts that have just been added to a maze even though it never traveled those shortcuts before. These outcomes provide the justification for speaking of *cognitive maps* (Tolman, 1948); the finding that organisms can locate an area even when approaching it from a new direction demonstrates that they learn spatial relations in addition to or perhaps instead of specific paths.

The problems of determining the environmental features to which organisms attend as they move from one place to another emerge on a grander scale in animal homing and migration. Wasps return to their nests, bees to their hives, salmon to their rivers of origin and birds to seasonal nesting grounds. Among the environmental features that may be important are landmarks, the location and movement of the sun and stars, polarized light, chemical gradients and magnetic fields, (e.g., Tinbergen, 1972; Walcott, Gould, & Kirschvink, 1979). In some cases organisms navigate in isolation to regions they've never visited before; in others, they do so in the company of other members of their species. Both cases must involve substantial phylogenic components, either with respect to important properties of environments or with respect to the social contingencies that lead to group migration, or sometimes both. With respect to long-distance migrations, in a phylogenic analog of shaping phylogenic contingencies may have selected individuals who could maintain orientation over very gradually longer journeys as the continents slowly drifted apart over geologic time, (Skinner, 1975). Large changes can accrue over millions of years if a transoceanic migration path becomes only an inch or so longer each year.

Chapter 13
Sources of Novel Behavior

> New *has a limited population of etymological relatives; it includes* novelty *and* innovate. *So also with* origin, *derived from the Latin* oriri, *to be born, which is related to* orient, *as the direction of the rising sun, and to* are, earnest *and* hormone. Emerge, *from the Latin* ex-, *out, plus* mergere, *to sink or immerse, implies a bringing forth or rising up;* emergence, *as an unforeseen occurrence, gives us emergency. A relevant cluster,* productivity, adduction, induction *and* educate, *is tied together by the common root* deuk-, *to lead or to draw out.*

Toward a Taxonomy of Novel Behavior
Reinforcement of Variations: Shaping and Fading
Emergence of New Responses: Higher Order Classes
  Equivalence Classes and Frames
Combining Classes: Adduction
  Serial Coordinations
  Coordinations in Parallel
  Joint Control
  Fluency and Teaching

Despite its initial support, Darwin's account of evolution in terms of natural selection soon encountered difficulty because of what was known about genetics in the nineteenth and early twentieth centuries (cf. Chapter 3). In particular, natural selection couldn't work without variations from which to select and it wasn't clear where those variations could come from. The problem was the source of the novel variations upon which the changes in species over evolutionary time depended. It was solved when data on mutations were added to what was known about genetics. Genetics was then seen as providing the transmission from one generation to the next and mutations as providing the variations. Even in this century the origin of novelty rather than natural selection itself has been a favorite target of Darwin's critics (Kirschner & Gerhart, 2005, p. 265).

The ontogenic selection of behavior by its consequences faced a similar challenge, especially with regard to the role of reinforcement in the acquisition of language by children. After the publication of *Verbal Behavior* by Skinner (1957), the linguist Noam Chomsky (1959) wrote a critique of the book in which he argued that Skinner could not account for the child's capacity to produce gram-

KEY TERMS: Novelty, Novel Behavior; Reinforcement of Variations; Emergent Behavior; Shaping and Fading; Higher-Order Classes; Adduction, Transfer of Function, Equivalence Classes, Mediated Generalization, Fluency, Joint Control; Multiple Causation.

matical sentences that the child could never have encountered before (cf. Chapter 25).

When we write a sentence, it is likely to differ from any others we've written. We could not have learned them individually because the number of ways in which words can combine into grammatical sentences is astronomical. This feature of language is called *productivity*, but language shares this property with nonverbal behavior. We can deal with novelty in terms of features that the new sentence shares with earlier ones; novel productions involve new combinations of already established syntactic and semantic classes (cf. Esper, 1973). New combinations have been important in evolution through natural selection (West-Eberhart, 2003). "Novelty in an organism's physiology, anatomy, or behavior arises mostly by the use of conserved processes in new combinations, at different times, and in different places and amounts, than by the invention of new processes" (Kirschner & Gerhart, 2005, p. 35).

Because children could not possibly have made contact with all possible sentences in the language they were mastering, Chomsky appealed to the *poverty of the stimulus* (Chomsky, 1959; Crain, 1991). Based their limited samples, how could children learn the difference between grammatical and ungrammatical sentences? Given that some ungrammatical forms were so unusual that children were hardly ever likely to encounter them, how could they know that they shouldn't use them? Chomsky concluded that children had to start out with some built-in language acquisition device or they would be unable to learn their native languages.

But a pigeon doesn't need to see all other wavelengths of a visual stimulus before learning to peck green but not red (Catania, 2003b). Neither does a child need to hear ungrammatical sentences before learning past and present tenses or active and passive voices. The selection of operant classes by their consequences does not depend on extensive sampling of negative instances. Evolutionary analogies are again helpful: populations are not selected from pools exposed to all possible environments, and not all variations are included in the pools upon which selection operates. It remains reasonable to consider structural constraints on what gets selected, but those constraints do not erase the genealogies of species or behavior classes.

TOWARD A TAXONOMY OF NOVEL BEHAVIOR

Where novel behavior comes from has been a theme in many examples we've considered so far in this text. Ontogenic selection must have variations on which to do its work. Thus, identifying the sources of novelty and variation is of special interest. In the history of the psychology of learning, many battles have been fought over whether some approaches can give better accounts of novel behavior than others, or even whether some can give any satisfactory accounts at all. Creativity in art or music or invention must deal not only with sources of inspiration but also with the skills that allow someone who is generating lots of ideas to tell the difference between the good ones and the bad ones. The issue has been particularly important in the area of verbal behavior, as we'll see when we get to learning with words.

The problem, however, is not whether we can identify ways of producing novel behavior. We can identify many varieties. The difficulty is in deciding how they should be organized and whether we've exhausted the possibilities. We'll sample a few in this chapter, but it is reasonable to assume that our taxonomy of sources of novel behavior is incomplete. Table 13-1 summarizes some sources of novel behavior, but it is probably best regarded as a work in progress.

We've already discussed some of these categories. Shaping and fading were considered in detail in Chapters 9 and 11 respectively. The direct reinforcement of novelty was considered in the context of outlining the nature of operant classes at the end of Chapter 10. Equivalence classes and higher-order classes were covered in connection with stimulus control in Chapter 12. From early on, examples of the coming together of separate response classes have been noted, especially in those involving response sequences, but the implications of such combinations were not treated explicitly. Now is the time to do so.

Table 13-1 Sources of Novel Behavior

Source	*Definition*
Differential reinforcement in shaping, fading and the creation of operants	Creating new behavior by differentially reinforcing approximations to a new response or a new stimulus class; creating new behavior by reinforcing novel instances defined relative to the populations of which they are members
Emergence based on higher-order classes	Setting the occasion for new instances of the members of higher-order classes, including extensions of equivalence classes and frames
Combining behavior classes: Adduction, transfer of function and their variations	Bringing the properties or members of different classes together in new ways; combining multiple sources of behavior so that new functional classes emerge

REINFORCEMENT OF VARIATIONS: SHAPING AND FADING

Raising the force of a response has been one of our shaping examples. If a rat presses the lever too lightly and that response isn't reinforced, a harder press next time might work. That means that along a quantitative dimension like response force something more may happen during shaping than just the strengthening of particular forces. Whenever a response is not quite forceful enough, a more forceful second response may be reinforced. Several instances of successfully pressing harder after an unreinforced press may shape up sequences of progressively harder presses, so that this kind of sequencing itself becomes a functional operant class. Once it is in place, shaping may progress more rapidly than if it depended solely on the relative likelihoods of different response forces. Instead of spreading to responding at both higher and lower forces, the effects of reinforcers would become directional, spreading more to higher than to lower forces. We might call this *facilitated variation*, to parallel similar processes that have been proposed at the phylogenic level (Kirschner & Gerhart, 2005).

Presumably a novice shaper who was worrying about extinction could accidentally move things in the other direction. If a forceful reinforced response happens to be followed by several weaker ones and the shaper quickly drops the reinforcement criterion to lower levels to avoid extinction, sequences of forceful presses followed by weaker ones might be reinforced. Once such sequences have been strengthened the progress of shaping may be slowed. This may not be an issue in fading procedures, because in that case the sequence of stimulus changes is under the control of the experimenter. In fading, the main problem is setting the pace for changing the relevant stimuli rather than setting a criterion for reinforcement.

Reinforcing variants that successively approximate some target response is a special case of reinforcing any variants drawn from a population or of reinforcing some subset, as in arranging differential reinforcement within the boundaries of a designated response class. Another special case is reinforcing variants meeting some criterion for variability or stereotypy. We've already seen that variability itself is a property of behavior for which contingencies can be arranged (e.g., Neuringer, 1986; Neuringer, 2004; Page & Neuringer, 1985; Pryor, Haag, & O'Reilly, 1969).

In one group of studies reinforcers were contingent on rats' sequences of presses on left (L) and right (R) levers (Neuringer, 2002, 2004). The differential reinforcement of a specific five-response sequence such as RLLRL was ordinarily unsuccessful, but the sequence was mastered by rats for which reinforcers had first depended on producing variable sequences. For example, each L and R sequence was compared with the sequences in the last N trials and was reinforced only if it differed from all of them. Rats with a history of reinforced variability in sequences of L and R lever presses were also more likely to explore when exposed to new environments. Similar findings have been obtained with the generation of random-number sequences by humans. In these cases, **no individual response can have the property on which the contingencies depend; the property depends on its relation to a population of earlier responses.**

When behavior is more variable, it provides more variants that may be reinforced and therefore makes shaping quicker and easier. Here too we can identify phylogenic parallels. The *genotype* is the genetic makeup of an organism, whereas the *phenotype* is the organism as it has developed within its environment. For example, two species of fish that live in highly similar lakes may have developed highly similar phenotypes in adapting to the common features of their environments, but the genotype of one might be far more variable than that of the other. Research in conservation biology has shown that the broader range of genotypic variations available to the former species make it more likely than the latter to survive if their lake environments change drastically (e.g., Reed, 2007; Vellend, 2006).

Variability on the stimulus side may also be relevant to fading procedures. We saw in Chapter 12 that interference in stimulus-control tasks can be reduced by using stimulus classes containing many members instead of limiting those classes to a very small number (e.g., just green versus red). Fading can be arranged along single stimulus dimensions, as when colors are faded out as forms are faded in. With photographs or other stimulus sets that allow categorization in many different ways, however, fading can be arranged along multiple dimensions (e.g., Verplanck, 1962/1992). Fading is procedurally demanding and it remains to be seen whether the greater variability allowed by larger and more diverse stimulus classes can produce dividends relative to fading that uses more limited classes, as it has in matching and related tasks (Wright, 2012).

EMERGENCE OF NEW RESPONSES: HIGHER-ORDER CLASSES

If individual lever presses make up an operant class and sequences of successively more forceful presses also become a functional unit, that functional class is a higher-order unit that contains individual presses within it. We introduced higher-order classes in our treatment of self-injurious behavior in Chapter 9, and Chapter 12 considered a higher-order class consisting of a child's imitations of the actions modeled by a puppet. This lever-press example suggests how pervasive higher-order classes can be.

These classes are usually described in terms of the relations that define them. Imitation is defined by correspondences between the behavior of a model and the behavior of the one who is imitating. Instruction following is defined by correspondences between verbal and nonverbal behavior. Whenever the contingencies maintaining a higher-order class differ from those maintaining one or more of its subclasses, we must ask which contingencies will prevail.

We can tell whether we've created a higher-order class if members of the class emerge that have never occurred before. New instances arise as special cases of the class, as when, in generalized imitation, a child successfully imitates some action the child has never seen before. Another example is instruction following, as when someone obeys an order that has never before been given or acted on.

Whether such novel instances decrease or persist long after they've been excluded from the reinforcement contingencies that continue for the others is an experimental question. If they decrease, we've demonstrated that we can pull them apart from the other members of the higher-order class. If they don't, the class may have some properties that maintain its integrity. These might be structural. For example, if a rat presses a counterweighted lever and we reinforce all presses between 10 and 200 grams, we shouldn't expect to see much happen if we exclude just presses between 87 and 91 grams from reinforcement. Presses at 88 or 90 grams are unlikely to extinguish.

As long as a higher-order class maintains its integrity, its subclasses are maintained along with the other members even if they're not consistently involved in the contingencies that maintain the others. When that happens, the lower-order subclasses will seem insensitive to the changed contingencies that have been arranged for them.

The overlap of the subclasses with other classes that share in other contingencies may also keep classes hanging together in spite of contingencies that might be expected to break them up. When we arrange new contingencies for subclasses of higher-order classes, we ordinarily expect those subclasses eventually to be differentiated from the others. But if the subclasses overlap in various ways with other classes, especially in interacting social contingencies, that may not happen. For example, following rules in high school is one thing; for students locked into social groups, however, doing what their friends suggest is something else.

Compartments place a crucial role in embryonic development at the level of biology (cf. Chapter 3), and higher-order classes may function as compartments at the level of behavior. Modularity in biological systems allows parts to be modified and become specialized independently of other parts (Carroll, 2005, p. 195; Kirschner & Gerhart, 2005, pp. 216-217); perhaps higher-order operants are the modules of behavior.

Equivalence Classes and Frames

Equivalence classes constitute one more example of higher-order classes. Once I've learned that a spoken *A* is equivalent to a written uppercase *A* or lowercase *a* and that the same is true for spoken and uppercase and lowercase *B*'s and *C*'s and all of the other letters of the alphabet, then for me each of the letters is a higher-order class with its various forms as the members.

Training with many specific instances may sometimes be a sufficient prerequisite for higher-order or generalized classes. For example, training with many symmetry problems may produce generalized symmetry, training with many transitivity problems may produce generalized transitivity, and so on. Such generalized classes have been called *relational frames* (Hayes, 1994; Hayes, Barnes-Holmes, & Roche, 2001) by extension from frames introduced by Skinner (1957) in his analysis of verbal behavior. They have been applied to therapies and emotional behavior (e.g., Friman, Hayes, & Wilson, 1998; Hayes & Wilson, 1994; Wilson & Hayes, 2000). As examples of higher-order classes, relational frames are relatives of equivalence classes, in the sense that they are defined by relations rather than by the exchangeability of members. For example, *equivalent to* is defined by the properties of reflexivity, symmetry, and transitivity, but only some of these properties hold for other frames, such as *opposite* or *greater than* (cf. Chapter 12).

But overlapping frames need not be consistent, as when true-false relations are pitted against value judgments. You may believe that something is true but you may still not like it. The overlap between these dimensions can matter a lot. For example, if in preparing for an exam you study what you need to know you'll probably be better off than if you study only what you like. Competing frames can be interpreted in terms of how contingencies at different levels might either support or conflict with each other, as when the contingencies maintaining a higher-order class differ from those maintaining specific instances.

In discussing frames it is easy to get caught up in structural descriptions of the relevant response classes, but defining frames in terms of networks of coordinations does not explain them. It is important to attend instead to the contingencies that create and maintain them. By itself the three-term contingency cannot accommodate nested and higher-order cases, so analyses of these more complex contingencies must move on from three terms to four or five or even more (Fields & Verhave, 1987; Saunders & Green, 1999; Sidman, 2000).

Our examples of equivalence in Chapter 12 mostly involved three-member classes, as when teaching AB matches, BC matches and identity matches creates two classes, $A_1 - B_1 - C_1$ and $A_2 - B_2 - C_2$. When some class members are connected via intermediate members rather than directly, as for A and C in this example, the intermediate member, in this case B, is called a *node* (Fields *et al.*, 1990). Expanding our two classes with CD and DE matches adds two more nodes, in $A_1 - B_1 - C_1 - D_1 - E_1$ and $A_2 - B_2 - C_2 - D_2 - E_2$. Here the AD and BE members are separated by two nodes and the AE members are separated by three. Training with AC, BC, CD and CE matches instead would also produce two five-member classes, but this time all the other members would be linked via the common single node at C. Increases in the number of nodes separating class members can slow the acquisition of equivalence classes and can affect their integrity (Fields *et al.*, 1995).

A further issue is the role played by verbal behavior in creating and maintaining these classes (e.g., Fields *et al.*, 2007). When a young child is performing unsuccessfully in an arbitrary matching task, it is sometimes enough simply to point at the sample and ask *Which one goes with this one?* to see an abrupt transition to accurate matching. In tasks like these, words may be involved in other ways than just as names for stimuli.

Given the possible role of verbal behavior in equivalences, can we say whether they are prerequisites for verbal behavior or verbal behavior is a prerequisite for them? On the one hand, verbal behavior includes so many correspondences—between things and their names, between words written or spoken, between things requested and things received, between letters in script or in type, and so on—that it is tempting to see equivalences as helping learners to master their languages. On the other hand, once a child can name a stimulus in a matching task, it's tempting to invoke the name as a mediator. I've vacillated on this question over the years, but as we learn more about the development of naming in children and the expansion of its functions (e.g., Horne & Lowe, 1996; Randell & Remington, 1999), I'm increasingly persuaded that verbal behavior has to come first.

Equivalences are typically established with arbitrary stimuli, such as visual shapes or nonsense syllables (cf. Chapter 21). Meaningful materials can enhance acquisition (Fields *et al.*, 2012), unless the relations run up against other well-established classes. Among the Catholics and Protestants of Northern Ireland, for example, names and preferred colors and flags and favored sports differ across the two groups. Catholic and Protestant learners were able to master equivalence classes involving only neutral stimuli but not those in which Protestant and Catholic stimuli were linked by a neutral node. For example, if A in an AB match consisted of a Catholic stimulus such as Irish green and C in a BC match consisted of a Protestant stimulus such as an image of a marching orangeman, so that either class always included both Protestant and Catholic stimuli, acquisition was slow if successful at all and learners typically failed AC transitivity tests; these effects were stronger in older than in younger children and were often locked in by age 11 or so (McGlinchey & Keenan, 1997). These outcomes were somewhat less likely among learners who attended universities with a substantial mix of Catholics and Protestants in the student body.

We probably should not be surprised that existing classes established over an extensive social history may be difficult to override with the relatively weak contingencies available in a laboratory setting. In these communities these stimuli were familiar and already seen as belonging to one or another of two incompatible classes. Recall the distinction between discrimination and prejudice in Chapter 11. To the extent that such classes are created by actual contingencies, they may illustrate discrimination. To the extent that they are created by verbal environments, they may illustrate prejudice. Either way, they make the point that we should be interested not only in what it takes to create equivalence classes; it would be good also to know what it takes to break them down.

Some overlapping class memberships may be subtle but pervasive in their effects. For example, the direction words, *right* and *left*, share one term with another pair, *right* and *wrong*. To the question *Do I turn left here?* the reply *Correct!* is probably less likely to create trouble than the reply *Right!* Small wonder that so many are directionally challenged in Western culture. In many other languages (e.g., Japanese) there is no such overlap. It is no accident that among English speakers directional confusions are less likely in nonverbal tasks than in tasks that involve directional talk (Maki, 1979).

These effects of overlapping classes are perhaps even stronger in languages with Latin origins. In Italian, for example, *right* and *left* are *destra* and *sinistra*, where *sinistra* shares its roots with *sinister*. It's easy to see how people who are different in handedness might be regarded with suspicion in some cultures. Political affiliations on the right or the left introduce still other potential complications.

COMBINING CLASSES: ADDUCTION

By definition, membership in a class implies that members of the class have some functions in common. But equivalence class membership needn't imply equivalence in every possible way, so it is usually an experimental question whether a stimulus that has become a member of an equivalence class by virtue of matching procedures will also share additional functions with the other class members (e.g., Bones *et al.*, 2001). One source of interest in these classes is the possibility that the existing functions of some members will transfer to others. For example, suppose a child is taught the equivalence of the spoken word *go* and green traffic lights as one class and of the word *stop* and red traffic lights as another, and the child already goes at the word *go* and stops at the word *stop*. If the functions of the words transfer to the traffic lights, the child will go on green and stop on red without having to be explicitly taught.

This case involves novel combinations of stimulus and response classes. Superficially the child who obeys the traffic lights after this procedure seems to be conforming to a three-term contingency, including green and red traffic lights as the antecedents, going or stopping as the behavior, and what may happen upon going or stopping as the consequences. But this behavior did not arise from exposure to the contingencies. It originated in the combination of existing response classes. We'll find that many examples of novel behavior involve the novel combination of existing classes. You may never have seen a purple cow, but if you've learned color words and animal words you'll be able to name one when you see one. The novel coming together of separate operants has been called *adduction* (Andronis, Layng, & Goldiamond, 1997; Johnson & Layng, 1992). Novel combinations can be coordinated serially or concurrently. As we'll see, creating conditions in which classes combine readily has special relevance for teaching.

Serial Coordinations

Let's begin with an example that involves combining phonetic units (Berko, 1958). Children from 4 to 7 years old read illustrated sequences of sentences that introduced a nonsense word. The last sentence was incomplete and prompted a different grammatical form of the word. For example, a sequence prompting a plural was: "This is a wug. Now there is another one. There are two of them. There are two ______." The child who says *wugs*, never before having encountered that word, has given a novel response.

The regular plural in spoken English is formed by *-z* after voiced endings such as *b*, *d* or *v*, as in *dogs*, by *-s* after voiceless endings such as *p*, *t* or *f*, as in *cats*, and by *-es* after s or z endings, as in *houses*. At this age children usually produce standard plurals with *-z* and *-s* (most children responded *wugz* to the sample sequence above), but only about a third give the *-es* plural to novel words like *tass* and *gutch* and *nizz* even though almost all give *glasses* as the plural of *glass*. Aspects of the standard English plural develop successively. Children first learn specific words and word sequences and then begin to master higher-order units such as classes of plurals.

The voicing of the final consonant of the noun might make a difference to a listener (e.g., *cat* versus *cad*), but the voicing of a plural *-s* probably wouldn't be noticed unless a listener was attending carefully. In other words, contingencies established by the verbal community presumably operate to maintain the distinction between the voicing and unvoicing of the final consonants of nouns, but the carryover of voicing to a plural *-s* is more likely a product of some constraints other than such contingencies (e.g., aspects of phonology and articulation). Can we create a nonhuman analog of such serial coordination?

In a trial procedure, the relevant stimulus dimensions were form (circle versus triangle) and number (one versus two, which we'll here call singular versus plural); the first peck after a 4-s trial of A or B was reinforced if responding during the trial satisfied appropriate contingencies. As illustrated in Figure 13-1, left-key pecking was established during singular stimuli on both bottom keys and right-key pecking during plural stimuli on both keys, as in B on the right. Next, high and low response rates were differentially shaped on the top key; after this shaping, top-key pecking during singular or plural circles was reinforced if 5 or fewer pecks had been emitted during the trial, and top-key pecking during singular or plural triangles was reinforced if 9 or more pecks had been emitted during the trial, as in A on the left. The two conditions were then arranged serially, starting with the top key. As illustrated in the figure (A followed by B), the bottom keys became available when pecking in a top-key trial met the top-key rate criteria. Throughout all conditions, no rate contingencies were ever arranged for bottom-key pecking.

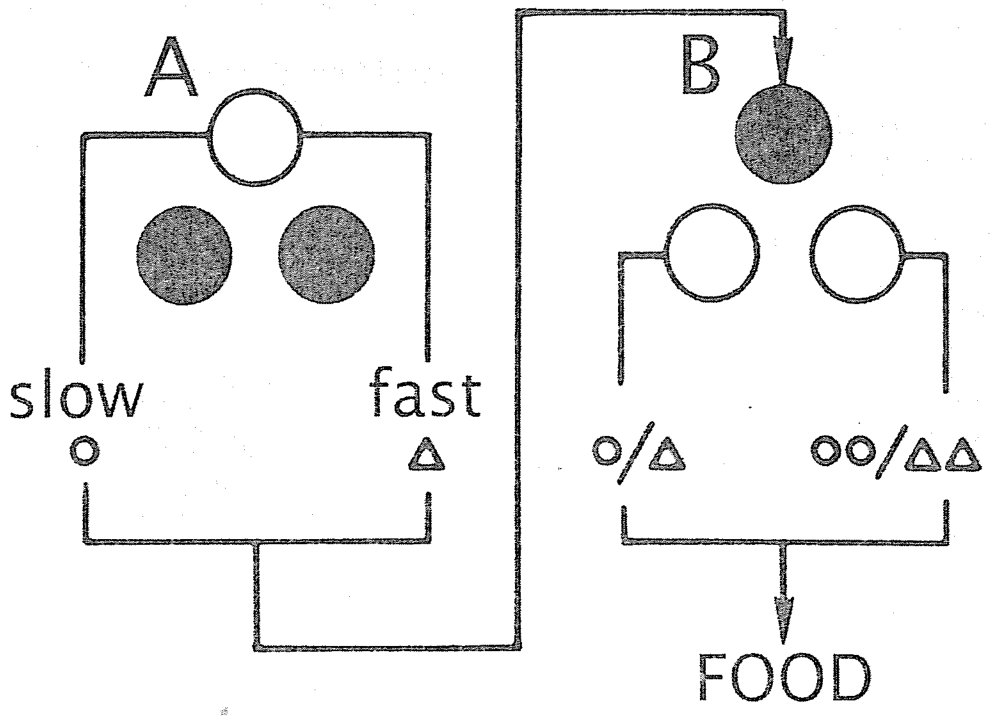

Figure 13-1 Schematic of a procedure demonstrating serial coordination. A pigeon saw one or two circles or triangles displayed on keys in a three-key chamber; filled circles represent dark keys. Training began with trials in which left pecks were reinforced during one form (singular) on each bottom key and right pecks during two forms (plural), as in B. Then, in trials on the top key, slow pecking was reinforced during circles on that key and fast pecking during triangles, as in A. Finally A and B were put together: when top-key rates met their high-rate or low-rate criteria during A, they produced B. The question was whether the top-key rate differences would appear on the bottom keys, where they had never been differentially reinforced. (Adapted from Catania & Cerutti, 1986, Figure 8)

Within about two weeks after A and B were combined, rate differences corresponding to those on the top key emerged on the bottom keys, as illustrated in Figure 13-2. Lower rates show as shifts of the distributions toward the left; higher rates show as shifts toward the right. The differences were smaller in magnitude than those on the top key, but they occurred even though differential reinforcement of bottom-key responding depended on response location and not on response rate. Reinforcement was maintained during all sessions in which the A and B components were combined, but the bottom-key rate differences developed only gradually. This development cannot be attributed to differential reinforcement and occurred despite continued reinforcement of roughly equal rates on the bottom keys across all stimuli during the sessions immediately after A and B were combined. Just as voicing persists in the plural *–s* in human verbal behavior, differential rates persist in the plural responding in the pigeon analog. These are of course small coordinations, but they illustrate how change can occur when different response classes are combined and how class properties can constrain sequential behavior.

Coordinations in Parallel

Never before having met a green dog, we would probably still say something like *Look at that green dog* if we met one today. This comes about through a new combination of the well-learned words, *green* and *dog*, in the face of the novel combination of color and species. Such novel combinations are a commonplace feature of human verbal behavior.

One account of productivity is that novel behavior can emerge from the novel intersection of two or more properties of behavior, each

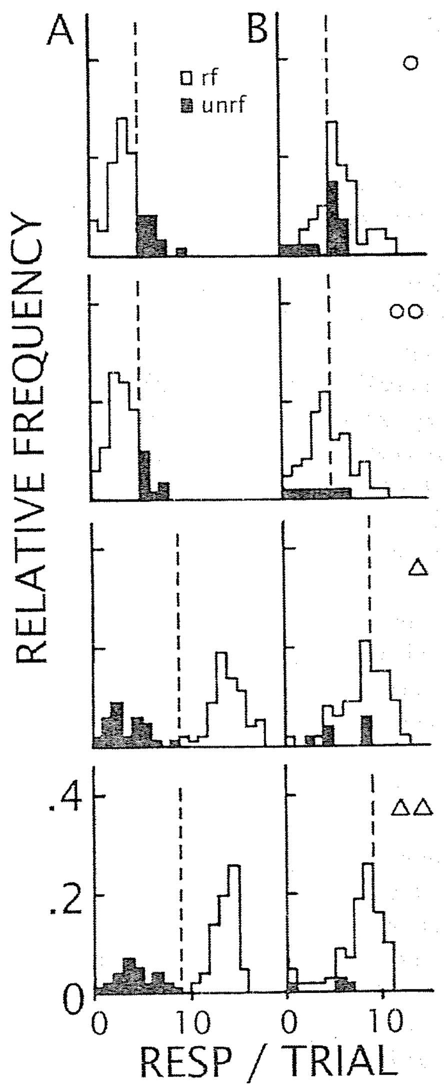

Figure 13-2 Distributions of responses per trial on top keys (A) and bottom keys (B) during singular or plural circles and triangles in the procedure diagrammed in Figure 13-1. Dashed lines show the top-key low-rate and high-rate cutoffs. Unfilled bars (rf) show data from trials in which pecking met the rate criteria in A and occurred on the appropriate key in B and therefore produced a reinforcer; filled bars (unrf) show data from trials that ended without a reinforcer. These data were obtained about two weeks after A and B were first combined. Rate differences corresponding to those on the top key emerged gradually on the bottom keys even though they had never been differentially reinforced. (Adapted from Catania & Cerutti, 1986, Figure 8)

determined by a different stimulus property (cf. Esper, 1973). This is an instance of multiple causation. We often speak of particular stimuli setting the occasion for particular responses. Strictly, discriminative stimuli do not occasion responses. Instead, particular stimulus properties occasion particular response properties. For example, the introductory example, *green dog*, involved a new combination of a color word and a species word, with each part of the verbal response depending on a different property of the nonverbal stimulus (cf. Skinner, 1957). Figure 13-3 illustrates a pigeon panel from experiments designed to provide a nonhuman analog of such relations (Catania & Cerutti, 1986; Catania, Ono, & de Souza, 2000).

Our objective was to separately establish control of response location by form and by number and of response rate by color, and then to combine these separate discriminations. For example, we asked whether a pigeon trained to peck the upper left key during singular circles and separately trained to peck rapidly during green would then peck rapidly on the upper left key given singular circles on green, a stimulus combination it had never seen before. In fact that didn't happen. Novel stimuli often disrupt behavior and they did so in this case. But control by all dimensions of the stimulus combinations was rapidly acquired once relevant contingencies were added for the new combinations.

The three stimulus and three response dimensions involved in this test for adduction are illustrated in Figure 13-3. Stimuli were presented on all four keys in 3-s trials, after which a peck on any key ended the trial. Stimulus number, singular versus plural, determined whether a peck on the top or bottom rows was reinforced. Stimulus form, circle versus plus, determined whether a peck on the left or right columns was reinforced. Stimulus color, green versus red, determined whether the peck was reinforced only after a high rate of pecking during the trial or only after a low rate.

The locations of the pigeon's key pecks conformed closely to the contingencies, and the differential reinforcement of respective high and low rates of pecking during green and red had gradually established two non-overlapping classes of responding, high rates of pecking during green and low rates during red, with distributions similar to those shown in Figure 13-2, Column A. After the shaping of high rates during green, the peck at the end of the trial was reinforced only if at least eight pecks had already occurred during the trial; otherwise an extended intertrial interval followed (timeout). Similarly, after the shaping of low rates during red, the peck at the end of the trial was reinforced only if fewer than three pecks had occurred during the trial.

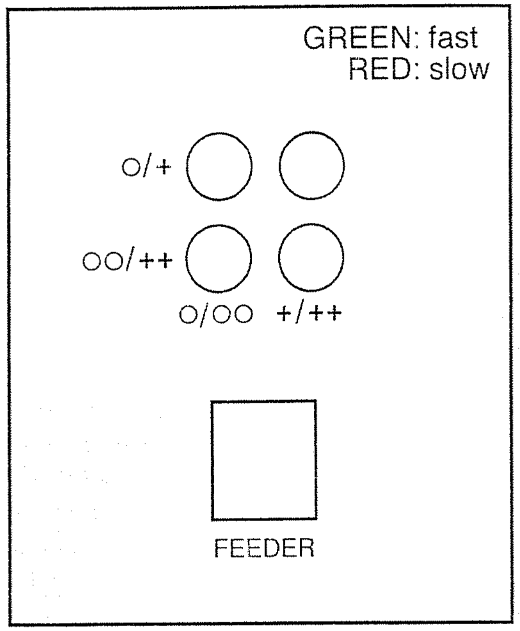

Figure 13-3 Four-key panel for a pigeon experiment demonstrating the control of multiple response dimensions by multiple stimulus dimensions. Identical stimuli were presented on each of the four keys within 3-s trials. Top-row pecking was reinforced given singular circle or plus and bottom-row pecking given plural circle or plus; left-column pecking was reinforced given circle and right-column pecking given plus; fast pecking was reinforced given green and slow pecking given red. For example, with plural plus on green on all four keys, the first bottom right peck was reinforced at the end of the trial if at least eight responses had occurred there during the trial; with singular circle on red, the first top left peck was reinforced if three or fewer pecks had occurred there during the trial (Adapted from Catania *et al.*, 2000, Figure 1).

This time the adduction test did not include extinction. Instead, for pigeons that were experimentally naive prior to training, incomplete contingencies were arranged for some stimulus combinations, so that differential reinforcement operated only for some of the response dimensions. All the contingencies of Figure 13-3 were in place, except that no rate contingencies were arranged during trials with singular circle on green or plural plus on red; let's call these the adduction stimuli. The question was whether an appropriate novel combination of response properties would emerge for these stimuli in the absence of differential reinforcement.

From the start, location contingencies worked for all keys and peck locations quickly came to conform to them. Rate differentiation proceeded more slowly, with response rate changes for the two adduction stimuli lagging behind those for which rate contingencies operated. Over about a month of daily sessions, however, the rates for the adduction stimuli caught up with the rest, as illustrated in Figure 13-4, even though no differential reinforcement operated for those two stimuli and even though intermediate rates of pecking had often been reinforced in the presence of the

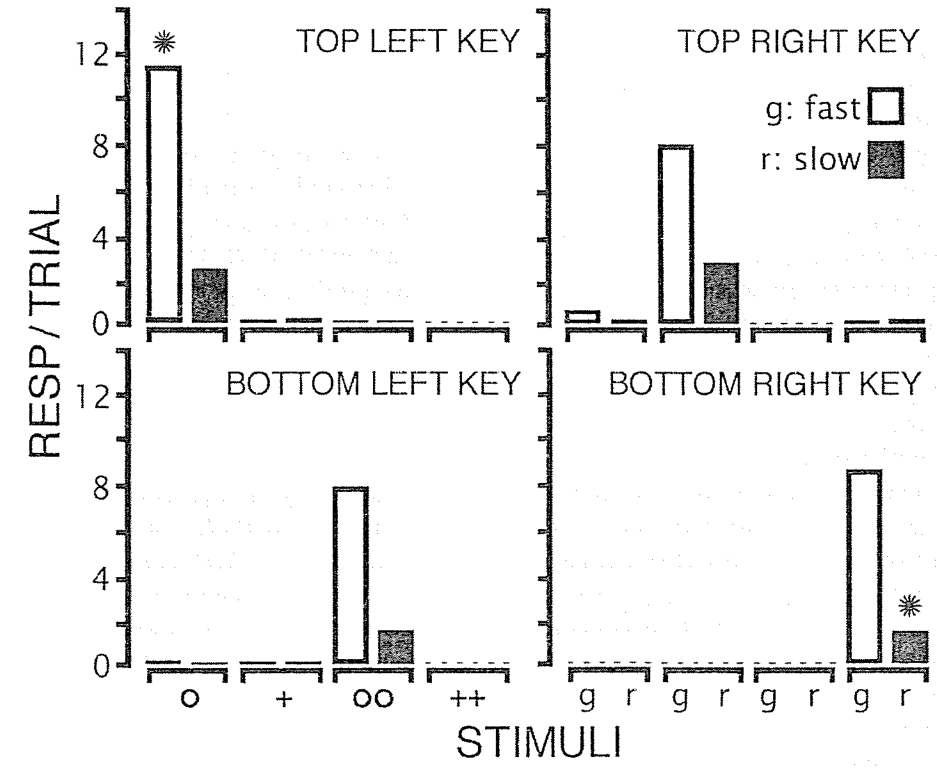

Figure 13-4 Responses per trial on each key and during each stimulus condition on the four-key panel diagrammed in Figure 13-3. Each graph represents one of the four keys and each x-axis bracket represents one of the four stimulus arrays, singular and plural circle and plus as labeled in the graph for the bottom left key. Within each bracket the bars represent responding during green (G) and red (R), as labeled in the graph for the bottom right key. Throughout training, rate contingencies were not applied to responding during singular circle on green and for plural plus on red (asterisks), so these rates had never been differentially reinforced given these stimuli. Responding conformed to the key and rate contingencies arranged for each possible stimulus combination, and responding given the two stimuli not participating in the rate contingency gradually came to conform to the differential rates maintained in the presence of the other stimuli.

adduction stimuli during the first weeks of the training sessions (Catania *et al.*, 2000). These patterns of responding were maintained for more than three months of subsequent sessions.

If even a pigeon can do this, should we be surprised if a child who has learned to say *one deer running, two sheep walking, one sheep running, two deer walking, two deer running* and *one sheep walking* will be able to properly say *one deer walking* or *two sheep running* when appropriate circumstances arise? Here too we have multiple causation: singular or plural determines the number, the creature determines the noun, and the fast or slow pace determines the gerund verb. (*Pace* developmental psycholinguistics and the poverty of the stimulus!)

But why didn't the novel combination in extinction work? Testing for stimulus control in extinction, as in determining generalization gradients (Chapter 11), has an extensive history in behavior analysis, but there may be problems with that kind of test. Consider the phylogenic contingencies. After its behavior during its first few encounters with a novel stimulus is unreinforced one organism stops engaging with the stimulus, but another persists in that behavior. The latter organism may be at a considerable disadvantage relative to the former, because in natural environments novel stimulus combinations are not necessarily correlated with corresponding combinations of contingencies. A food with some combination of properties such as color, size and shape may have become well-established as a reinforcer based on its nutritive properties, low toxicity, and so on. A new food may share some of these properties and is likely to be sampled when encountered, but if it is not particular nutritive and also has delayed toxic effects, the organism that persists in consuming it even though its consummatory behavior is no longer effectively reinforced may soon face both ontogenic and phylogenic extinction.

It is all too easy to make up evolutionary just-so stories, but perhaps phylogenic contingencies like these have selected organisms to be more like the former finicky sort than the latter obstinate sort. If so, many procedures for testing novel stimulus arrangements might be fundamentally flawed even if they include a history of intermittent reinforcement prior to extinction. If novel stimulus combinations lead organisms to behave in this way given extinction contingencies during tests of transfer or emergence, then novel stimulus combinations will quickly lose control and we may often incorrectly conclude that adduction does not occur. Perhaps when we fail to observe its emergence that is simply because we arranged testing procedures that operated against it. The emergence of novel combinations of response rates and locations in this procedure but not in extinction tests is at least consistent with these speculations.

These experiments do not explain behavior; they are demonstrations. To some extent they are nothing more than thought experiments realized in practice (Catania, 1980a). In fact, when I first became involved in discussing them they hardly seemed worth doing because it seemed obvious that they could be made to work. They needed to be done only to show those who found the analogies unconvincing that nonverbal organisms could indeed produce such behavior. Nevertheless, some things were learned along the way, such as the effects of extinction versus relaxed contingencies during adduction tests.

Joint Control

Now it's time to ask how adduction may work in our own behavior. *Joint control*, for example, occurs when common responses to different stimuli mediate judgments of equivalence or other relations, as when a child matches one arbitrary stimulus to another after having been taught to give the same name to each (Lowenkron, 1997). "Joint control is a discrete event, a change in stimulus control that occurs when a response topography evoked by one stimulus (e.g., the sample) and preserved by rehearsal, is emitted under the additional (and thus joint) control of a second stimulus (e.g., the comparison)" (Lowenkron, 1998, p. 332). Joint control, by the way, is about two stimuli jointly determining the behavior of single individual; it should not be confused with joint attention (see Chapter 11), which involves one individual attending to a stimulus that someone else is pointing to or otherwise indicating.

As an instance of joint control, Lowenkron (2004) offers the example of someone trying to find a number in a list, as in checking a set of winning lottery numbers. Say the number is *135476*, as read from the lottery ticket. If I were doing this, I'd probably repeat the number to myself as I ran my finger down the list. What else must happen before I can shout *I won!* if I find *135476* in the list? I'm already saying the number to myself but now I'm also saying it as I see it. Together, these two sources of saying the number have come together so they jointly control my saying it. But here is the critical part of the argument: before I can say *I won!* I must be able to distinguish this instance, in which there is a match, from all of the other cases in which the number I'm saying and the number I'm reading don't match.

In other words, for joint control to work we must be able to discriminate between instances in which we find a correspondence and those in which we do not. These correspondences have sometimes been said to exhibit *parity*, a match between two sources of behavior that arise from different sources (Palmer, 1996, 1998), as when babbling children react to having produced sounds that resemble those made by their caregivers. But on what sort of history could the discrimination between jointly controlled responding and responding that doesn't involve joint control be based?

Instances in which some response is occasioned by two different events at about the same time are sometimes notable. Skinner (1977a) remarked on hearing his teacher say a word at the same time he was reading it. Our attention is often captured by unusual events, including coincidences: we notice if someone says something at the same time we say it or think it; in the cocktail-party phenomenon, we start attending to someone else's conversation upon overhearing our own name in it; in our lottery-ticket example of joint control, we react to the coincidence between the number on our ticket and the winning number on the list.

We've learned to look at parents and teachers and others when they've called us by name, so attending upon hearing it has had consequences. It is reasonable to think of attending to our name when we hear it in the context of a new source, as in the cocktail party phenomenon, as a product of that history. This implies that pre-attentive behavior can be reinforced (some might prefer to invoke search images that govern our attention; this is no problem if they work like rehearsing a number to oneself). Aside from the source, how different is this from listening for your name as attendance is taken or searching for it on a list? And if this works for hearing your name, how about hearing a word at the same time you're reading it?

That too has a history. If a child is shown a letter in class and is asked by the teacher to say it, the child must attend to both the seen and the spoken letter together. There will be an enormous number of occasions for these sorts of combinations of events as the child matures. Maybe the more appropriate question is not why joint control becomes important but instead what happens if a child fails to become responsive to it? Could this be a component of some developmental disabilities? If it is absent, can it be taught?

Joint control involves discriminative responding to an aspect of our own behavior. In that respect it shares the properties of behavior called autoclitic that we'll encounter later (Chapter 22). We can no more pin down the stimulus dimensions on which this discrimination is based than we can say what defines redness or time or beauty (cf. Chapter 11). We do not appeal to physics before deciding whether some stimulus class can enter into a discriminated operant. Once we've seen a few instances, we can usually identify joint control when we see it. We'll return to issues like these when we get to verbal behavior in Part IV.

Fluency and Teaching

You can teach a child number facts until the child reaches a level of proficiency of 75% or 80%, and this might be accepted as a passing performance. If you could get the child up to 85 or 90 or maybe even 95%, you'd probably be pleased. It would be even nicer to get the child to 99% or better. Certainly those are the levels we'd like to see if these children ever become professionals in engineering or medicine or other disciplines that demand quantitative skills. But let's settle for 95%. Even so, it might be inappropriate to decide that your high-performing child was doing well enough to move on to something else. Suppose you start asking the child questions like *How much change will you get from ten dollars if you buy three toys that cost two dollars each?*, where the arithmetic needed to get the answer is well within this child's competence. If you were asking many questions like this and the child has not been taught to translate such word questions so they can be dealt with as arithmetic, you might see the child's performance drop below passable levels. So now you teach that skill to a proficiency level of 95% or better and you try the test again. The child still does not come close to 95%.

This is a failure of adduction. The two components, translating the word problem and doing the arithmetic, must combine if the child is to answer correctly, but here they mostly didn't. Just putting the two pieces together as above isn't good enough to produce adduction. The child must be taught to a level of *fluency*, the "combination of accuracy and speed that characterizes competent performance" (Binder, 1996, p. 164). The child must answer both accurately and rapidly, and this means continuing to practice beyond the level of high accuracy alone. After all, this is what we expect of fluent language. It matters in other skills too (e.g., Polson, Grabavac, & Parsons, 1997). Once the child becomes fluent in arithmetic and translating word problems, adduction typically follows (Binder, 1996; Johnson & Layng, 1992). Given the question above, the child will quickly reply *Four dollars*. Fluency training is a powerful educational tool, and students who have been taught this way come to excel in reasoning and putting facts together and the various other skills that make up critical thinking. When confronted with the pieces of larger problems, adduction comes easily to them (e.g., Binder, 1996; Greer & Ross, 2004; Johnson & Layng, 1992; Twyman, 1998).

One way of thinking about how this works is to assume that the correct answer in cases like these is given by the product of the probabilities of each component. If the child can do the arithmetic part at a 90% level and the word translation part at an 80% level, the chance of getting both right is 0.9 times 0.8, which equals 0.72. Even with both at 90%, the probability works out to be 0.81. That suggests that the child who starts out at 99% and 99% should end up with an accuracy of about 0.98 on the combined problem. But that doesn't happen. In cases like these, children do far worse than the combined probabilities suggest they should. Why is the overtraining for both speed and accuracy necessary?

One possibility is that when two reinforced classes of behavior come together, each has an inhibiting effect on the other. If so, the function of the overtraining for speed and accuracy may bring each class back up to the level it would have reached in the absence of the other. It would also make them stronger, in the sense of being more resistant to change than other classes that had not been learned to a fluency criterion (cf. Nevin, 1992, and Chapter 5). This is consistent with what we know about accuracy in other settings. For example, correct responses and errors in a matching or oddity problem will ordinarily differ in strength, and introducing reinforcers for other responses will typically reduce the errors more than the correct responses (Catania & Dobson, 1972). Paradoxically, instead of competing with or degrading the matching or oddity performance, introducing some other concurrently reinforced response class makes the performance more accurate. We'll explore related interactions among response classes when we consider reinforcement schedules in Chapter 16.

Fluency is about the behavior of the student rather than about the behavior of the teacher. It therefore has special relevance for educational practices. Clearly it cannot be managed by teachers working with large classes and for some well-defined skills, such as mastering basic arithmetic facts, computers may function better as teachers than humans. Teaching programs built around fluency have become increasingly sophisticated and we're learning more about how it works (Binder, 1996). The foundations upon which it is based have been around for a long time.

During my senior year at Columbia, I took a seminar on teaching taught by Fred Keller. Educational applications were in the air. Skinner had given a lecture on teaching machines at Barnard College. The power of computers was not yet available, so books with special designs or machines that could successively present questions and feedback were the alternatives to what later became programmed instruction (Barrett, 2002; Skinner, 1968). It would be a while before those early efforts paid dividends in the precision teaching of Ogden Lindsley (Potts, Eshleman, & Cooper, 1993), Fred Keller's paced instruction (Keller, 1968; Twyman, 1998), and other innovations (e.g., Goldiamond & Thompson, 2002; Greer & McDonough, 1999; Greer & Ross, 2004).

I had been an only child for thirteen years when I acquired a brother. Robert was about eight by the time I became a student in Keller's seminar and as my course project I took on teaching him some algebra, and the distributive law in particular: $a(b+c) = ab+ac$. One thing I learned was that it mattered to teach the law in as many different ways as possible: solving specific number problems, working with unknowns, saying what the law meant. By the time the semester ended Robert understood the distributive law, in the sense of being able to work with it in different ways. I cannot say with any certainty whether the project had any long-term effects, but his subsequent career has consistently involved numbers and mathematics and computer programs.

When I left Columbia for the graduate program at Harvard, one of my main interests was working on programmed instruction and related educational applications. But I soon became caught up in the pigeon laboratory, so that interest was deferred until I began teaching my own courses in a faculty position some years later. When I eventually arrived at UMBC I was fortunate that my colleague Eliot Shimoff had already been applying what we knew to education. We collaborated in both the laboratory and the classroom, so that the line between teaching and research was sometimes blurred (e.g., Shimoff & Catania, 1995, 2001). Various examples involving computer programs and class experiments and demonstrations appear in this text. That blurring is probably as it should be, for what is the point of research if you don't learn from your colleagues and teach them in turn about what you've learned?

This has been said before by others and I'll say it here and again later. For too long, educational systems have focused on what the teacher does instead of what the student does. **What the student does is what the student learns.** Students are not vessels into which knowledge is to be poured. What a teacher does matters little if no contingencies are placed on the behavior of students. Too many educational innovations simply change the curriculum. Yet changes in what a student is supposed to learn are largely irrelevant to changing how the student is taught.

Students should be able to learn in many different ways, but there is no evidence for significant effects of individual learning styles (Pashler *et al.*, 2008). The student who learns by rote and the one who learns by discovery and the one who learns by understanding will each learn something, but the one who learns in all these ways and others besides will learn more. The more time students spend with the material to be learned, the more they learn. Whatever a teacher can do to get students to put in more course time—exams, reading assignments, class discussions, papers, online exercises, lab projects, computer simulations—is likely to be a good thing, provided the contingencies don't become aversive.

I was once surprised when a graduate student complained that I hadn't provided outlines of the topics I was discussing in class. When I pointed out that each chapter in the text already had an outline, the student wasn't satisfied. This student was having the same problem as those educators who think that the behavior of the teacher matters more than the behavior of the student. Students who produce their own outlines learn much more than those who depend on outlines produced by their teachers. PowerPoint presentations are popular, but students who depend on them are having things done for them rather than learning to do things themselves. That is really not what education is about. For that reason I was sorely tempted to omit the chapter outlines and key words in this edition; I thought that I might be doing no favors for my readers by including them.

In those days when programmed instruction seemed to offer so much promise, some concentrated on teaching one or another specific skill. The targets ranged from reading and basic arithmetic through university-level courses. But some had a different dream: an educational system constructed from simple units like letters and numbers to the reading of words and sentences, and from there to whole paragraphs, not simply at the level of reciting the contents but instead answering questions and expanding on them. Critical thinking would be introduced early on. From there the next step would be pages and sections and eventually whole chapters. These could be contained in a convenient teaching device called a book (today we'd prefer a portable electronic tablet that could hold many such devices). Students who reach this point have been weaned from their teachers. They have learned how to learn; they provide examples of learning set writ large. I hope that some of my readers qualify. Perhaps this is what Skinner intended when he wrote his oft-quoted sentence, "Education is what survives when what has been learnt has been forgotten" (Skinner, 1964/1980).

Chapter 14

Motivating Variables and Reinforcer Classes

Responses, depending on their relations to eliciting stimuli, consequences, establishing operations and discriminative stimuli, are variously said to be elicited, emitted, evoked *or* occasioned. *All four terms have Latin roots. The first three share a prefix abbreviated from* ex-, *out:* elicit, *derived from* laqueus, *noose or snare, is related to* delight *and* latch; emit, *derived from* mittere, *to send or let go, is related to* omit *and* intermittent; evoke, *derived from* vocare, *to call, is related to* vocal, evoke, *and* invoke. Occasion, *derived from the prefix* ob-, *against, and* cadere, *to fall, is related to* case, accident, chance *and* coincidence.

Assessing Reinforcers
Reinforcer Classes and Reinforcer-Specific Effects
Conditioned or Conditional Reinforcement
Pseudo-Reinforcement
  Bribes
  Intrinsic Reinforcers and the Hidden Costs of Reward
Addendum A: Motivating Events in Escape and Avoidance

The semester after my class on discrimination with Schoenfeld (Chapter 11), I took his next course in the Experimental Psychology sequence. It was called *Motivation*. The text in my first semester had been Woodworth's *Experimental Psychology* (1938). In my second semester we used two texts: Skinner's *Behavior of Organisms* (1938) and Hull's *Principles of Behavior* (1943). Motivation was included in these books, but that was just one piece of their coverages of reinforcement and related topics. This too was a laboratory course, but this time the experiments were with rats rather than humans. One involved arranging schedules of what were then called conditioned reinforcers.

Courses with this title now would probably cover the topic by providing a systematic survey of motivational systems—hunger, thirst, sex, social drives, and so on—at some times concentrating on physiological mechanisms and at others perhaps speculating on the evolutionary or social precursors that might have created them. The fields are so broad that the coverage could hardly be exhaustive. For example, being able to breathe involves another such system, but it is seldom given much attention in the typical motivation course; contrary to common assumptions, by the way, there is scant evidence that breathing is triggered by buildup of carbon dioxide (Provine, Tate, & Geldmacher, 1987).

KEY TERMS: Motivational or Establishing Operation; Reinforcer Class or Consequence Class, Reinforcer-Specific Effects; Motivation and Emotional Effects.

When I took my course, we considered the meaning of hunger in the context of discussions of a hypothetical caveman who caught and caged a small animal and then began to observe its eating habits. I worried over that treatment of the word *hunger*. It seemed so much a term for a private felt state that even with Schoenfeld's help it took me a long time to acknowledge the significance of the argument for how the term must have evolved from observations of behavior (Catania, 1997). The point was that the term could not have acquired any function in everyday language except to the extent that it was based on public observations of the likelihood of eating as a function of time since the last meal and related variables.

Food is in our environment and to deal with how we behave with respect to it we should be looking there and not within for drives or hunger pangs or other inner causes. No one looks inside us to decide whether we're hungry. Though we think we know better than anyone else what's inside us, we often judge our own hunger by our behavior: *I didn't realize I was so hungry*, we say, while scarfing down a meal. Of course we're different when we're hungry than when we've been well-fed and it's worthwhile to find out how we're different. But hunger is not something that someone will someday pull out of organisms and stick under a microscope.

Here is a quotation from Skinner: "We say that an animal is hungry if, when we give it food, it eats.... In our everyday use of the word...we ordinarily attribute hunger to an animal only because it eats or because it exhibits behavior that we have frequently observed to be followed by eating" (Skinner, 1932, p. 22). He went on to say "the use of the word depends on the conditioning of the experimenter" (the emergence of the vocabulary of shaping and reinforcement would follow only some years later). Skinner had identified the problem as one of verbal behavior, and here he anticipated some critical arguments about the origins of our vocabulary of private events that we'll encounter later (Chapter 22).

ASSESSING REINFORCERS

The topic of motivational or establishing operations is about those events that make stimuli more or less effective as reinforcers or punishers. Chapter 2 gave them introductory consideration and the treatment of the relativity of reinforcement in Chapter 6 showed how the effectiveness of reinforcers depended on the probability of responding relative to the probability of the behavior they occasioned. We've been devoting attention to reinforcers for several chapters now, and that is as it should be. To get to know other people, one of the most important things you can look for is the reinforcers that maintain their behavior. Some are pretty obvious, as when we notice what people spend lots of time doing: are they mostly involved in reading, or music, or athletics, or social events, or politics? The questionnaires on values that are sometimes used to advise people about the careers for which they may be best suited often ask directly or indirectly how they like to spend their time. Another issue related to the relativity of reinforcers is whether the effectiveness of reinforcers depends on the response class upon which they are contingent (e.g., Reichle, Lindamood, & Sigafoos, 1986). Does it matter, for example, whether social behavior is maintained by social reinforcers or by nonsocial ones?

When I was close to officially becoming a major in psychology I discussed my options with Schoenfeld. One of the questions he asked was whether I wanted to go into the field because I thought I could use it to solve my own problems. I was puzzled by the question, though I knew that then, as now, clinical psychology was attractive to many students. I probably said something about being fascinated by the quantitative data I'd seen in his experimental classes and in Fred Keller's introductory course. I think Schoenfeld was probing whether my reinforcers were fundamentally about finding out how behavior works. Given where things went from there he must have approved of my answer.

In my own interactions with students after I began my own teaching, I often took Schoenfeld's lead and probed potential majors about why they wanted to enter the field. In giving students advice it made a difference to learn whether they mostly wanted to help others or to learn about behavior or to stay close to friends who had already chosen that direction or to make money or to have a chance to meet well-known scholars such as Skinner. Some were fairly single-minded about one or another but more often than not it was a mix. An inventory of the reinforcers that maintain someone's behavior provides a personality test of sorts. Your judgments of other people can change dramatically if you believe they're doing things on the basis of friendship or love and then you discover that they're instead doing it for other reasons: to advance professionally, to get even with someone, to gain financially, to arrive at a position of power, or to escape from aversive responsibilities. Reinforcers can be diagnostic.

Reinforcers can also be changeable. Behavioral pediatrics deals with both the common and the unusual behavior problems of children (Friman & Piazza, 2011). Sometimes such problems require interventions and sometimes they take care of themselves. For example, a parent might be seriously concerned with a child's persistent thumb-sucking when suddenly it drops out. This might happen as a result of emerging social reinforcers as the child gets older: at some point the social consequences of thumb-sucking may become more important than the oral sensory ones.

Identifying reinforcers is crucial to the applications of behavior analysis. Descriptive analysis is concerned with methods for assessing and comparing reinforcers (e.g., Vollmer, Borrero, & Wright, 2001). For example, one reinforcer may be masked by another that is more potent, so its effectiveness must be assessed not only when the more potent reinforcer is available but also when it is unavailable (Francisco, Borrero, & Sy, 2008). A major component of behavioral economics is assessing the relative potency of different reinforcers in terms of the behavior they maintain; a measure called *value* summarizes that potency.

Negative reinforcement should not be overlooked (Flanagan, Goldiamond, & Azrin, 1958). For example, aggressive or fearful behavior in dogs and other pets is often interpreted as elicited or emotional behavior, but it may be more appropriately viewed as operant behavior reinforced by keeping aversive events at a distance. A growling dog is much less likely to be approached than a placid one. A fearful cat is more likely to run and hide when a stranger visits than one that is more relaxed.

Rosales-Ruiz and colleagues (Rentfro, 2010, Nov/Dec; Rosales-Ruiz, 2011) exploited contingencies like these to shape calmer behavior in dogs and other animals whose behavior included aggression, excessive barking and other topographies usually characterized as emotional. With an aggressive dog, for example, another dog is gradually brought closer until aggressive behavior begins. The intruder dog is withdrawn only when the target dog shows reduced aggression. Soon the intruder dog is again brought closer and the procedure is repeated. Gradually the intruder approaches come closer and the criterion reduction in aggressive responses becomes stricter. Aggressive dogs become friendly ones. These dramatic changes in behavior can be shaped in a few days or even hours and can be maintained by similar contingencies. The distancing that shaped these changes in behavior involved negative reinforcement: the approach of the intruder was aversive, and its withdrawal reinforced behavior that was more acceptable to the pet owners.

Many reinforcers have obvious biological significance and presumably became so through evolutionary contingencies: e.g., food, water, sexual contact, air. Others probably emerged as byproducts of those motivational systems. These are sometimes called addictive. Some criteria exist for classifying some reinforcer dependencies as addictions; in Chapter 17 we'll consider the development of drug tolerance and some aspects of drug dependencies. But the term is often used more widely. As a long-distance runner, can you get addicted to endorphin highs? Can you become addicted to a person? To a political stance or a religion? To dependence or co-dependence? Aren't we already addicted to food and liquids? What then do we say if food is ingested but then rejected, as in bulimia, or is no longer reinforcing and perhaps even becomes aversive, as in anorexia? In addiction overconsumption can be fatal, but these cases are just the opposite.

These are all questions about reinforcers. It can be useful to know how they become established, as in the transition from casual to habitual drug use. It can also be useful to identify the components of such reinforcers. For example, cigarettes are a vehicle for nicotine, but nicotine by itself is not addictive (Killeen, 2011); people don't get hooked on nicotine patches. Cigarette manufacturers may have exploited the interaction of nicotine with other products of smoking as a route to addiction. Cigarette smoking is relevant to other topics we'll consider later. Self-control situations, for example, involve an immediate consequence pitted against an opposed later one, as when the immediate effects of inhaling cigarette smoke are pitted against remote effects such as eventual lung disease (cf. Chapter 16). In abstinence programs, smokers must commit to nonsmoking, and what we know about verbal behavior suggests that they are more likely to be successful if their verbal commitment has been shaped as a result of their own talk about smoking than if it has mainly grown out of instructions given by others (see Chapter 23 on verbal shaping and verbal governance). Where verbal behavior comes from sometimes may be as important as what it says.

REINFORCER CLASSES AND REINFORCER-SPECIFIC EFFECTS

In three-term contingencies, the first term corresponds to motivating or establishing operations and the stimulus classes that define discriminated operants. The second corresponds to the response classes called operants. The third corresponds to consequences and these too are classes of events. (Cuvo, 2000). Successive reinforcers arranged in experimental settings may be similar but they are not identical. Pellets produced by a rat's lever presses will never be totally uniform. Pieces of grain in a pigeon's feeder will always differ somewhat in shape and color. A parent's reinforcing hugs or smiles or positive comments will inevitably vary from one instance to the next.

In our earlier discussions of higher-order classes, one example was behavior that might produce attention and thereby maintain the self-injurious behavior (SIB) of children with severe developmental disabilities. Because it shares its consequences with other responses such as shouting obscenities or throwing things, the self-injurious behavior may be part of a larger class, attention-getting behavior. Within this class some types of responses may be more probable than others or may be differently available in different settings (Lalli *et al.*, 1995). For example, a child might be more likely to engage in self-injury if nothing to throw is close at hand, or the child might be more likely to shout obscenities given one audience than given another. Nevertheless, their membership in a common class makes it likely that these responses will often vary together.

But what if attention from staff members on this child's hospital unit does not function like attention from the child's mother when she visits? If we find that one kind of attention cannot substitute for the other, we might best treat attention from these two different sources as two separate reinforcer classes. This is important to know, because the behavior shaped by attention from the staff during therapeutic interventions on the unit may be incompatible with the behavior that has been shaped by the mother's attention at home. An effective treatment program must shape the mother's behavior as well as the child's or the treatment gains realized on the unit will be lost soon after the child's discharge.

Recall that operant classes are created by common consequences (cf. Chapter 9). These consequences are the glue that holds the classes together. But what then are the implications of using a common reinforcer across the many variations of a complex task. In a typical matching study, for example, all correct responses, whether to one comparison stimulus or the other, produce the same reinforcer. If a reinforcer enters into a single functional class along with the stimuli and the responses of an arbitrary-matching task, using this procedure may be a mistake. While the contingencies may work to separate the different matching classes, such as green pecks given a square sample and red pecks given a circle sample, the common reinforcer class may work instead to keep them together, just as a common reinforcer keeps all of the forms of pressing a lever together in an operant class.

The problem can be addressed by arranging different outcomes correlated with different classes of discriminative stimuli and responses. Acquisition of stimulus-control tasks such as matching-to-sample is typically enhanced when this is done (Urcuioli, 1985, 1991, 2005; Urcuioli & DeMarse, 1994). Some studies have also compared effects of differential outcomes with effects of differential responses to the sample, but the latter are issues about attention to the samples and not about reinforcer classes (cf. Zentall & Sherburne, 1994)

Studies with nonhumans (e.g., Dube *et al.*, 1993; Fedorchak & Bolles, 1986; Miyashita, Nakajima, & Imada, 2000) have been extended to arbitrary matching by children (Dube & McIlvane, 1995; Dube *et al.*, 1987; Dube *et al.*, 1989; Pilgrim, Jackson, & Galizio, 2000; Schomer, 2002). Even with correction procedures children may learn slowly in these tasks. If the tasks are modified so that responses from the different problem classes each produce a different visual reinforcer (e.g., different cartoon pictures displayed on a computer screen), the acquisition of accurate arbitrary matching usually proceeds more rapidly than when all responses produce the same reinforcer (Pilgrim, 2004). The moral is that whenever possible we should arrange different reinforcers rather than a single reinforcer for the shaping and maintenance of different response classes.

This needn't mean two reinforcer classes must differ qualitatively; it is probably sufficient for them to be discriminated. For example, if a pigeon's pecks on one key produce a green light inside its feeder and those on a second key produce a red light, with the lifting of the feeder tray requiring a peck on a left switch inside the feeder given green and on a right switch given red, the two routes of access to the food are likely to function as two different reinforcer classes (cf. Catania, 1969). It is also important to distinguish between common consequences and common contingencies. Different reinforcement schedules arranged for two different response classes will typically maintain different performances even if both produce the same reinforcer (Chapters 15 and 16 provide many examples); here it is the different contingencies arranged for the two response classes that matter most.

Reinforcer-specific effects provide another reminder that what may seem to be limits on the capacities of a species may depend instead on procedural details. For example, despite the mixed success and limited transfer from same or different stimulus sets to others in experiments on identity and oddity matching in pigeons (see Chapter 12), substantial transfer of same-different responding to novel stimulus pairs can occur after pigeons have learned to respond differentially to pairs of visual stimuli depending on whether the pair members are the same or different (Blaisdell & Cook, 2005; see also Katz & Wright, 2006).

CONDITIONED OR CONDITIONAL REINFORCERS

Once some reinforcers are available they can be used to create new ones. We saw how that could happen when we considered chaining in Chapter 10. We also saw that some reinforcers, called conditioned or conditional, become effective through their relation to other reinforcers. Maybe some of those, called unconditioned or unconditional, can be traced even back further to other sources; for example, foods may seem fundamental but many are acquired tastes. Some that have been related to many different reinforcers, of which money is a prominent example, have been called generalized reinforcers. But verbal behavior rather than reinforcement may be a more critical determinant of the functions of money; for example, the response most likely to be reinforced by a direct deposit into your bank account is the checking of your account balance. Just as variable reinforcers may maintain more behavior than fixed ones (Steinman, 1968), conditional reinforcers based on a broader range of reinforcers are likely to be more potent than those based on a narrower range. In the clicker training of pets (Pryor, 1999, 2009), it is easy to limit the relation of the clicker to edible treats, but the clicker becomes far more effective if it is broadly generalized to a variety of other reinforcers, such as playful interactions or access to new environments.

Widely used procedures for assessing the potency of conditional reinforcers are concurrent-chain schedules (Herrnstein, 1964b). We'll consider them in more detail in Chapter 16, but it is appropriate to introduce them here. These schedules include concurrent *initial links*; a consequences of initial-link responding is entry into a *terminal link*, in which another schedule operates. For example, one of two equal initial links might give access to a fixed-ratio (FR) schedule, in which the last of a fixed number of responses is reinforced, while the other might give access to a fixed-interval (FI) schedule, in which the first response after a fixed time since some event is reinforced. More responding in the first initial link than the second would tell us that the FR is preferred to the FI; the opposite outcome would tell us that the FI is preferred to the FR.

Suppose I'm standing in front of two doors behind each of which is an area where some activity is available and I must knock on one of the doors to gain access to the area behind it. Once I've finished the activity in that area I can come around and knock on the doors again. After a while I might find that the activities behind the left door are more interesting than those behind the right door, so I end up knocking more often on that door than the other. Knocking on the doors are my initial links and what I do after entering an area is my terminal link. If I knock most often on the left door you'll know that its activities are the ones I prefer.

Figure 14-1 illustrates how the procedure can be arranged on two pigeon keys. The initial-link schedules are typically equal and relatively short variable-interval schedules that operate concurrently (e.g., VI 30-s VI 30-s). Each terminal link operates separately, and the pigeon returns to the initial links after a reinforcer has been delivered. In this example, the conditional reinforcers are the green and red terminal link stimuli. In their presence responses produce food, so we can ask how potent each one is in maintaining responding during the initial links.

This procedure separates the reinforcing effectiveness of the terminal link from the contingencies that maintain responding in that link. This is a crucial feature of the procedure. Schedule A might produce a different rate of pecking than schedule B, but this rate difference couldn't tell us how hard the pigeon might work on the white initial-link keys to get to one or the other terminal-link. If slow responding in A produced as many reinforcers as rapid responding in B, maybe the opportunity for slower responding is what the pigeon would prefer.

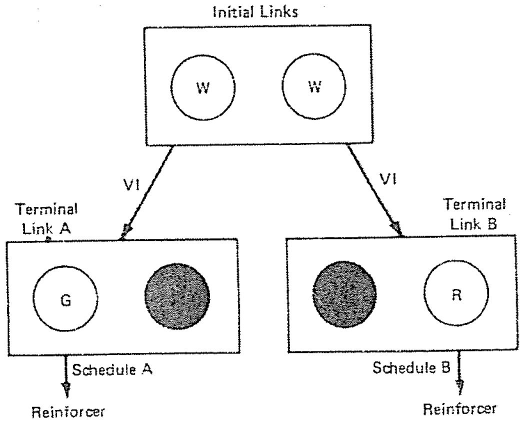

Figure 14-1 Schematic diagram of a concurrent-chain procedure as it might be arranged for a pigeon's key pecks. In initial links, equal but independent VI schedules operate on each white key (W). According to its VI schedule, left-key pecks produce terminal link A, in which the left key is green (G), the right is dark, and pecks on green produce reinforcers according to schedule A. Similarly, according to its VI schedule, right-key pecks produce terminal link B, in which the right key is red (R), the left is dark, and pecks on red produce reinforcers according to schedule B. The relative rates of pecking the two initial-link keys define preferences for the respective terminal links. For example, if a pigeon pecked the left white key more often than the right white key, the pigeon would be said to prefer schedule A to schedule B.

FREE-CHOICE PREFERENCE

Concurrent-chain schedules have demonstrated the effects of many variables on the potency of reinforcers. Let's see how they can be used to determine whether an opportunity for free choice can function as a reinforcer. Questions about freedom are questions about whether organisms prefer to have alternatives. We may believe that it is good to have options, but can we demonstrate such preferences and are they important to pigeons as well as to people? By making two keys available in one terminal link and only a single key in the other, we can ask whether pigeons prefer free choice to forced choice (Catania & Sagvolden, 1980; Sran & Borrero, 2010).

Figure 14-2 illustrates how this procedure works. Equal VI 30-s schedules operate in the initial links. Four keys are lit in both terminal links, but in the free-choice terminal link, pecks on any

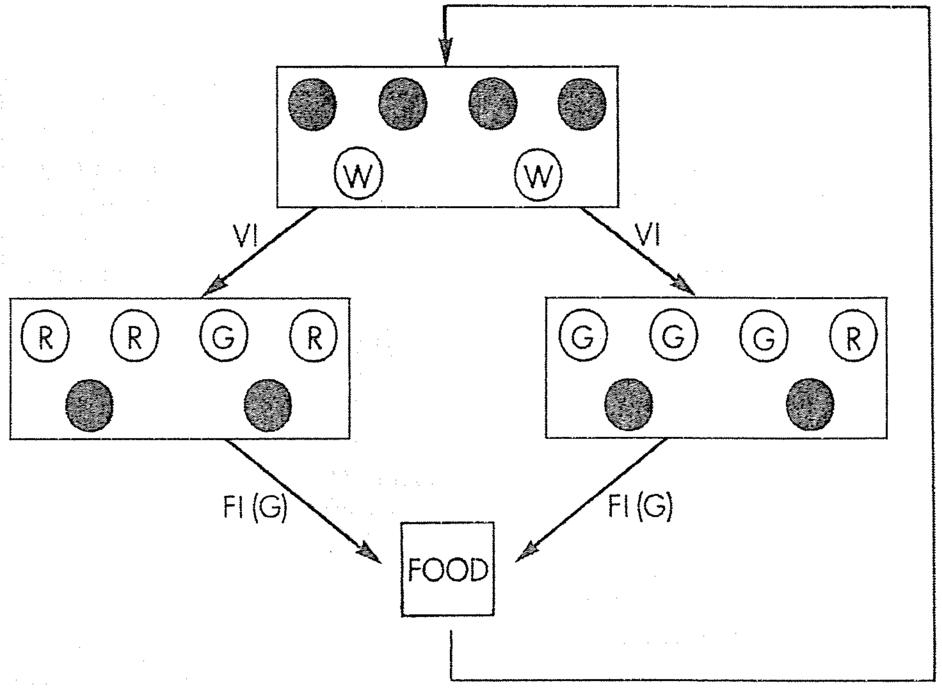

Figure 14-2 Schematic diagram of concurrent-chain schedules for studying free-choice versus forced-choice preference in the pigeon. In a six-key chamber, initial links were arranged on two bottom white keys (W). According to their VI schedules, left-key pecks produced the terminal link shown on the left and right-key pecks produced the one shown on the right. In terminal links, pecks on a green key (G) produced food according to an FI schedule and pecks on a red key (R) did nothing. The terminal link on the left provided a forced choice because only one green key was available on which a peck could produce food. That on the right provided a free choice because a peck on any of the three available green keys could produce food. Several variations on this procedure controlled for locations of the colors, green-key response rates maintained at different locations and other properties of performance. (Adapted from Catania & Sagvolden, 1980, Figure 1)

one of three green keys produce a reinforcer at the end of a 30-s fixed interval and pecks on the red key are ineffective; in the forced-choice terminal link, the same FI 30-s schedule operates for pecks on a single green key and pecks on any of the three red keys are ineffective. The pigeon will earn reinforcers after about the same amounts of time and responding whether it does so on any one of three keys in free choice or always on just one key in forced choice.

As shown in Figure 14-3, in these schedules pigeons prefer free choice to forced choice. Preferences shifted left when free-choice was arranged in the left terminal link and to the right when it was arranged in the right terminal link. Furthermore, various control procedures demonstrate that the preferences do not depend on terminal-link differences such as responses per reinforcer or the distribution of responses to different keys or particular colors and their locations (Catania, 1975a). Could it be that the pigeon has learned that different contingencies exist in free and in forced choice? If one key fails during free choice, the other is available as an alternative; if the single key fails during forced choice, no other is available to fall back on. If the free-choice preference is learned in this way, we should be able to reverse it by making more reinforcers available during forced choice than during free choice. But if we do so, the effects are only temporary; the free-choice preference returns when the reinforcers in the two terminal links become equal again. We cannot produce a durable forced-choice preference (Catania, 1980b).

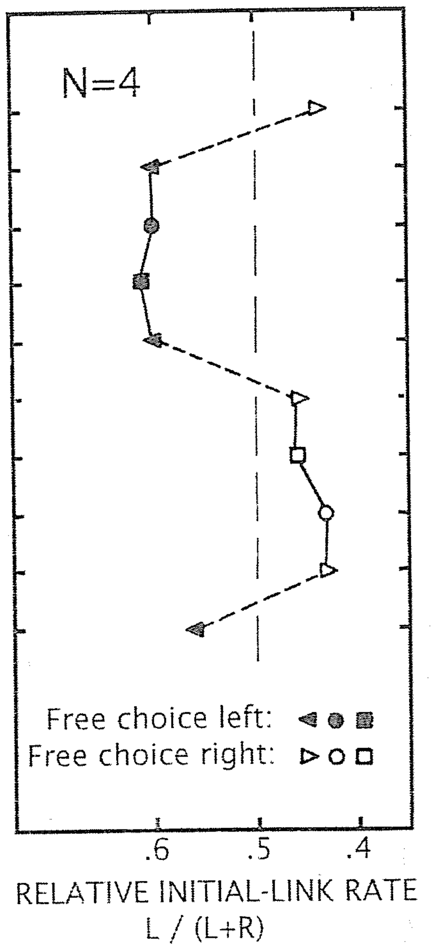

Figure 14-3 Preferences for free-choice over forced-choice terminal links averaged across four pigeons in the procedure illustrated in Figure 14-2. Preference is given by the relative initial-link rate, the proportion of left initial-link pecks, L / (L+R). The x-axis is scaled so that a data-point shift to the left corresponds to an increasing left preference. The triangles point in the direction consistent with a preference for free choice (three keys on which a peck may be reinforced) over forced choice (one key on which a peck may be reinforced). Preferences consistently shifted in the direction of the free-choice terminal link. The three symbols (triangle, circle, square) represent different arrangements of terminal-link key colors, such as fixed versus variable locations along the four keys; these arrangements had no substantial effects. (Adapted from Catania & Sagvolden, 1980, Figure 3)

Perhaps the free-choice preference has a phylogenic basis. For example, given that food supplies are sometimes lost to competitors or disappear in other ways, an organism that chooses environments with two or more food sources will probably have advantages over one that chooses environments with only a single food source. If such preferences exist even in the behavior of pigeons, they cannot be mere products of human cultures. They may occur because evolutionary contingencies have selected organisms that prefer free choice to forced choice. If so, we may be able to mask it temporarily, as by punishing responding during free choice but not forced choice, but we won't be able to eliminate it. This conclusion is based on data from pigeons, but maybe that makes free-choice preference more fundamental, because it seems to imply that at least this aspect of our human concept of freedom has biological roots. If so, no political entity can take it away.

PSEUDO-REINFORCERS

The concept of reinforcement has often been criticized. Audiences have seen talk-show hosts bring parents before them in tears because in school *My child was reinforced!* Reinforcement in the classroom has been criticized because it is said to undermine learning for its own sake. It has been charged with failures to manage the behavior of children in classrooms, where instead the rule *Spare the rod and spoil the child* is said still to hold. The building of self-esteem has been offered as a substitute for reinforcement, as if it could be a replacement for achievement rather than one of its products. Differential attention to competence is seen as unfair relative to noncontingent praise: *Good job!* Contingencies are obscured by the language of reward, which too often fails to specify what has been rewarded. Many purported problems with reinforcement come about because contingencies are mixed with verbal interventions. Reinforcement may be obscured when human verbal and nonverbal behavior interact (e.g., Shimoff & Catania, 1998). Here we explore some relevant examples.

Daniels (2009) has spelled out how consequences presumed to be reinforcing in the workplace may instead have damaging effects (c.f. Gilbert, 1978). Designating an employee of the month leaves out everyone else, and at the end of the month that lucky employee returns to the same old hunt for a parking space anyway. Meeting goals that have been set might have as one consequence a shift in the goalposts the next time around. Getting a good appraisal periodically may be fine, but what maintains the employee's behavior in the times between? Raises or bonuses that are evenly shared give as much to those who have done the least as to those who have worked the hardest. What are the contingencies and where are the reinforcers in these examples? Business provides many examples in which what superficially seem to be reinforcing practices go seriously wrong. And we haven't exhausted the examples, so behaviorally based business practices are likely to benefit not only employees and management but also the bottom line.

Bribes

Reinforcers have sometimes been equated with bribes (e.g., Kohn, 1993), but it is unlikely that the arrangements described as bribes by critics of the practice of reinforcement involve the direct effects of reinforcers. The language of bribery has a long history in law and ethics as an offer of goods or favors in exchange for favorable treatment in law enforcement or business or other endeavors. Critics of reinforcement have extended this language to the parental practice of specifying some consequence when asking a child to do something (e.g., *Clean up your room and you can play a computer game*). There are good reasons to advise parents against bribery in this sense, but the reasons are different from those offered by the critics. They correctly recognize the potentially different effects of natural and artificial consequences, but they also seriously confuse cases of verbal stimulus control with those involving other kinds of contingencies.

Parents sometimes complain that their child only cooperates with requests when there is an explicit immediate payoff. This is a problem of stimulus control. The parent may sometimes say *It's time to do your homework* and at other times *You can text your friends when you finish your homework*. Unless the child who has complied with the request gets an opportunity to text whether or not the contingency has been explicitly stated, the child will learn to comply only when the parent states it.

Here the opportunity to text is the bribe offered in exchange for completing homework. The offer may become a discriminative stimulus. Given that a bribe specifies behavior and its consequences, offers of bribes function as stimuli that set the occasion for the behavior desired by the parent. The child who is frequently bribed in this sense will learn to discriminate between conditions in which bribes are in effect and those in which they are not. The parent who often uses bribes will probably soon find that the child complies only when bribes are offered.

The child will not learn to initiate appropriate behavior if the initiation rests with the one who offers the bribe. Over the long run, therefore, compliance with bribes will probably interfere with the effects of more constructive contingencies. If reinforcement works at all in such cases, it is in strengthening compliance with bribes, which is hardly the best way to make use of reinforcers. When such unintended stimulus control develops, it is important to teach the parent to reinforce compliance without explicitly stating the contingency, or at least to reinforce compliance both when the contingency is explicitly stated and when it is not.

On the other hand, some parents have heard the language of bribes applied to the practice of reinforcement and are therefore reluctant to deliver reinforcers. This is a different problem, and these parents may resort to other contingencies, perhaps including aversive ones, if natural contingencies don't shape appropriate behavior. The appropriate course here is to educate these parents about reinforcers. In doing so it is crucial to teach them not to routinely accompany the arrangement of contingencies for the behavior of their children with statements of those contingencies. Statements of contingencies can be effective and once verbal governance is in place they can efficiently produce important behavior changes in home and school settings. But reinforcers underlie verbal governance, so they must not be neglected. This is probably good advice for teachers and clinicians as well as parents.

Intrinsic Reinforcers and the Hidden Costs of Reward

Verbal behavior is involved in the distinction between *intrinsic* and *extrinsic* reinforcers (e.g., Lepper & Greene, 1978). Some reinforcers are intrinsically effective, whereas the effectiveness of others must be established. An intrinsic reward or reinforcer is one that has a natural relation to the responses that produce it whereas an extrinsic one has an arbitrary relation to those responses. For example, music is an intrinsic consequence of playing an instrument but the music teacher's praise is an extrinsic one. Events presumed to function as reinforcers because their function has been instructed have been called extrinsic reinforcers, as when a child is told that it is important to earn good grades, but labeling them so does not guarantee their effectiveness. It has been argued that extrinsic consequences undermine the effectiveness of intrinsic ones. Despite much evidence to the contrary, this argument has persisted and continues to have impact on the use of operant contingencies in schools and other settings (Cameron, Banko, & Pierce, 2001; Cameron & Pierce, 1994; Eisenberger & Cameron, 1996).

In one experiment (Lepper, Greene, & Nisbett, 1973), one group of children received gold stars for artwork such as finger painting. After the gold stars were discontinued, children in this group engaged in less artwork than did those in a second group who never received gold stars. The gold stars, extrinsic reinforcers, were said to have undermined the intrinsic reinforcers, the natural consequences of painting.

The children were told to earn the gold stars, however, and the experiment didn't test their effectiveness as reinforcers. There were no data to show that children painted more when they got gold stars. If the stars were reinforcers at all, they were reinforcers established by instructions. To produce such effects, it may be sufficient merely to get the children talking about the relevant behavior and contingencies (cf. Wilson & Lassiter, 1982). Thus, the results had nothing to do with a difference between intrinsic and extrinsic reinforcers; instead, they were probably a demonstration of the insensitivity of instructed behavior to contingencies; Schwartz (1982b) provides another example of stereotypy determined by verbal behavior but attributed to reinforcers.

The claimed deleterious effects of so-called extrinsic rewards are only inconsistently demonstrable and are small and transient when they do occur (Cameron *et al.*, 2001; Cameron & Pierce, 1994). Problems are more likely to arise with extrinsic reward that is not contingent on performance than with contingent reward (Eisenberger & Cameron, 1996). In any case, if there is an effect its transience and small size is hardly consistent with the argument that extrinsic reinforcement may ruin the lives of children. Nonverbal effects of reinforcers must be distinguished from the social contingencies that maintain the verbal governance of behavior. When situations involve verbal behavior there is a good chance that verbal governance will override more direct effects of reinforcement.

We don't ordinarily expect behavior to occur in educational settings if it has no consequences. What would have happened to performances in my courses if I had simply guaranteed grades of A to all my students at the start of the semester and then assumed that they all would complete every assignment and study hard for every exam? The answer is obvious, and yet comparable contingencies have sometimes operated in elementary and high school education, with students promoted to higher grades without regard to the completion of earlier ones.

Grades and other institutional consequences are typically remote from what happens in the classroom, so let's consider instead the child who is learning to read. What are the natural consequences of learning to name the letters of the alphabet, or even of learning to read whole words? It is only when the child can do at least that much that the natural consequences of reading can start to take hold. Only then is the child perhaps ready to read a story, so that reading can become "its own reward." Until that happens, the teacher has no choice but to arrange artificial contingencies, using consequences such as praise or other extrinsic reinforcers.

Responsible teaching will add extrinsic reinforcers only when there aren't any intrinsic ones or when intrinsic consequences aren't working. As we have seen, even aside from this argument the evidence from the literature on the hidden costs of reward is simply unconvincing, and the consequences of its application are likely to be unfortunate. The relevant literature shows, among other things, that problems with reward are typically small and short-lived if they occur at all, and are most likely to arise not with contingent reward but rather, as we might have predicted, with rewards that are noncontingent.

ADDENDUM 14A: MOTIVATING EVENTS IN ESCAPE AND AVOIDANCE

The presentation of aversive stimuli makes their removal reinforcing. Even more so than with positive reinforcement, these establishing effects must be distinguished from discriminative, eliciting and other effects of stimuli. When we deal with positive reinforcement, we usually speak of motivating or establishing operations in terms of deprivation and related procedures. Food is less likely to reinforce the behavior of a rat that has very recently eaten than of one that has not eaten for some time. The analogous event for negative reinforcers is *presentation*; it would be called *satiation* were the stimulus food instead of shock. The motivating event in the case of negative reinforcement is the aversive stimulus itself, but why that is so isn't obvious.

The aversive stimulus is the motivating event because there is no reason to escape or avoid an aversive stimulus unless it is either actually or potentially present. It's tempting to think of the aversive stimulus as signaling a contingency, but contingencies in which responses turn off a shock cannot exist in the absence of shock. This is different from positive reinforcement; when responses produce food, the contingency can be signaled whether or not the rat has been food-deprived.

Here's an example. Shock is delivered to a rat when a light is either on or off; a lever press removes the shock for a while when the light is on but not when it is off (Bersh & Lambert, 1975). Under these circumstances the rat comes to press the lever when the light is on but not when it is off. The discriminative stimulus here is the light, because the contingency between lever presses and shock removal is signaled by whether the light is on or off. The shock does not function as a discriminative stimulus because it does not signal the operation of the contingency. It makes shock-free periods reinforcing and its presentation is therefore a motivating operation.

The contingencies operating when the light is off in this example are properly called extinction contingencies. Lever presses remove shock during light but not during dark, but shock absence would be an effective reinforcer during either one given appropriate contingencies. This wouldn't be so if shock were never present when the light is off. And even without correlated lights a rat might come to discriminate periods during which its presses turn off shock from those in which they do not, but that would depend on its sensitivity to relations between its behavior and consequences rather than to the presence of shock per se (cf. Chapter 7 on the discriminative effects of punishers).

The aversive stimulus is the motivating or establishing event because there is no reason to escape or avoid an aversive stimulus unless it is either actually or potentially present. A similar case for shock presentations as motivational can be made for avoidance. It should then be a straightforward exercise to extend the argument to cold or bee stings or loud noise or other aversive stimuli. In all of these cases **contingencies are about the consequences of responding whereas motivating events are about whether those consequences are important enough to serve as reinforcers.**

Chapter 15
Reinforcement Schedules

Schedule *is derived from the Middle English* sedule, *a slip of parchment or paper, which is in turn derived from the Latin* scheda, *papyrus leaf, and the Greek* skhizein, *to split. The Indo-European root,* skei-, *to cut or split, links* schedule *to* schizo-, *as in* schizophrenia, *and to* science *and* conscious, *from the Latin* scire, *to know, in the sense of being able to separate one thing from another.*

Not all classes of responses have consistent consequences. The reinforcement of some responses but not others, sometimes called *intermittent* or *partial* reinforcement, is a general feature of behavior. I don't always find what I'm looking for when I shop and I don't always get a satisfactory reply when I ask a question. Suppose I phone someone. Sometimes I'll get to talk to that person, but other times I'll get no answer or a voicemail greeting or someone else will answer. Games of chance by definition provide intermittent consequences. *Continuous* or *regular* reinforcement, the reinforcement of every response within an operant class, is the exception rather than the rule. For this reason, we must examine the effects of *schedules of reinforcement*, arrangements that specify which responses within an operant class will be reinforced.

Three basic types of schedules are (1) those that allow a response to be reinforced after some number of responses (ratio schedules), (2) those that allow a response to be reinforced after some time has elapsed since some event (interval schedules), and (3) those that allow a response to be reinforced depending on the rate or timing of prior responses (schedules that differentially reinforce rates or interresponse times). Such number, time and rate requirements can also be combined in diverse ways to produce more complex schedules. We've already briefly considered some schedules (e.g., interval schedules in Chapter 11 and differential-reinforcement schedules in Chapter 10). This chapter examines these and other schedules

KEY TERMS: Ratio and Interval Schedules; Fixed and Variable Schedules; Schedule Contingencies; Delay of Reinforcement; Yoked Schedules; Limited Hold.

in more detail. We'll be concerned with the effects of reinforcement schedules not only as valuable experimental tools but also as ubiquitous properties of behavior in their own right (cf. Ferster & Skinner, 1957; Schoenfeld & Cole, 1972).

Let's return to the telephone. Suppose I call a cousin whose voicemail is inactive and get no answer. My chances of getting an answer later will depend on when and not on how many times I call. My cousin will answer only if I call at a time when my cousin is available; if not it won't make any difference how many times I try calling. Some variable time will pass that depends only on how long it is until my cousin is available. I have to call at the right time to get an answer. This is an everyday approximations to *variable-interval* or *VI* reinforcement schedules. These schedules reinforce a single response that occurs after a specified time has elapsed, and the time varies from one instance to the next; earlier responses do nothing. A VI schedule is designated by the average time to the availability of a reinforcer.

Now consider a different situation. I'd like to get something from a vending machine, but the machine won't accept the change I have. I can start asking those who pass by for change and maybe I'll get what I need. In this case, whether I get change doesn't depend on when I ask. Instead, only a few of those I ask are likely to be able and willing to make change for me. I have to keep asking until I find such a person. I might succeed after asking just one or two people or I might have to ask many. In other words, my getting change depends on the number of times I ask, and this number will vary from one occasion to another. Such schedules are called *variable-ratio* or *VR* schedules. They're designated by the average number of responses required per reinforcer or, in other words, the average ratio of responses to reinforcers.

What follows will concentrate on responses such as pigeons' key pecks and reinforcers such as food deliveries. It is important to remember, however, that accurately applying the language of reinforcement schedules to settings outside the laboratory demands that we carefully specify the responses and reinforcers that enter into these contingencies. For example, consider phoning to get people to pledge to a fund drive for a charity or a political campaign. Whether any particular call is answered will depend on when you make the call, but the number of pledges you get will depend on the number of calls you make. At the level of whether your call is answered, the contingencies are those of interval schedules, but at the level of whether you get a pledge, the contingencies are those ratio schedules. It also makes a difference whether you're calling someone who keeps an irregular schedule or a business that opens promptly at a given time. **Descriptions of schedules should include the response to be reinforced, the reinforcer it produces, and the specific contingency that relates them.**

VARIABLE-RATIO AND VARIABLE-INTERVAL SCHEDULES

In a variable-ratio or VR schedule, the delivery of a reinforcer depends on a variable number of responses without regard to the passage of time. In a variable-interval or VI schedule, the delivery of a reinforcer depends on the passage of a variable time and then a single response; responses that occur earlier do nothing. Some properties of the contingencies arranged by VR and VI schedules are illustrated by hypothetical cumulative records in Figure 15-1. Three records that might be generated by a 100-response VR schedule (VR 100) are shown on the left (A, B and C); three that might be generated by a 1-min VI schedule (VI 1-min) are shown on the right (D, E and F).

A VR schedule is usually arranged by a computer that randomly selects responses for reinforcement; a VR schedule that randomly selects some fraction of responses for reinforcement is sometimes called a *random-ratio* or *RR* schedule. Before computers were available, VR schedules were arranged by a loop of tape that was driven across a switch by responses. Each response moved the tape a small constant distance. The tape had holes punched in it and whenever the switch sensed a hole the next response was reinforced. For illustrative purposes, a portion of a tape for a VR 100 schedule is shown vertically just to the right of the response scale that accompanies the VR records in Figure 15-1. In a VR 100 schedule, one response is reinforced per 100 responses on the average, but the number varies from one reinforcer to the next. The figure shows correspondences between the holes on the tape and reinforced responses, indicated by pips on the records.

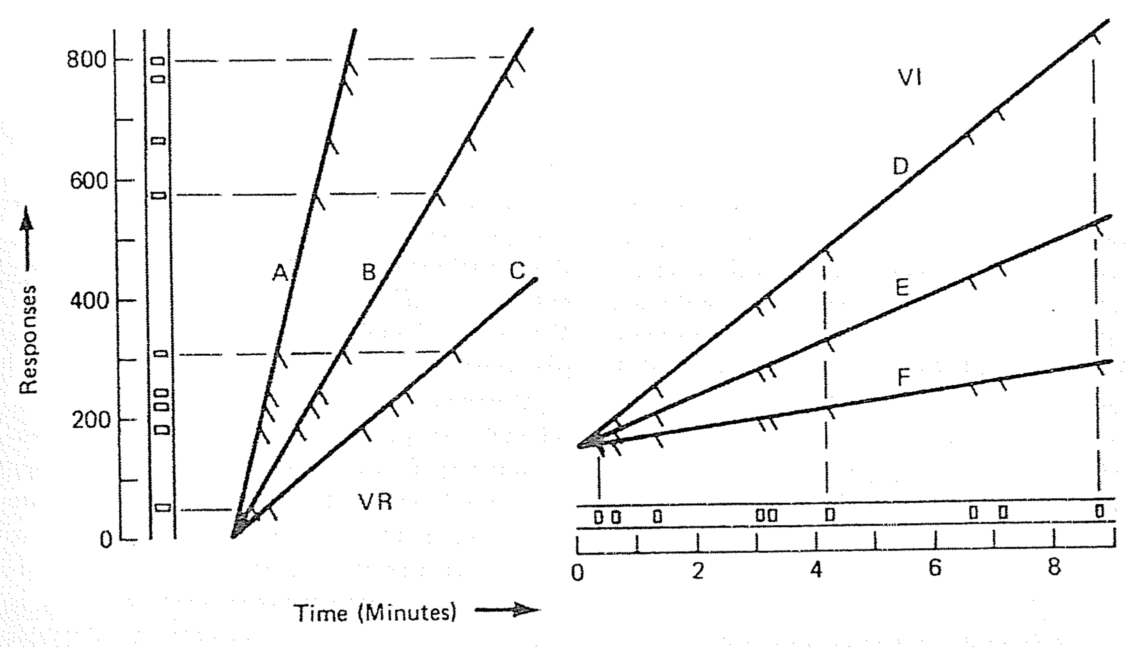

Figure 15-1 Hypothetical segments of cumulative records of responding maintained by a 100-response variable-ratio schedule (VR 100) and a 1-min variable-interval schedule (VI 1-min). The vertical strip to the left of the VR records and the horizontal strip below the VI records represent portions of punched tapes used to select responses for reinforcement (but contemporary laboratories would use computers to arrange these schedules). The three left records (ABC) show that VR response rate affects reinforcement rate but not responses per reinforcer; the three right records (DEF) show that VI response rate affects responses per reinforcer but not reinforcement rate.

In the examples in Figure 15-1, the highest response rate, A, produces reinforcers most rapidly; the lowest, C, produces them least rapidly. In other words, **in VR schedules higher response rates produce higher reinforcement rates.** With a pigeon's key pecks, an actual VR 100 schedule would probably produce a performance that looked most like record A. With moderate ratios of responses to reinforcers, VR schedules ordinarily generate high and roughly constant response rates between reinforcers. When the ratio becomes very large (e.g., hundreds or thousands of responses per reinforcer), response rate decreases; this is not so much because responding slows down overall, but rather because continued high rate responding is often interrupted by pauses.

Like VR schedules before computers were available, VI schedules also were typically arranged by a loop of tape driven across a switch. For these schedules, the tape was driven by a motor at constant speed instead of by responses. Whenever the switch sensed a hole the next response was reinforced. This arrangement allowed a response to be reinforced after some time had passed instead of after some number of responses. A portion of such a tape is shown just above the time scale on the right in Figure 15-1. Again, the figure shows correspondences between the holes on the tape and reinforced responses, indicated by pips on the records.

Once the switch operated in a VI schedule, the tape usually stopped until the reinforcer was delivered; at this point, the reinforcer was said to be *set up*, in the sense that the next response was eligible to produce a reinforcer. An alternative method now commonly used when VI schedules are arranged by computer is generating pulses at a fixed rate and randomly selecting some proportion of them to set up a reinforcer for the next response; for example, if pulses at a rate of 1/s are selected with a probability of .1, a setup would be created about once every 10 s on the average, thereby arranging a VI 10-s schedule. Schedules arranged in this way are sometimes called *random-interval* or *RI* schedules. Such schedules sometimes accumulate setups if those reinforcers have not been collected, in effect keeping the schedule timing going even if a response hasn't occurred; stopping the tape in the earlier method was an artifact of the electromechanical equipment of the times, and continuing the timing even after setups corrected for it (Catania, 2002, 2011a).

Records D through F of Figure 15-1 illustrate an important property of VI schedules: Even though response rates differ considerably across the three records, they all include the same number of reinforcers. **The VI schedule provides a relatively constant reinforcement rate over a substantial range of possible response rates.** But reinforcers aren't delivered unless responses occur, so reinforcers are produced less often than specified by the VI schedule if the time between responses becomes long relative to the time between scheduled reinforcers.

In practice, with response classes such as a pigeon's key pecks, the difference between scheduled and actual or obtained reinforcement rates is usually small (especially if the schedule timing doesn't stop with setups) . This relative independence of reinforcement rate from response rate, coupled with the relatively constant response rate that it generates, makes VI reinforcement a preferred baseline schedule, a schedule that can be used to study the effects of other variables such as drugs or chemical pollutants. With a pigeon's key pecks, an actual VI 1-min schedule would probably produce a moderate-rate performance that looked most like record D. Notice that the moderate rate of record D is not very different from the rate in VR record C. A DRL 10-s schedule, as discussed in Chapter 7, would ordinarily produce a low-rate record like VI record F. The point is that the several records in this figure show what different response rates do to reinforcement rates given VR or VI schedules; they don't represent typical performances produced by each schedule.

Response rates maintained by VR and VI schedules are presented in Figures 15-2 and 15-3. Both figures show data obtained with pigeons' key pecks reinforced with food. In Figure 15-2, rate of pecking is plotted as a function of the ratio, responses per reinforcer, arranged by a VR schedule. At VR 1 (leftmost point), every response is reinforced. Even with reinforcement duration excluded, responding only slightly exceeds 25 responses/min. Response rate is substantially higher at VR 10 (second point), and reaches its maximum, more than 200 responses/min, at VR 50. Response rate then decreases gradually with increasing VR size. But this apparently gradual decrease with larger ratios actually comes about mostly because responding begins to be interrupted by long pauses (the appearance of long pauses in ratio performance is sometimes called *ratio strain*).

In Figure 15-3, rate of pecking is plotted as a function of the reinforcement rate provided by a VI schedule (e.g., VI 1-min provides a maximum of 60 reinforcers/hr). Rate of pecking increased as a function of reinforcement rate, but the function was negatively accelerated; the change in response rate produced by a given change in reinforcement rate became smaller as reinforcement rate increased.

The two figures show that VR and VI schedules differ considerably. For example, the different y-axis scales indicate that VR response rates are higher than VI response rates over most of the range of values for each schedule; with pigeons, VR rates often exceed 200 responses/min, whereas VI rates rarely exceed 100 responses/min. Even so, the detailed effects of each schedule can vary with other conditions. For example, the shapes of the functions can be affected by whether the organism receives all of its food within experimental sessions or receives some outside the sessions. These two situations have been called *closed* and *open* economies (Hursh, 1980).

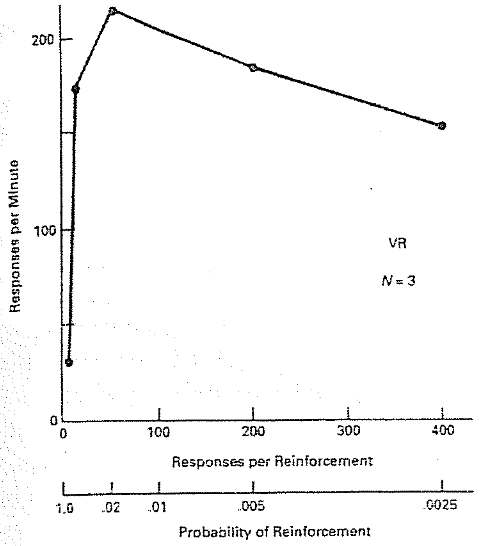

Figure 15-2 Rate of key pecking as a function of VR schedule for three pigeons. The two bottom scales show correspondences between responses per reinforcement and probability of reinforcement. (Adapted from Brandauer, 1958, Table 2)

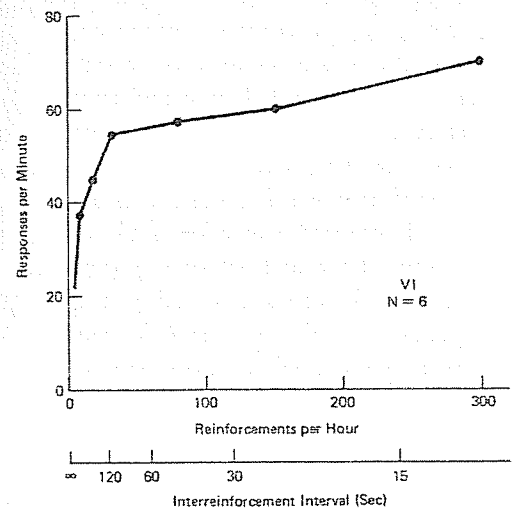

Figure 15-3 Rate of key pecking as a function of rate of VI reinforcement for six pigeons. The two bottom scales show correspondences between rate of reinforcement and average interreinforcement interval in VI schedules. (Adapted from Catania & Reynolds, 1968, Figure 1)

Ratio and interval schedules also differ in the way responding changes when reinforcement is reduced or discontinued. The idealized cumulative records of Figure 15-4 illustrate responding maintained by VR and VI reinforcement and responding during extinction after VR or after VI reinforcement. With VR reinforcement, responding decreases with larger ratios (cf. VR 50 and VR 1000). As with the decreased response rates at larger ratios, **extinction after VR reinforcement usually produces abrupt transitions from high response rates to periods of no responding** (a *break-and-run* pattern of responding). With VI schedules, on the other hand, a high rate of VI reinforcement produces more responding than a low rate of VI reinforcement (cf. VI 1-min and VI 15-min), but in both cases responding is distributed fairly uniformly in time. In addition, in contrast to the break-and-run pattern of extinction after VR, **extinction after VI usually produces gradual decreases in the rate of responding.**

Recall, though, that response rate is not necessarily related to resistance to change (Nevin *et al.*, 2001a); with regard to resistance to change, a moderate-rate VI performance might be considerably stronger than a high-rate VR performance. What makes VR and VI schedules generate such different performances? It seems reasonable that VR schedules should produce higher response rates than VI schedules. After all, when VR responding increases reinforcers are delivered more often, but this doesn't happen when VI responding increases. Yet the separation between successive reinforcers is variable in both schedules, so how do the different ratio and interval relations between responding and reinforcers act on behavior? We'll come back to that question later.

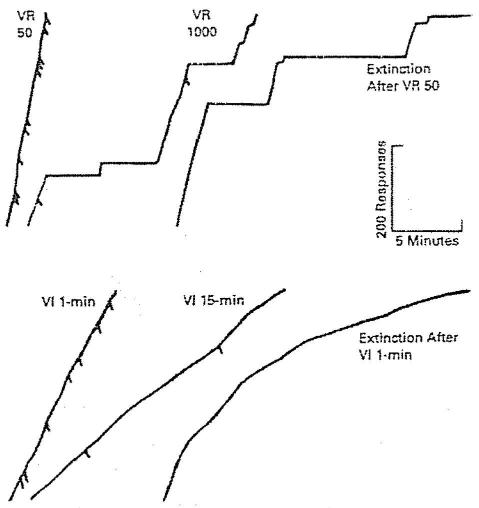

Figure 15-4 Patterns of responding during VR and VI reinforcement and during extinction after VR or VI reinforcement. These somewhat idealized cumulative records contrast effects of the two types of schedules: with VR, a higher maintained response rate, and abrupt transitions between high rates and long pauses with large response requirements or during extinction; with VI, a relatively constant response rate that decreases gradually during extinction. Both schedules generate substantial amounts of responding in extinction.

Limited Hold

One temporal contingency sometimes added to schedules is called the *limited hold* or *LH*. With a limited hold, a setup or scheduled reinforcer remains available only for a limited time; if no response occurs within that time, the reinforcer is lost. For example, suppose I have a message that I don't want to deliver by voicemail. If I'm trying to call a colleague who spends a lot of time on the phone, getting an answer may not become more likely as time passes since my last try, because my colleague may never be off the phone for very long. If my colleague makes a new call within only a few seconds after hanging up on the last one, the schedule is a VI schedule with a limited hold. I'm most likely to succeed in these circumstances by trying often, in the hope of getting through before my colleague makes the next call. In this phone example, the limited hold will vary in duration; in the laboratory, it is usually constant. A limited hold typically produces increased response rates, but a very short limited hold may allow so few responses to be reinforced that the schedule fails to maintain responding (Hearst, 1958).

Reinforcement Schedules and Causation

The effects of reinforcers depend on the responses they follow, but reinforcers can follow responses either when produced by responses or when delivered independently of responses. Does responding change in the same way when it produces a reinforcer as when it happens by accident to be followed by a reinforcer? We considered this question briefly in Chapter 5 (see also Catania, 2005c, 2006; Katz & Catania, 2005). The following suggests that causal relations between responses and reinforcers may affect behavior differently than coincidental temporal contiguities:

> Is it possible that the accidental correlations in time among responses, stimuli, and reinforcers do not exert control over behavior?... One of the characteristics of accidental correlations between behavior and environmental events is *variability*. Every aspect of behavior may vary and yet be contiguous with a reinforcer that is independent of the behavior. In contrast, behavior that is instrumental must have at least one aspect that has a more or less fixed correlation with the reinforcer. Were animals sensitive to this difference, they could detect those events over which their behavior has no real control. (Herrnstein, 1966, pp. 42-43)

In one experiment (Lattal, 1974), pigeons' key pecks were reinforced according to a VI schedule. Once VI performance was stable, the schedule was changed: At the end of some intervals, the reinforcer was delivered immediately, without regard to responding. With decreasing percentages of response-produced reinforcers, response rate decreased. When the percentage of response-produced reinforcers was zero, so that food was completely independent of behavior, response rates approached zero. This decrease in response rate has interesting implications. For example, when a third of the reinforcers were response-produced, response rates were roughly half those when all reinforcers were response-produced, but even the pigeon with the lowest response rate responded at about 30 responses/min, or a response every 2 s. A third of the reinforcers (the response-produced ones) followed responses immediately, and all the rest were likely to follow within a second or two of the last response if responses were occurring every 2 s or so. Why then didn't these accidental temporal contiguities between responses and reinforcers maintain responding at a rate close to that when all reinforcers were response-produced?

When reinforcers are delivered independently of responses, the time between the most recent response and the reinforcer is likely to vary from one reinforcer to the next. Lattal's data suggest that this variability counteracts the effects of accidental contiguities. But what about cases in which the time between the most recent response and the reinforcer varies even though the reinforcer is response-produced? This occurs, for example, when delay of reinforcement is added to a schedule (Dews, 1960).

A case is illustrated in the event records of Figure 15-5 (left). Vertical lines represent responses; arrows represent reinforcers. The top record shows a segment from a standard VI schedule: The interval ends at the dashed line and the next response, a, is followed immediately by a reinforcer. A segment from a VI schedule to which a 3-s delay of reinforcement has been added is shown in the middle record: The interval ends at the dashed line and response b produces a reinforcer 3 s later. Because other responses occur during this time, the time from the last response to the reinforcer (c) is shorter than the scheduled delay (d); it varies depending on the spacing of responses during the delay. The bottom record shows a segment from a schedule that delivers reinforcers independently of responses (technically, this is a *variable-time* or *VT* schedule): At the dashed line, the reinforcer is delivered. It is independent of responses, so the time between the last response and the reinforcer (e) varies.

The bar graph in Figure 15-5 shows data obtained with these schedules (Sizemore & Lattal,

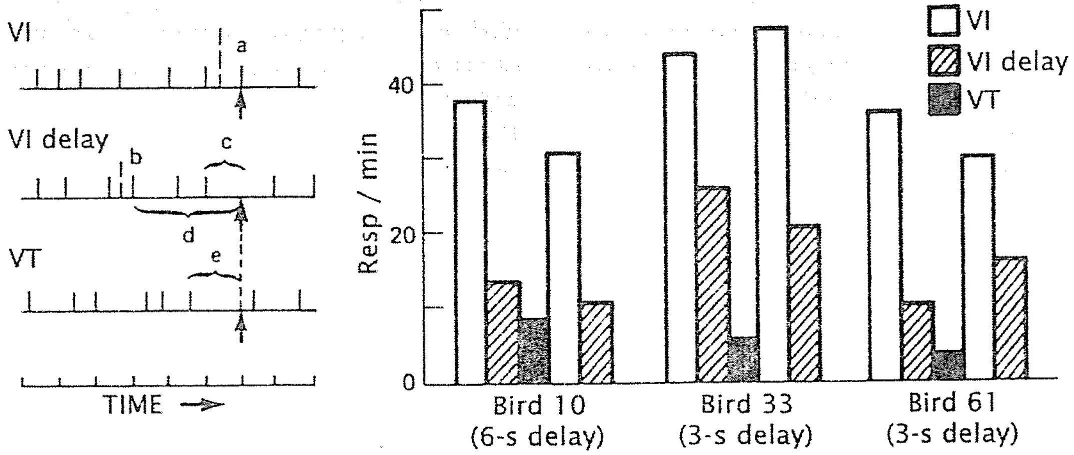

Figure 15-5 Hypothetical segments of event records from VI, VI with delay and VT schedules (left), and three pigeons' rates of key pecking maintained by these schedules (right). The highest response rates were maintained by VI reinforcement and the lowest by VT. (Adapted from Sizemore & Lattal, 1977, Table 1)

1977). The rates of pecking maintained by VI with delay were between the moderate rates maintained by the VI schedule and the lower rates maintained by the response-independent reinforcers of the VT schedule. If the difference between VI and VT depends in some way on the variable times between the last response and the reinforcer (c and e), what makes the rates with delayed VI reinforcement higher than those with VT reinforcement?

These issues remain open to experimental analysis. The different effects of response-dependent and response-independent reinforcers depend in a complex way on how correlations among events are integrated over time. The complexity is perhaps best illustrated with a human example (cf. Catania & Keller, 1981). If a lobby elevator has no floor indicator, the only consequence of my pressing the elevator button will probably be the arrival of the elevator after a variable delay. But the elevator might also arrive because someone pressed the button earlier and then took the stairs instead, or because someone already in the elevator pressed the button for the lobby floor, or because this elevator always returns to the lobby after calls to other floors. One day the elevator door might open just as I reached out to press the button; on another it might arrive and leave and arrive again, even though I didn't press the button in the meantime because I was standing there talking with someone. When the elevator arrives, I can never be sure my button press brought it to my floor. Nevertheless, I'll probably continue to press elevator buttons.

There are several elevators I use in the buildings I typically visit when I'm on campus. The programs for when the door closes and the elevator starts after someone has pressed a floor button vary from one elevator to another. In one the door closes automatically a few seconds after the press of a floor button and subsequent presses do nothing. The door in another elevator also closes automatically but after a couple of seconds a second press on any floor button will close the door. Still another imposes no delay between a press and the door closure. The variations are sufficient that I have occasionally suggested that students work out the arrangements for one or more elevators as an exercise in the analysis of contingencies. Their judgments about causation in these cases, like mine, will depend not on constant conjunctions of events but rather on a sampling of outcomes based on complex contingencies.

Research on the effects of contingencies shows that organisms are sensitive to the consequences of their own behavior. This sensitivity may depend sometimes on detailed or molecular relations between responses and reinforcers, and sometimes on overall or molar properties of response and reinforcer rates, and sometimes even on both in combination. The distinction between events caused by behavior and events accidentally correlated with behavior is central to our concepts of causation. For this reason, the study of reinforcement schedules is fundamental. Demonstrating that humans may be similarly sensitive to the consequences of their own behavior may be of practical as well as theoretical significance. For example, in vigilance tasks such as monitoring a radar screen, the detection and report of a signal can be used to reinforce the behavior of looking at the screen. If an observer has a button that lights up a radar screen and button presses produce detectable signals according to a VR schedule, a high rate of observing is maintained; this observer will more accurately detect real targets not scheduled by the experimenter than will an observer sitting in front of a continuously lit radar screen without such a schedule in operation (Holland, 1958, but see also Chapter 23).

FIXED-RATIO AND FIXED-INTERVAL SCHEDULES

If the probability that a response will be reinforced is greater at some times than at others, rate of responding is likely to be higher at those times than at others (Catania & Reynolds, 1968). For example, suppose most intervals in a VI schedule are 10 s long, and the rest are between 50 and 100 s. Response rate will probably be high at about 10 s after reinforcement, and if a response isn't reinforced at that time it might then decrease over the next 20 or 30 s before increasing again later on. We've mostly considered VR and VI schedules designed to hold the probability of reinforcement roughly constant over number of responses (in VR) or over time (in VI). But schedules can be arranged in which the number of responses per reinforcer or time to the availability of a reinforcer is constant from one reinforcer to the next; such schedules are respectively called *fixed-ratio* or *FR* schedules and *fixed-interval* or *FI* schedules (in addition, schedules in which the time between response-independent reinforcers is constant are called *fixed-time* or *FT* schedules). One important property of fixed schedules is that they introduce discriminable periods during which no reinforcers occur.

Let's turn first to fixed-ratio or FR schedules. In an FR schedule, the last of a fixed number of responses is reinforced. For example, if a vending machine takes only dollars and the item you want costs three dollars, inserting the third dollar completes a fixed ratio of exactly three responses (let's assume the vending machine is in good working order). Each fixed ratio ends with a reinforcer; the count doesn't start over if the FR responding is interrupted. The first response of the ratio is never reinforced, so FR responding typically consists of a pause followed by a high response rate. As FR size increases, the average duration of the pause increases (Felton & Lyon, 1966). The pause is called a *postreinforcement pause* or *PRP*, but it may be more appropriate to think of it as a preresponding pause; for example, pauses in FR occur as consistently when a ratio starts after a stimulus onset as when it starts after the reinforcer produced at the end of the last ratio.

Once responding begins after the pause, it typically occurs at a high rate and without interruption until the next reinforcer (FR pauses and high rates are another example of break-and-run responding; cf. VR schedules). The cohesiveness with which FR responding is maintained once it begins suggests that we should regard the FR sequence not simply as a succession of responses but rather as a behavioral unit in its own right.

The treatment of schedule performances as response units in their own right led to the development of higher-order schedules of reinforcement. For example, if reinforcement depends on completing a variable number of fixed ratios, the arrangement is a second-order schedule in which successive FRs are reinforced according to a VR schedule (e.g., Findley, 1962). An analysis of the component performances in such schedules is analogous to examining the properties that define an operant class.

Ordinarily, FR responding develops so rapidly that we can't easily see the details. Figure 15-6, however, shows FR responding developing gradually. In a two-key pigeon chamber, the pigeon's pecks on one key were reinforced according to an FR 100 schedule while a VI schedule operated concurrently on the other key. The VI schedule on the other key retarded the development of the typical FR performance and in effect made it possible to examine it in slow motion; the change in temporal patterning that usually emerges in two or three sessions took place over 150 sessions. In early sessions, responding was fairly uniformly spaced between reinforcers. With continued exposure, lower response rates and eventually pauses began to follow reinforcers. The portion of the ratio consisting of uninterrupted high-rate responding gradually became longer, so that in later sessions responding continued with few if any interruptions after each pause. With successive sessions, the FR run was built up backwards from the reinforcer.

Now let's turn to the fixed-interval or FI schedule: A response is reinforced only after some constant interval has passed since some environmental event, such as the last delivery of a reinforcer or a stimulus onset; responses before this interval ends have no effect. An example is looking at a clock as time passes during a lecture. Here we're assuming that the reinforcer is seeing that the time has come at which you can leave the classroom. Looking at the clock before then doesn't make it run any faster. Responding maintained by FI schedules usually occurs at zero or low rates early in the interval and increases as the end of the

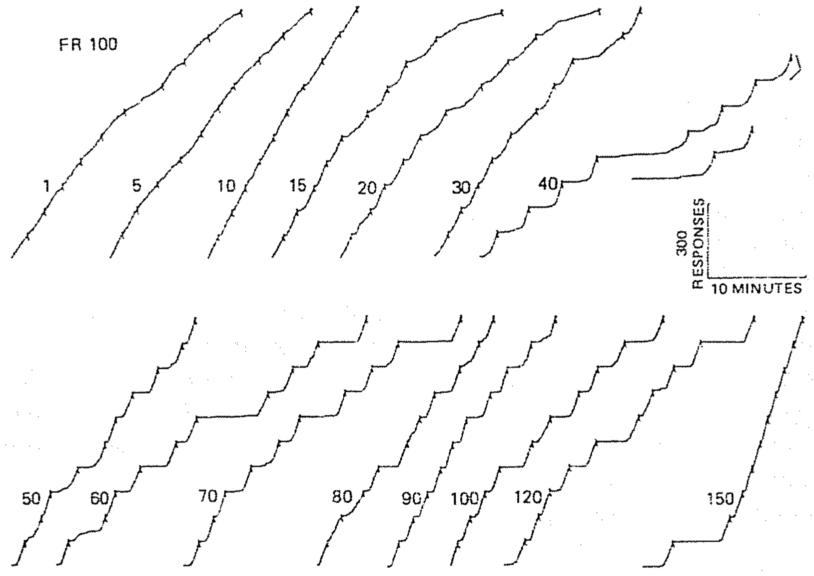

Figure 15-6 Development of a pigeon's performance under an FR 100 schedule of food reinforcement. The typically rapid development of FR responding was slowed down by the concurrent operation of a VI schedule (not shown) for pecks on a second key. The cumulative record segments are from early portions of the numbered sessions of FR reinforcement.

interval approaches, so we'd expect you to look at the clock rarely at the start of the lecture but more often as it nears its end. (We'd obviously have different things to say about the reinforcers involved if a lecture kept you so absorbed that you never glanced at the clock until it was over.)

Two sample cumulative records of FI performance in Chapter 5 (see B and C of Figure 5-6) show sequences of fixed intervals. The concave-upward pattern of such records is sometimes called FI *scalloping*. Depending on the sessions of exposure to FI schedules and other variables, FI scallops may show relatively abrupt transitions from no responding to a roughly constant rate of responding, as in B, or gradually increasing rates after responding starts, as in C. The pattern of FI responding tends to be consistent over relative rather than absolute time in the interval. For example, if responding reaches half its final or terminal rate 40 s into a 100-s fixed interval, it's likely to do so 20 s rather than 40 s into a 50-s fixed interval.

Treatments of FI performance must consider the finding that the FI scallop survives repeated interruptions. The phenomenon is illustrated in Figure 15-7 (Dews, 1962). Key pecks of four pigeons were reinforced with food according to FI 500-s schedules. The left graph shows average rates of pecking in successive 50-s portions of the interval. In another procedure, the houselight was turned on and off in successive 50-s portions of the interval; after the last 50 s of the interval, it remained on until the peck at the end of the interval was reinforced. Pecking while the houselight was off decreased (dark bars in Figure 15-7). Nevertheless, when the houselight was on, the FI responding increased throughout the interval (shaded bars) just as it had when it was uninterrupted during the interval.

Findings like these have raised questions about the relative contributions to the FI scallop of a gradient of temporal discrimination, in that the organism's different rates at different times in the interval imply discrimination of the elapsed time, and a gradient of delayed reinforcement, in that responses at various locations in the interval are consistently followed by the reinforcer at a particular delay (e.g., in FI 500-s, responding at 400 s is typically followed not less than 100 s later by a reinforced peck). The two gradients must work together, because responding at any time in an interval depends on what happened after that time in previous intervals; it can't depend on the reinforcer at the end of the current interval, because that reinforcer hasn't even been delivered yet.

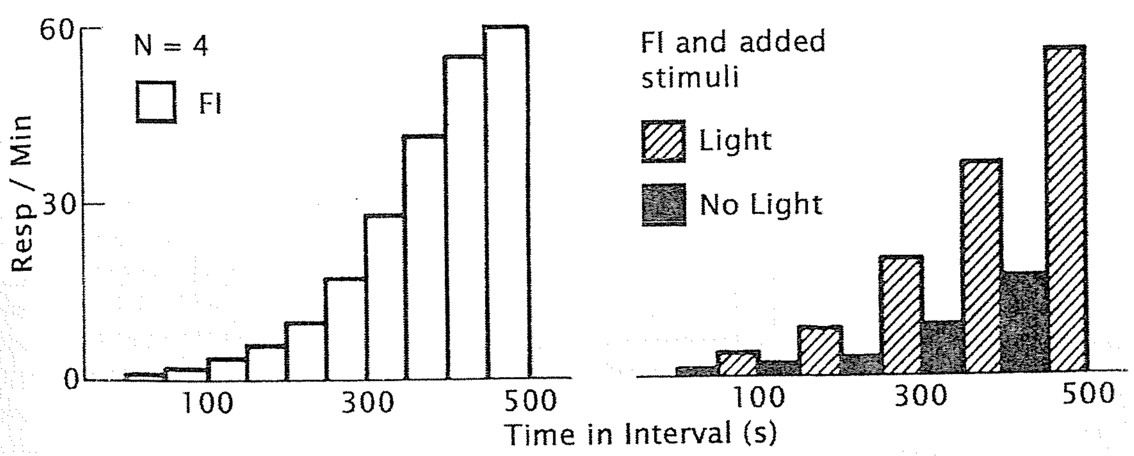

Figure 15-7 Rate of pecking in 50-s periods during a standard FI 500-s schedule of reinforcement (FI, left), and during an FI 500-s schedule in which a light and its absence alternated every 50 s (FI and added stimuli, right), for 4 pigeons. With the added stimuli, periods when the light was off produced low response rates and therefore interrupted FI responding, but rate when the light was on increased in much the same way as in the standard FI. (Adapted from Dews, 1962, Figure 2)

Delay of Reinforcement

The suggestion that the FI scallop involves some kind of gradient of delayed reinforcement carries the implication that the reinforcer produced by the last of a sequence of responses has effects that depend on its relation to all of the preceding responses and not just the one that produced it. In the early days of reinforcement schedules, schedule effects were discussed as *partial reinforcement*, and it was seen as a paradox that more behavior could be generated by reinforcing some fraction of the total responses than by reinforcing every response. But looking at schedules in terms of the delayed reinforcement of all the responses that precede the reinforced response suggests that intermittent or partial reinforcement works as it does because it allows each reinforcer to reinforce many responses instead of just one.

Figure 15-8 makes the point schematically. Both I and II show a time sample including an early response, *b*, and a later reinforced response, *a*. In I, no other responses intervene; in II, responses *c* through *h* occur between them. One way of interpreting these two cases is to say that in I the second of two responses was reinforced, so one-half (0.500) were reinforced, whereas in II only the last of eight was reinforced, so only one-eighth (0.125) were reinforced. But that interpretation assumes that the effect of the reinforcer was limited only to *a*, even though it followed the other responses too. An alternative interpretation is to assume that the effect of the reinforcer depends on its relation to all of the responses that preceded it, some after longer delays than others. From that perspective, only two responses were reinforced in I, whereas eight were reinforced in II. In this view, the partial reinforcement effect needn't be regarded as paradoxical. Instead, reinforcement schedules allow each reinforcer to act on several responses at a time instead of just the most recent one.

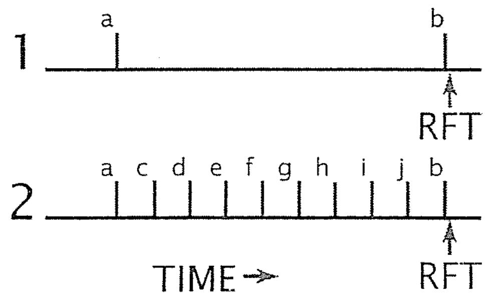

Figure 15-8 Schematic of reinforced response sequences. In 1, the last of two responses is reinforced; in 2, the last of eight is reinforced. Reinforcement probabilities are respectively 1/2 and 1/8, but only two responses are followed by the reinforcer in 1 as opposed to eight responses in 2. Even allowing for effects of delay of reinforcement, fewer responses are reinforced in 1 than in 2.

The earlier responses in a sequence that ends with a reinforcer contribute less to future responding than the later ones, because of the longer delays that separate them from the reinforcer (Dews, 1962). This means that in interpreting effects of schedules, we need to know the form of the delay gradient. The delay gradient has entered successfully into some mathematical models of operant behavior (e.g., Catania, 2005d, 2011a; e.g., Killeen, 1994), but technical problems complicate its experimental determination. For example, if we arrange delayed reinforcers only some fixed time after a response, we must either allow additional responses to occur in the interim, in which case the actual time from the last response to the reinforcer will often be shorter than the one we scheduled, or we can reset the delay with each subsequent response, in which case the resetting of the timer will differentially reinforce pauses at least as long as the delay interval, and that differentiation will be confounded with the delayed effects of the reinforcer (e.g., Catania & Keller, 1981). And we can't avoid these problems by presenting a stimulus during which the delay operates, because then we've just substituted an immediate conditional reinforcer for the delayed one.

One way to get around the technical problems is to arrange a VI schedule on two pigeon keys, as illustrated in Figure 15-9. The reinforcer is always produced by pecks on the right key, but only if preceded by a particular sequence on both keys (e.g., four left pecks followed by four right pecks, as in the sample sequence in the figure). The left pecks can be maintained only by the reinforcer produced by the last right peck, but they are always separated from that reinforced right peck by a delay determined by the required sequence. The delay can then be varied by changing the number of pecks on the right key while holding the number on the left key constant. The function relating rate of responding to the delay between the last left-key peck and the reinforcer can then be taken as a delay-of-reinforcement gradient.

Some practical implications may be more important than these technical problems. We know that the effects of delayed reinforcement on responses that precede the one that produces a reinforcer aren't restricted to responses in a single operant class. For example, a pigeon's pecks on one key may be maintained because they are followed later by reinforced pecks on another key (Catania, 1971). Now let's consider a task that involves correct responses and errors over successive trials (e.g., an experiment on matching-to-sample: cf. Chapter 12). We'll use a procedure in which we reinforce every correct response and repeat any trial with an error until the pigeon gets it right. But now we've guaranteed that any string of errors will be followed, after some delay, by a reinforced correct response. Correct responses will probably dominate eventually, because they are the ones most closely followed by the reinforcer, but errors may diminish only slowly and may even continue at a modest level even though they never actually produce the reinforcer, because they are reliably followed after a delay by a reinforced correct response too. (Errorless learning procedures may be effective because errors that never occur can never be closely followed by the later reinforcers produced by correct responses: cf. Chapter 11).

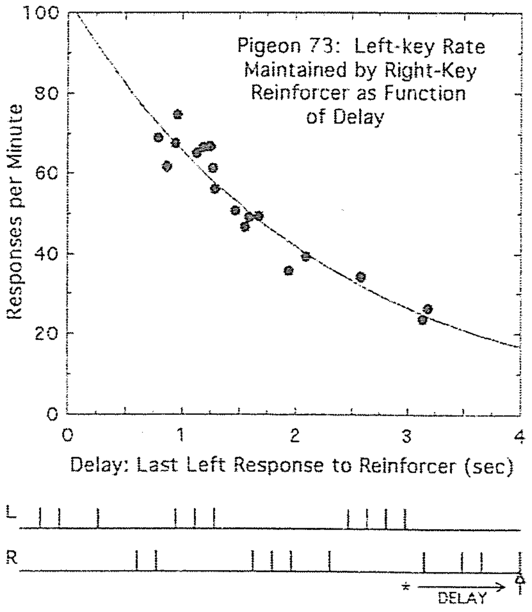

Figure 15-9 A delay-of-reinforcement gradient produced by using a varying number of responses on a right key to separate left-key responding by a delay (asterisk to vertical arrow) from the right-key reinforcer at the end of the response sequence. A sample sequence of left and right pecks is shown below the graph; the two-key sequence was maintained by a VI 60-s schedule. (adapted from Catania, 1996)

The moral of this story is that teachers must be alert for sequences in which a student's errors are followed by corrections, so that they don't strengthen incorrect responses along with the correct ones that they reinforce. A reinforcer that follows a sequence of correct responses will probably do a lot more good than a reinforcer that follows a single correct response after several errors. The best teacher will be the one who can judge whether correct responses are so infrequent that they should be reinforced even though preceded by errors or frequent enough that the reinforcer can wait until the student has made several correct responses in a row. Examples like these should remind us that shaping is often more art than science (but see Galbicka, Kautz, & Jagers, 1993).

DIFFERENTIAL REINFORCEMENT OF LOW RATE

The DRL or differential-reinforcement-of-low-rate schedule was introduced in Chapter 10. A DRL schedule arranges a reinforcer for a response that is preceded by some minimum time without responding; that time is usually constant. For example, a DRL 20-s schedule will reinforce any response preceded by at least 20 s of no responding. In DRL performance, responding is unlikely to extinguish, because decreasing low rates produce even more reinforcers. Once some responding has been reinforced, response rate typically increases, but once it has done so responses follow each other more closely so fewer meet the scheduled criterion, and response rate decreases once more, and so on. Typically the responding stabilizes at some value, oscillating between increased rates accompanied by decreased reinforcement and decreased rates accompanied by increased reinforcement. A pigeon's pecking can be maintained over long periods by such contingencies. Thus, DRL performance is highly resistant to change, though its rate remains relatively low. It therefore again illustrates the independence of response strength and response rate (Nevin & Grace, 2000).

Figure 15-10 shows pigeon data relating the DRL value (the minimum reinforced interresponse time or IRT) to the modal interresponse time each value produces (Malott & Cumming, 1964). The coordinates in the figure are logarithmic (equal distances give equal ratios, not equal intervals; note the spacings between 1, 10 and 100). The data fall roughly on a straight line in these coordinates, which means the function approximates a power function. It is well fit by the equation $T = 1.3t^{0.9}$, where T is the pigeon's modal IRT and t is the IRT required by the schedule. In this function, IRTs are longer than or overestimate the required IRT at values below 10 s or so, they are about equal to it at about 10 s, and they are shorter than it or underestimate it at values above 10 s or so.

Figure 15-11 shows comparable data from a procedure that reinforced long latencies in a trial procedure rather than in a DRL schedule (Catania, 1970). This figure plots the data in a different way to show distributions of latencies at each schedule value. A pigeon's peck on a lit key was reinforced only if some minimum time had elapsed

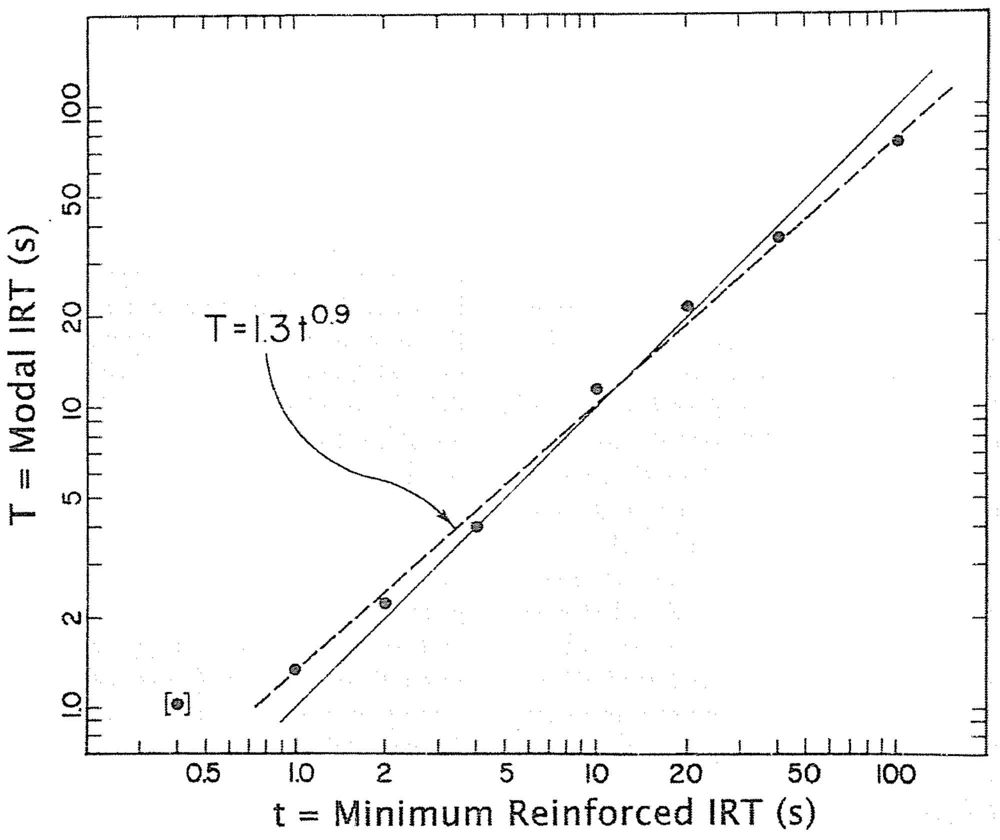

Figure 15-10 Modal pigeon interresponse times given the minimum reinforced interresponse times of various DRL schedules. The data are presented in logarithmic coordinates, in which power functions appear as straight lines. (Adapted from Malott & Cumming, 1964, Figures 11 and 12)

since a light turned on. Here again the pigeon overestimated short-duration requirements and underestimated long ones. Latencies eligible for reinforcement are shown as the unfilled areas of each distribution; those ineligible for reinforcement, i.e., all that were shorter than the criterion latency, are shown as the filled areas of each distribution. At short latency requirements (the five leftmost distributions), most responses were reinforced; as the requirements became longer a smaller proportion was reinforced.

This tendency for responding to exceed a short latency criterion and to fall below a long one is a common feature of the differential reinforcement of temporal properties of behavior in human as well as nonhuman performances. It is also relevant to the question of whether these performances should be regarded as discriminations of duration, i.e., time since the last response or times since the onset of a stimulus. Discriminations of duration by pigeons can be differentially reinforced. But these distributions are broad, so if timing of a sort is involved it is imprecise. It may be that the alternate strengthening and weakening of responding as reinforcement rates fluctuate lead directly to oscillations above and below the schedule criterion and that the oscillations are sufficient to produce data such as these. If so, invoking a theoretical construction such as timing is unnecessary.

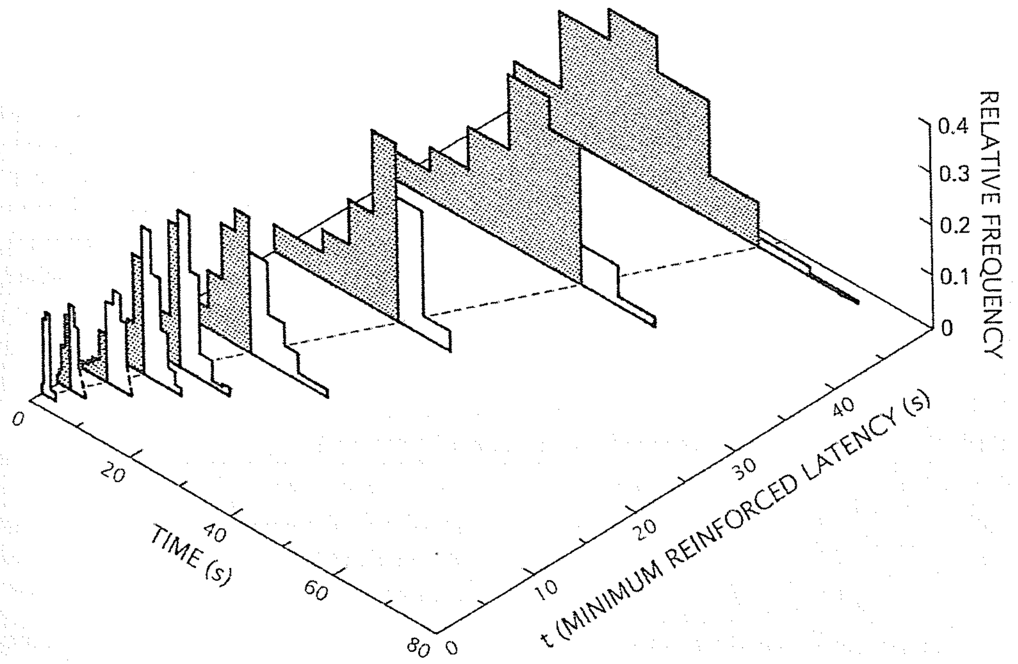

Figure 15-11 Pigeon latency distributions as a function of required latency at several criterion values in the differential reinforcement of long latencies. Shaded areas of each distribution show latencies too short to be eligible for reinforcement; unshaded areas show reinforced latencies. As with interresponse times, latencies overestimated short requirements and underestimated long ones. (Adapted from Catania, 1970, Figures 1-5 and 1-6)

REINFORCEMENT SCHEDULES: A TAXONOMY

We've considered a variety of reinforcement schedules. Some parts of the vocabulary of schedules are logical, but others are admittedly idiosyncratic. For example, the names for FI and VI versus FT and VT schedules are arbitrary (either pair could instead have been fixed-duration and variable-duration, presumably abbreviated FD and VD). Still, the different names are correlated with very different contingencies. Even though FI, DRL and FT schedules all require the passage of a constant time, what responses can do is different in each. An FI schedule imposes no conditions on responding during the interval but one response must occur at its end. Some time passes during which responses do nothing, and then the next response is reinforced. A DRL schedule, however, requires a specified time without responses, and then the next response is reinforced. Responses that occur too soon do something: They start the time over. And in an FT schedule, the reinforcer is delivered at the end of the specified time without regard to responding. In this schedule, responses never do anything. These schedule names emerged incidentally as research evolved, and are now so well established they would be hard to change.

Of several noteworthy attempts to classify reinforcement schedules more systematically (Schoenfeld, 1970; Schoenfeld & Cole, 1972; Snapper, Kadden, & Inglis, 1982), none has attained general usage. Table 15-1 summarizes some major schedules. The definitions apply whether reinforcers are arranged successively and without interruption, or occur within separate trials (e.g., an FI is usually timed from the last reinforcer, but if other events are arranged between successive intervals, timing can begin with the onset of some stimulus, such as a color presented on a pigeon key). The glossary provides additional details about schedules.

The first two columns of Table 15-1 provide schedule names and their standard abbreviations. In practice, designations of time or number usually accompany the abbreviations (e.g., VI 30-s, LH 5-s, DRL 10-s, FR 50). The third column describes schedule contingencies, the conditions under which responses are eligible to produce reinforcers (cf. FI versus FT and DRL versus DRO). The last column briefly comments on each schedule. The vocabulary of this table, presented in terms of reinforcement schedules, can also be extended to punishment schedules (e.g., Azrin, 1956). The symmetry of reinforcement and punishment, illustrated in Chapter 7 (Figure 7-1), applies also to scheduling effects. For example, superimposing an FI schedule of punishment on maintained responding produces an inverted scallop, a gradually decreasing response rate as the end of the interval approaches, instead of the increasing rate that an FI schedule of reinforcement ordinarily produces.

Table 15-1 Basic Schedules (t s = time in seconds; n = number of responses)

Name and Abbreviation		*Contingency*	*Comment*
Variable Interval *Random Interval*	VI *RI*	t s, then 1 response	t varies; with random intervals, response rate is roughly constant
Fixed Interval	FI	t s, then 1 response	t constant; generates FI scallops
Variable Ratio *Random Ratio*	VR *RR*	n responses	n varies; high constant rates, but large n may produce ratio strain
Fixed Ratio	FR	n responses	n constant; generates postreinforcement pauses and high rate runs
Variable Time	VT	t s	t varies; reinforcers are free or noncontingent
Fixed Time	FT	t s	t constant; reinforcers are free or noncontingent
Continuous Reinforcement	(FR 1)	1 response	All responses reinforced; also abbreviated CRF
Extinction	EXT	–	As procedure, often a label even if response has never been reinforced
Limited Hold	LH	Reinforcer cancelled if no reinforced response within t s	t constant if not otherwise specified; LH, added to other schedules, cannot stand alone
Differential Reinforcement of low rate (or long IRT)	DRL	t s without response, then 1 response	Maintains responding easily; decreased responding increases reinforcement and thus prevents extinction
Differential Reinforcement of high rate	DRH	1 response within t s or less of last response	Alternatively, at least n responses within t s; sometimes difficult to maintain, because decreased responding reduces reinforcement
Differential Reinforcement of paced responding	DRP	1 response within t and t' s of last response	Sets both upper and lower limits on reinforceable response rates
Differential reinforcement of other behavior	DRO	t s without response	A negative punishment or omission procedure; ordinarily decreases rate of designated response

Schedules of reinforcement are powerful determinants of behavior. It is one thing to describe them and another to watch the effects of changing the parameters of a reinforcement schedule. "In one moment, psychology became converted, for these students, from something read about passively to a practical and powerful activity that influenced large magnitudes of behavior instantly and in orderly ways.... For some, there is a magical sense of playful achievement when a procedure is converted to orderly behavioral control" (Ferster, 1978, p. 348). Reinforcement schedules have long seemed ripe for theoretical accounts, but theories have sometimes been criticized for their remoteness from the data upon which they were based. In an article called "Are Theories of Learning Necessary?" Skinner (1950) contrasted behavioral theories of learning with those that appealed to nonbehavioral dimensions or systems. His conclusion was not to reject theories in general but rather only certain types of theories:

> Perhaps to do without theories altogether is a *tour de force* that is too much to expect as a general practice. Theories are fun. But it is possible that the most rapid progress toward an understanding of learning may be made by research that is not designed to test theories.... This does not exclude the possibility of theory in another sense. Beyond the collection of uniform relationships lies the need for a formal representation of the data reduced to a minimal number of terms. A theoretical construction may yield greater generality than any assemblage of facts. But such a construction will not refer to another dimensional system and will not, therefore, fall within our present definition. It will not stand in the way of our search for functional relations because it will arise only after relevant variables have been found and studied. Though it may be difficult to understand, it will not be easily misunderstood, and it will have none of the objectionable effects of the theories here considered. (Skinner, 1950, pp. 215–216)

We may be ready for theory in this sense.

ADDENDUM 15A: YOKED SCHEDULES

The *yoked-chamber* procedure (Ferster & Skinner, 1957) lets us study some variables that operate within schedules. In yoked chambers, an organism's performance in one chamber determines the events that occur in another organism's chamber. The procedure can be used to equate reinforcement rates in VR and VI schedules: Each reinforcer produced by a VR schedule for one pigeon's key pecks schedules a VI reinforcer for the next peck of a second pigeon. In other words, the second pigeon's pecks are maintained by a VI schedule in which successive interreinforcer intervals match those produced by the first pigeon's VR performance. In these circumstances, the two schedules differ in responses per reinforcer but not in time between successive reinforcers.

Conversely, yoked schedules can equate responses per reinforcer for the two schedules by arranging things so that the responses emitted per VI reinforcer by one pigeon determine the ratios of a second pigeon's VR schedule. In this case, the second pigeon's pecks are maintained by a VR schedule in which successive ratios match those produced by the first pigeon's VI performance.

Cumulative records from both types of yoking are illustrated in Figure 15-12 (Catania *et al.*, 1977). For one group of pigeons, pairs were assigned to yoked schedules in which one pigeon's pecks per reinforcer in a VI 30-s schedule generated a second pigeon's VR schedule, as illustrated by Pigeons 402 and 410. For another group, pairs were assigned to yoked schedules in which one pigeon's interreinforcer intervals in a VR 25 schedule generated a second pigeon's VI schedule, as illustrated by Pigeons 414 and 406. Independent of whether VR is yoked to VI or VI is yoked to VR, VR schedules generate higher response rates than VI schedules. These differences emerge quickly: The records are from the last of just three sessions of 50 reinforcers each. The phenomenon is so reliable that the records in Figure 15-12 were obtained from a routine laboratory exercise in an undergraduate course in experimental psychology.

The yoking experiment shows that the rate difference between VR and VI schedules can't be

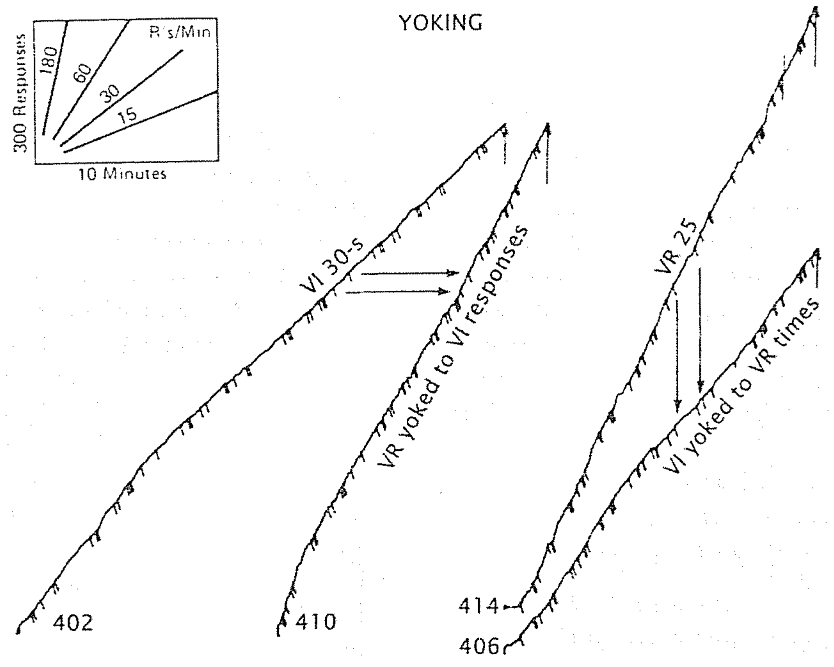

Figure 15-12 Cumulative records from the third session of yoking for two pairs of pigeons. Responses per reinforcer from Pigeon 402's VI performance generated a yoked VR schedule for Pigeon 410's pecks. Interreinforcer times from Pigeon 414's VR performance generated a yoked VI schedule for Pigeon 406's pecks. Horizontal arrows connecting the left records show correspondences of responses per reinforcer for that schedule pair; vertical arrows connecting the right records show correspondences of interreinforcer intervals. In both cases, VR response rate was higher than VI response rate. (From Catania *et al.*, 1977, Figure 1)

attributed to responses per reinforcer or time per reinforcer, because the rate difference remains even when these are the same in both schedules.

ADDENDUM 15B: INTERRESPONSE TIMES AND DELAYS

How does the difference between interval and ratio schedules generate its effects? It is correct that going faster given a ratio schedule will produce more reinforcers per unit time whereas going faster given an interval schedule will not, but procedures like the yoked scheduling of Addendum 15A rule out reinforcement rate and other input-output alternatives as candidates. The question has been around since the earliest days of schedule research (Skinner, 1938).

One suggestion was the differential reinforcement of interresponse times (Anger, 1956). Figure 15-13 illustrates how it might work. On the left are shown some possible VI sequences ending in a reinforced response after 5 s and some possible VR sequences ending with a reinforcer after the last of 5 responses. The VI sequences show that the more time passes without a response, the more likely it is that the interval has ended and the next response will be reinforced; reinforcement probability grows with IRT, as in the VI function in the graph on the right. The VR sequences show that reinforcement probability remains independent of IRT; reinforcement probability is constant, as in the VR function in that graph. Comparing the two schedules, when IRTs are short reinforcement probability is higher for VR than for VI schedules, but when they are long it is higher for VI than for VR schedules. The higher reinforcement probability for long than for short IRTs in VI should result in the differential reinforcement of long IRTs; an increase in those longer IRTs will mean a decrease in response rate. If that happens only in VI schedules, VI rates should generally be lower than VR rates.

The IRT argument seems reasonable, but the rapidity with which the VI-VR rate difference emerges works against it. The acquisition of DRL performance usually takes a while, whereas the VI-VR difference emerges quickly. Attributing an outcome that develops quickly to a process that occurs more slowly is not persuasive.

Skinner's early research had concentrated on the acquisition and extinction of responses that were either always or never reinforced. In that context he had developed an account in terms of a reserve that was augmented by the reinforcement of responses and depleted by their unreinforced emission. He started out calling his reserve a Reflex Reserve but later referred to it as an Operant Reserve. By then he had discovered reinforcement schedules and was beginning to explore the effects of intermittent reinforcement (Skinner, 1938, 1956). But when his reserve could not handle the fact that reinforcing some but not all responses produced more rather than less behavior, he abandoned it. On the output side it worked pretty well, but there seemed no way in which reduced input could produce enhanced output.

If instead all responses contribute to the reserve because each is followed after some delay by the reinforcer, however, the reserve can work. It requires that each response depletes the reserve by some amount, and that each response leading up to a reinforcer augments or replenishes it by a quantity weighted by the delay, the time that separates it from the reinforcer (Berg & McDowell, 2011; Catania, 2005d, 2011). Now consider the left side of Figure 15-13 again. With a VI schedule, if response rate increases more responses contribute to the reserve, as in the top VI sequence relative to the other two, but each of these responses also depletes the reserve a bit. With a VR schedule, if response rate increases all

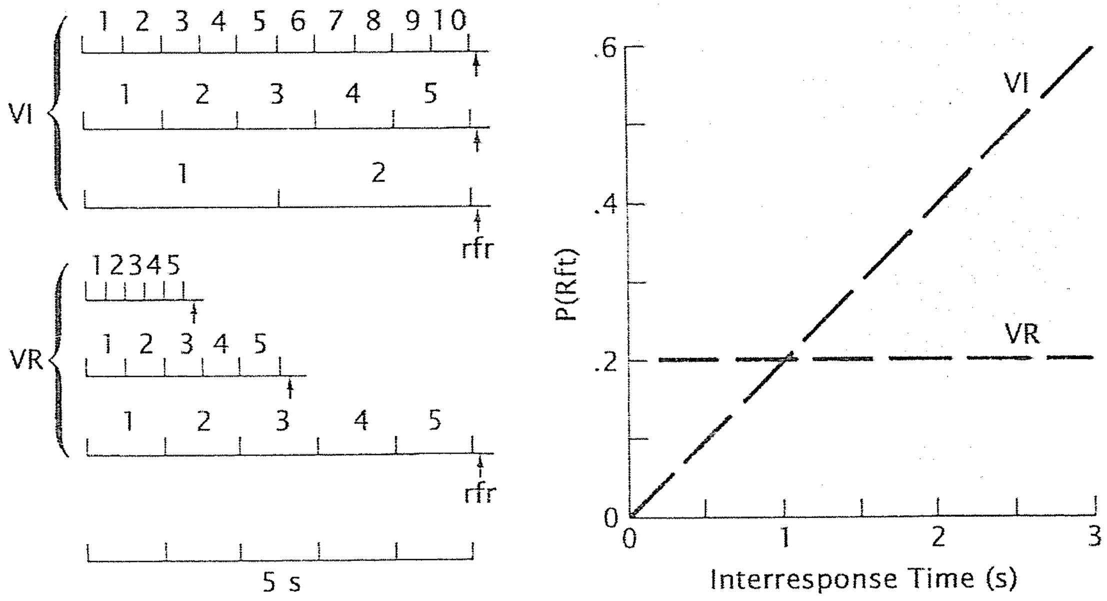

Figure 15-13 The relation between interresponse times (IRTS) and probability of reinforcement for VI and VR schedules. In the event records on the left, vertical lines represent responses, arrows represent reinforcers, and IRTs are numbered consecutively. The VI records illustrate three different response rates ending in a reinforcer after 5 s; the VR records illustrate three response rates ending in a reinforcer after 5 responses. The graph on the right convert these relations to reinforcement probability as a function of IRT. Reinforcement probability grows with increasing IRT in VI schedules but is constant in VR schedules.

of the responses get closer to the reinforcer, as in the top VR record relative to the other two; because their delays to the reinforcer are now shorter, their contributions to the reserve are larger, but their number stays the same so the amount by which they deplete the reserve is unchanged. These differences make the reserve grow larger with VR schedules than VI schedules. Note also that DRL schedules constrain the growth of the reserve, thereby keeping response rates low, because the DRL contingency guarantees that the closest response to the reinforced response is still separated from it at least by the DRL value.

This account is promising, but though it is consistent with the effects of these schedules of reinforcement it remains theoretical. It would become more persuasive if it were shown to be consistent with other behavioral facts, such as the form of the delay gradient as determined by experimental data.

Chapter 16

Schedule Combinations: Behavior Synthesis

The vocabulary of schedule combinations offers a mixed etymological bag. Multiple, *from the Indo-European* mel-, *strong or great, plus* pel , *fold, is related to* meliorate, imply *and* complicate. Mixed, *from the Indo-European* meik-, *to mix or mingle, is related to* promiscuous *and* miscellaneous. Chain *comes from a Latin word of obscure origin,* catena, *and may be related to* enchant. Tandem, *applied to bicycles after the Latin* tandem, *lengthwise, is related to* tantamount *and* though. Concurrent *and* conjoint *share the Latin prefix* con-, *with. From the Latin* currere, *to run,* concurrent *is related to* carry, intercourse *and* curriculum; *from the Latin* iungere, *to join,* conjoint *is related to* juxtapose, conjugate, yoga *and* yoke.

Schedules don't operate in isolation. They can alternate with each other, either with correlated stimuli (*multiple* schedules) or without them (*mixed* schedules). The consequence of completing one schedule can be the onset of another schedule, either with correlated stimuli (*chained* schedules) or without them (*tandem* schedules); in such contexts, one schedule can be the unit of behavior upon which another schedule operates (*higher-order* schedules). Schedules can operate at the same time, either for different responses (*concurrent* schedules) or for the same response (*conjoint* schedules), and schedules operating concurrently can each produce other schedules (*concurrent-chain* schedules). But these schedule combinations aren't of interest just for their own sakes. Instead, they bear on problems of historical significance, such as discrimination learning (multiple schedules), conditioned or conditional reinforcement (chained and second-order schedules) and choice (concurrent and concurrent-chain schedules).

This all probably seems daunting, especially if unlike me and some of my colleagues you are not

KEY TERMS: Multiple and Mixed Schedules, Schedule Interactions, Observing Responses and Attention; Chained, Tandem and Second-Order Schedules; Concurrent Schedules and Concurrent-Chain Schedules; Preference, Choice, Matching, the Matching Law.

a reinforcement-schedules junkie. But the point of what we'll be doing here is not to explore these complex schedules for their own sake. A proper exploration would take far more than one chapter in any case. Instead, this chapter will use representative complex schedules to show how they can shed light on topics of more general interest, such as choice, procrastination, and self-control. A colleague may build a pigeon procedure hoping to capture some properties of complex human behavior such as procrastination or may see something reminiscent of procrastination in a pigeon's behavior in a procedure designed for another purpose (e.g., Mazur, 1996). Either way, when my colleague shows me what the pigeon does I might be persuaded to say the pigeon was procrastinating, in which case we all might conclude that the procedure has captured some of the contingencies the operate for us when we say we are procrastinating (as I did earlier today before I finally got to work on this chapter). On the other hand, maybe the pigeon's behavior doesn't look quite right as an example of procrastination, in which case we all might review what the pigeon does in that procedure and what we do when we're procrastinating to try to figure out what contingency was left out of the pigeon procedure or, in other words, how to synthesize a procrastinating pigeon. If we come up with something and try it out in the procedure and discover that what the pigeon now does seems a pretty good analog of human procrastination, we'll have learned something about the contingencies that lead us to procrastinate.

An experimental analysis involves taking complex behavior apart to find out what it's made of. The taxonomy of behavior that this book has been developing provides the behavioral units for our analysis. Once we've finished an analysis by teasing the pieces apart, we can validate it by seeing if we can put them back together again. The reversal of an analysis is a synthesis, as when, in chemistry, a compound is synthesized after an analysis has determined its elements and its structure. We'll be using various schedule combinations to synthesize complex behavior. The success of our syntheses will tell us something about the adequacy of our analyses.

MULTIPLE AND MIXED SCHEDULES: OBSERVING RESPONSES

In multiple schedules, two schedules alternate, each correlated with a different stimulus; we speak of stimulus control when the performance appropriate to each schedule occurs during the corresponding stimulus. For example, if an FI schedule operates for a pigeon's pecks during green and a VI schedule for its pecks during red, FI scallops in green may alternate with roughly constant VI response rates in red; reinforcement can alternate with extinction in multiple schedules, but in this example a different reinforcement schedule is correlated with each stimulus.

Multiple schedules often serve as baselines in studies of variables that affect behavior. With multiple FI FR schedules, for example, drug effects on both FI and FR responding can be obtained with a single set of doses; drug effects often vary with the schedule that maintains responding. In *behavioral pharmacology* and *behavioral toxicology*, such baselines sometimes reveal large behavioral effects of substances such as pollutants at concentrations with only small physiological effects. Substances that act on behavior can be classified on the basis of such effects (e.g., Dews, 1970; Weiss & Laties, 1969).

Discriminative stimuli are effective only if the organism observes them. We can get a pigeon to observe stimuli by requiring it to peck a key to produce them. Let's start with the irregular alternation of VR reinforcement (RFT) and extinction (EXT) of a pigeon's key pecks, in a procedure illustrated in Figure 16-1. In this mixed schedule, we'll keep the same yellow keylight on during both RFT and EXT, so the pigeon usually pecks throughout both. But now we add an observing key, a second key lit white on which pecks produce stimuli correlated with the component schedules. During VR pecks on this key turn the first key from yellow to green for a while; during EXT they turn it from yellow to red. In effect, pecks on the observing key change the mixed schedule to a multiple schedule. With the multiple-schedule stimuli present, the pigeon comes to peck at high rates during green, the RFT stimulus, and at near-zero rates during red, the EXT stimulus. As in matching-to-sample, in which pecks on the sample allow the pigeon to observe the comparisons, observing responses here allow the pigeon to observe the stimuli correlated with the components of the multiple schedule.

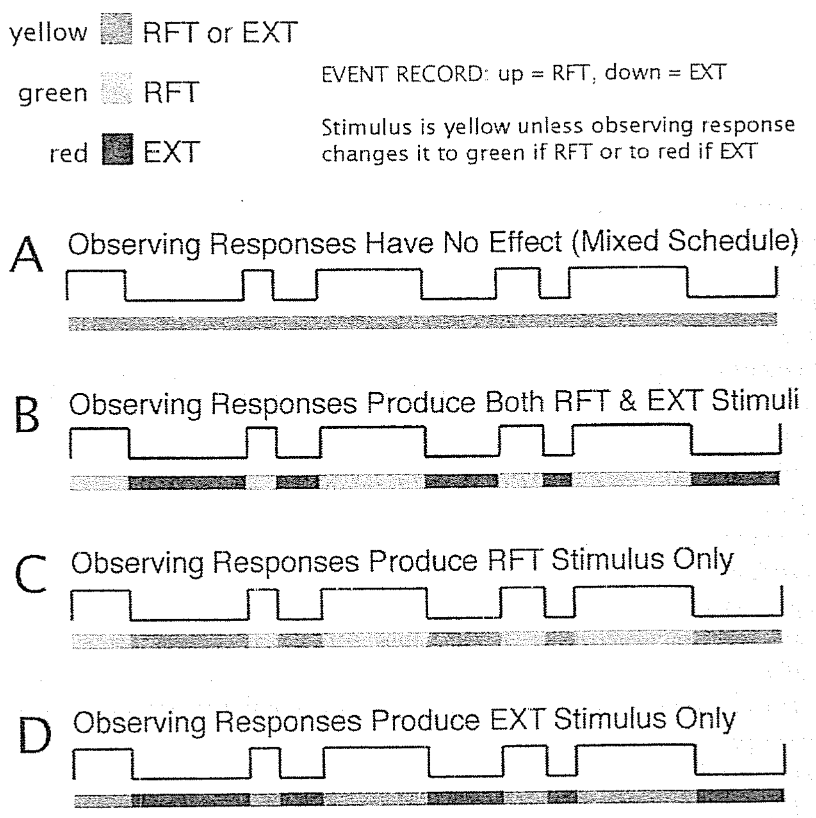

A Observing Responses Have No Effect (Mixed Schedule)

B Observing Responses Produce Both RFT & EXT Stimuli

C Observing Responses Produce RFT Stimulus Only

D Observing Responses Produce EXT Stimulus Only

Figure 16-1 Schematic of four procedures designed to explore maintenance of observing responses by stimuli correlated with reinforcement (RFT) and those correlated with extinction (EXT). The event record is up during a RFT schedule and down during EXT. Ordinarily (A), the pigeon key is lit yellow whether the current schedule is RFT or EXT (mixed schedule). In B, an observing response, a peck on a separate key, produces green during RFT or red during EXT, converting the mixed schedule to a multiple schedule. Two other procedures allow observing only during RFT, when the schedule key stays yellow if the current schedule is EXT (C), or only during EXT, when the schedule key stays yellow if the current schedule is RFT (D).

Pecks on the observing key are maintained by the observing-response procedures just described (Kelleher, Riddle, & Cook, 1962). But what keeps them going? One possibility is that the reinforcing effectiveness of the discriminative stimuli depends on their relation to the food reinforcers arranged by their schedules; another is that the pigeon observes because they allow it to behave more efficiently with respect to the component schedules. For example, the pigeon doesn't stop pecking during EXT when that key is always yellow, but it does stop when an observing response turns that key red. The question can be posed as one of whether observing is maintained because discriminative stimuli are conditional reinforcers or because they are informative.

If information is involved, then the two discriminative stimuli should be equally informative, even though one is correlated with RFT and the other with EXT. Suppose observing responses just turn on green during RFT and do nothing during EXT; the pigeon may learn that when observing responses do nothing, pecks on yellow are no longer reinforced. Similarly, suppose observing responses just turn on red during EXT and do nothing during RFT; the pigeon may learn that pecks on yellow are no longer reinforced when observing responses do nothing. In other words, if observing responses allow the pigeon to behave more efficiently with respect to the component schedules, turning on green only or turning on red only are both just as useful as turning on either red or green as appropriate.

But it turns out that observing behavior is well maintained when it produces just the RFT stimulus but not when it produces just the EXT stimulus (Dinsmoor, 1985, 1989, 1995). Similarly, stimuli correlated with differential punishment (reinforcement in one component and reinforcement plus punishment in the other) do not maintain observing responses very well; if there are any informative effects, they are overridden by the aversiveness of the stimulus correlated with reinforcement plus punishment, even though that stimulus would allow the pigeon to respond more efficiently by slowing down only when it appeared (Dinsmoor, 1983).

In other words, **the reinforcing effectiveness of a discriminative stimulus depends not on its informativeness but rather on the particular consequences with which it's correlated.** It works well if it is correlated with good news but not if it is correlated with bad news. Thus, a central problem in discrimination learning may simply be that of getting the organism to observe the relevant stimuli. We considered this issue in Chapter 11 in discussing feature-positive and feature-negative discriminations and other procedures.

As we'll see in Chapter 24, this isn't just about pigeons; we humans are not immune to such contingencies. The effectiveness of a message depends more on whether its content is reinforcing or aversive than on whether it's correct or complete. For example, it is consistent with this finding that people often hesitate to have medical symptoms diagnosed. The phenomenon has long been recognized in folklore, as in accounts of the unhappy treatment of messengers who bring bad news. Organisms do not work for information per se. That is just one more fact about behavior that has been derived from behavior analysis, and sometimes it is a fact that is resisted. When that happens, it is a fact that illustrates itself.

If our analysis so far is correct, why do we ever attend to stimuli correlated with extinction or with aversive events? Why does the mother check on the crying child? Why attend to warnings about floods or tornadoes or wildfires? It seems plausible that stimuli differentially correlated with avoidance or escape contingencies should maintain observing behavior, but relevant data in the literature on aversive control are hard to find.

An appropriate experimental preparation might include multiple schedules in which aversive stimuli were avoidable during one component and unavoidable during the other (as in Bersh & Lambert, 1975), with appropriate procedures for yoking shock densities (e.g., Hineline, 1970). By analogy with contingencies involving reinforcement and extinction, would observing responses be maintained if they produced only the avoidance stimulus but not if they produced only the stimulus during which aversive events were unavoidable? We may be more likely to observe stimuli correlated with aversive events when we can do something about them, or have had a history of doing something about them, than when they are inevitable. This is a reminder that we still have much to learn about aversive contingencies. In a world so filled with aversive events that enter into contingencies with behavior (e.g., Perone, 2003), can we safely extend our applications without expanding our experimental analyses of aversive contingencies?

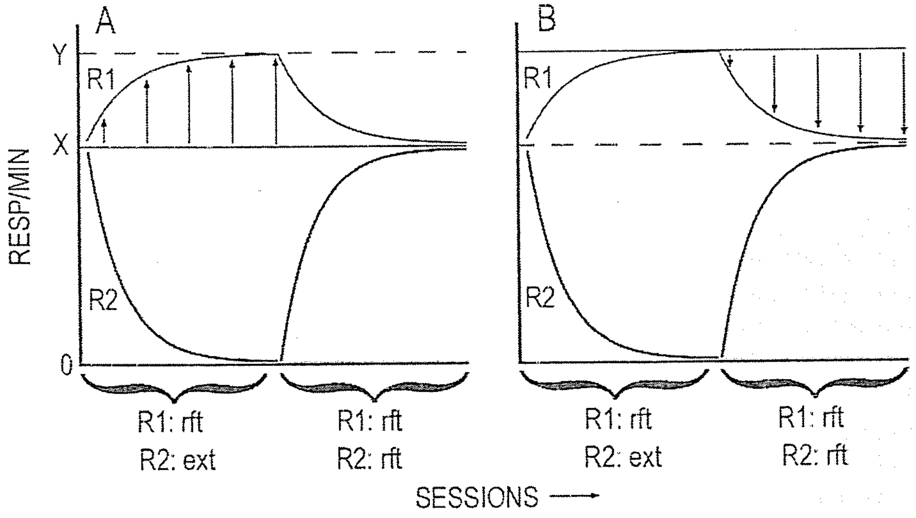

Figure 16-2 Behavioral contrast as an increase in R1 produced by extinction of R2, with VI reinforcement of both R1 and R2 taken as the baseline (A), or as a decrease in R1 produced by the introduction of VI reinforcement of R2, with R1 reinforced alone as the baseline (B). The data are the same in these two cases, but various lines of evidence argue for Case B as the appropriate interpretation. (Adapted from Catania, 1969, Figure 4)

MULTIPLE SCHEDULES: INHIBITORY INTERACTIONS

In multiple schedules, behavior in one component is often affected by what happens in the other. For example, if the schedule that maintains a pigeon's pecks during one stimulus changes from VI reinforcement to extinction while VI reinforcement continues during a second stimulus, decreased key pecking during the first stimulus is often accompanied by increased pecking during the second, even though the schedule that operates during the second is unchanged. This phenomenon has been called *behavioral contrast* (Reynolds, 1961b). Contrast effects vary with responses, reinforcers and organisms (e.g., Hemmes, 1973) and range from sustained increases to increases that last only briefly after a schedule change (e.g., Catania & Gill, 1964). They have sometimes been interpreted as the summation of two types of pecking: operant pecking maintained by food reinforcers, and respondent pecking produced by the correlation of discriminative stimuli with reinforcers, as in the delivery of reinforcers during the VI but not the EXT component of multiple VI EXT (e.g., K. J. Keller, 1974; see also autoshaping, Chapter 17).

A problem, however, is whether contrast effects should be interpreted as an effect of extinction on reinforced responding or on reinforced responding on other reinforced responding. Figure 16-2 illustrates these alternatives. Contrast effects have typically been discussed in the language of inhibition, but that language can be ambiguous in specifying what is inhibited and what does the inhibiting. Early treatments of extinction argued that responding decreased because something inhibited it (cf. Chapter 5); only later came the argument that one response can be inhibited by the reinforcement of another, as in B in the hypothetical data set of Figure 16-2.

At the start of both graphs, each of two responses (*R1* and *R2*) has been maintained by reinforcers. In both graphs, one (*R2*) is subjected first to extinction and then to the reinstatement of reinforcement. In A, the starting performance at *X* is taken as the baseline, and the decrease of *R2* in extinction is taken as an inhibitory process that is accompanied by an excitatory effect on *R1*.

When reinforcement is reinstated for *R2*, both responses return to their baseline levels at *X*. In B, the extinction of *R2* is taken as the removal of an effect of *R2* reinforcers on *R1* responding, allowing *R1* to return to its baseline level at *Y*. When reinforcement is reinstated for *R2*, it again reduces *R1* responding. If the language of inhibition is applied to A, it is not clear what is doing the inhibiting; if it is applied to B, the interaction involves specified and observable units: The reinforcement of *R2* has inhibited *R1*.

These schedule interactions occur both with concurrent schedules and with successive conditions arranged in multiple schedules. Based on various lines of evidence, they are more appropriately treated in terms of interactions among reinforced responses, as in B, than in terms of side-effects of inhibitory processes during extinction. In Chapter 5 we explored the alternative accounts of the various phenomena that had been taken as evidence that in extinction responding was somehow there all the time but inhibited. It is also useful to consider how inhibitory interactions work in other biological systems. Figure 16-3 illustrates them in the context of interactions among the rates of firing of receptors in the visual system.

The top section of each of the graphs in Figure 16-3 schematically depicts an array of ten photoreceptors showing for each a light level ranging from zero to nine. In 1a, for example, the third and seventh cells are lit at a level of 8 and the rest at a level of 0. The graph below each cell array displays the activity of each cell as a rate of firing on an arbitrary scale. In 1a, the stimulated cells show substantial activity and the other cells remain inactive. In 1b through 1d, the light level is held constant at 8 for the third cell but the light is moved closer, to the sixth (1b), fifth (1c) and finally to the adjacent fourth cell (1d). The closer together the lit cells, the lower their rates of firing; the firing of each cell reduces the rate of firing of the other. This mutual reduction is an example of reciprocal inhibition (Hartline & Ratliff, 1957; Hartline, Wagner, & Ratliff, 1958).

Graphs 2a through 2c show effects of lighting groups of two, four or six adjacent cells. With two cells, each is equally active. With four or six, however, the end cells are more active than the ones in the middle. This is because the activity of the cells on the end is inhibited only by neighbors on one side, whereas the activity of those in the middle is inhibited by neighbors on both sides (Ratliff & Hartline, 1959). The enhanced activity difference between unlit and lit cells at the boundary is one kind of contrast effect. Note, however, that all of the differences are created by inhibition, which reduces the firing of a cell relative to what it would have done if lit all by itself.

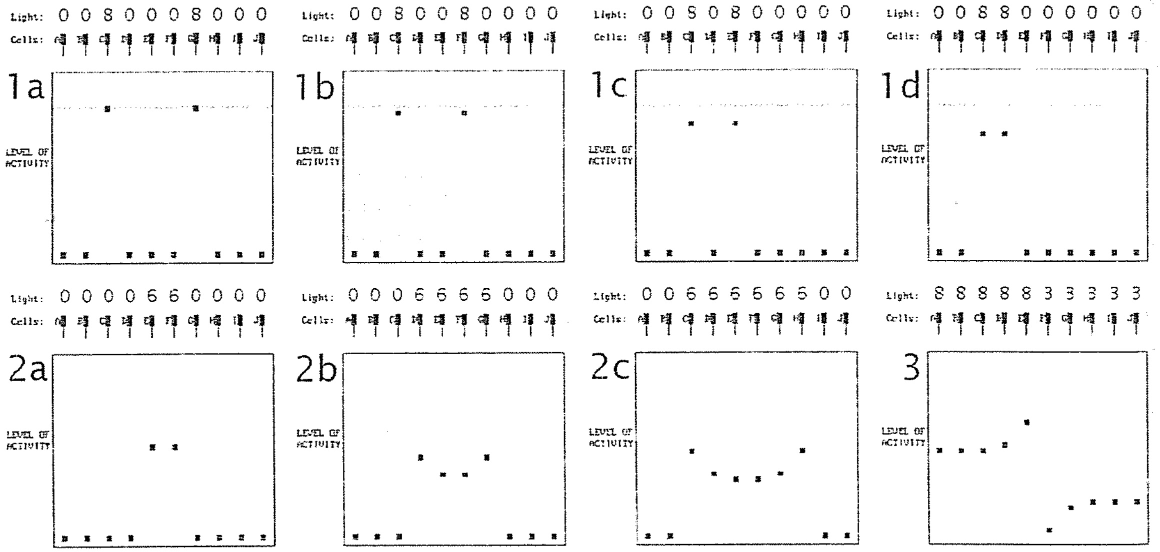

Figure 16-3 Schematic display of rates of firing of an array of ten photoreceptors given various levels of light stimulation for each; light levels are given by the numbers above each array. The activity of each cell reduces the activity of neighboring cells; this reciprocal inhibition falls off with distance, as illustrated in 1a through 1d. When a group of cells is lit, those at the ends are inhibited only by cells to one side whereas those in the middle are inhibited by neighbors on both sides, so the endmost cells are more active than the other identically lit cells. These spatial effects are illustrated in 2a through 2c, in which groups of two, four or six cells are lit. When a discrete boundary occurs between cells lit at one level and cells lit at another, as in 3, the cells at the boundary show an amplified difference relative to the cells further from the boundary. These contrast effects are all reductions relative to a cell's activity when it is lit all by itself.

Another effect, not shown here, is that the initial firing produced by lighting a cell is typically substantially higher than its later firing under continued light (Hartline, Wagner, & Ratliff, 1956). The cell may not be immune to inhibitory effects from its own activity. More important, this means that these interactions enhance differential responding to temporal boundaries as well as spatial ones. Effects like these of are great generality across a range of species, from the visual systems of horseshoe crabs, flies, and fish to those of people, and from the visual system to hearing and touch and other senses (von Békésy, 1967).

Behavior, too, is a biological system, so we should expect to find behavioral units interacting in corresponding ways. Behavioral contrast is an example limited just to two response classes, analogous to graphs 1a through 1d of Figure 16-3. Figure 16-4 shows effects analogous to graph 3 when a discontinuity between reinforcement and extinction was arranged along a spatial array, and effects analogous to the greater activity after stimulus onset than during continuous stimulation when reinforcement and extinction were arranged in a repeating temporal sequence (Catania & Gill, 1964). In each case, a pigeon's pecks were reinforced according to an FI schedule in discrete trials. In the spatial sequence (Figure 16-4, left), the

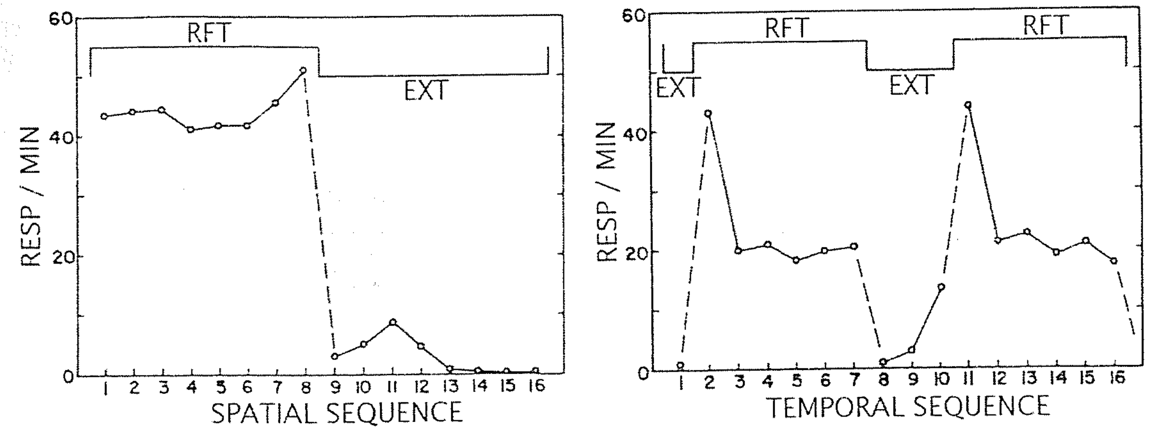

Figure 16-4 Spatial and temporal effects on reinforced responding analogous to those seen in sensory systems. In a spatial sequence consisting of FI (RFT) and extinction (EXT) trials (left), the top eight of an array of sixteen lamps were correlated with RFT and the bottom eight were correlated with EXT; lamps were lit one at a time in irregular order. In the temporal sequence (right) RFT trials given one top lamp and EXT trials given one bottom lamp were repeatedly presented in a consistent order. The spatial sequence showed a contrast effect at the boundary between RFT and EXT; the temporal sequence showed transient contrast, in that RFT rates were highest immediately after an EXT trial. (Adapted from Catania & Gill, 1964, Figures 2 and 3)

lamps in a vertical array of sixteen lamps were lit in irregular order on successive trials. The top lamps, 1 to 8 were correlated with the FI schedule (RFT); the bottom lamps, 9 through 16, were correlated with extinction (EXT). As for inhibitory interactions in sensory systems, rate differences were amplified at the boundary between RFT and EXT. In the temporal arrangement (Figure 16-4, right), one of the EXT lamps and one of the RFT lamps were arranged in a repeating sequence of sixteen trials. As in sensory systems, response rates were consistently higher when the FI stimulus (RFT) followed immediately after an extinction stimulus than when it followed another FI stimulus. **Response classes are strengthened by the reinforcers they produce, but those reinforcers also inhibit other response classes.**

CHAINED, TANDEM, AND SECOND-ORDER SCHEDULES

Chained schedules have been used extensively to study *conditioned or conditional reinforcers*, reinforcers that acquire their capacity to reinforce through their relation to other stimuli already effective as reinforcers (Kelleher & Gollub, 1962). For example, the feeder light becomes a reinforcer only through its relation to food in the feeder, and the clicker becomes a reinforcer only through its relation to the various reinforcers arranged by the pet owner (cf. Chapter 14). Because the opportunity to engage in highly probable responses reinforces less probable responses only if the opportunity is signaled, the conditional reinforcing functions of stimuli have something in common with their discriminative functions.

Extended Chains

Assume that we've gradually increased the size of an FR schedule arranged for a pigeon's pecks on a white key to a value of FR 240. In this situation, with only the food produced by pecking, the pigeon easily maintains an adequate daily ration. In fact, since a pigeon can eat its daily diet in as little as 4 or 5 min and each food delivery lasts only 4 or 5 s, including the time to get from key to food hopper, the pigeon may be done for the day after just 60 reinforcers. At FR 240, the pigeon must peck 14,400 times/day to earn this much food. We can expect typical FR performance: postreinforcement pauses followed by rapid and uninterrupted pecking. Nevertheless, our pigeon will earn enough food daily to keep itself alive and healthy for an indefinite time. Its life expectancy may be as high as 20 years, and with this kind of diet a veterinarian might even judge the bird to be slightly overweight (unless they live at sites where they are handsomely fed by tourists, pigeons in the wild tend to maintain themselves at 70% or less of free-feeding weights, so maintenance at 80% of free-feeding weight is hardly a starvation diet).

But pigeons can't count very well, and we might wonder whether the pigeon is at a disadvantage each time it goes through its 240 pecks. Suppose we try to help it estimate how many pecks are left in the ratio by changing the color on the key after every 40 pecks. Each ratio starts with the key blue for the first 40 pecks; it turns green for the next 40, yellow for the next 40, orange for the next 40, red for the next 40, and purple for the final 40 pecks, the last of which produces the reinforcer. The pecks per reinforcer remain the same; only the key colors have changed.

With no distinctive stimuli, the pigeon worked well enough, earning an adequate daily ration. Have we helped by providing the colors? The stimuli have the surprising effect of slowing the pigeon down. The pause at the start of the ratio, when the key is blue, lengthens. When the pigeon finally responds, the pecks that once came in rapid succession now occur sporadically. When the key turns green after 40 pecks, the pigeon may pause again before starting the next 40 pecks. When the key was always white, the pigeon didn't pause at this point. In green, another 40 pecks and the key turns yellow, and then another 40 and it turns orange. This time the pigeon is less likely to pause: 40 more pecks turn the key red and then 40 more purple, and the pigeon now quickly completes the last 40 pecks and food is delivered. But then the

key is blue again and another long pause begins. The stimulus changes in the chained schedules had some reinforcing effects, but they were mostly restricted to the late components, close to the reinforcer.

The added colors didn't help. This chain will never become a chunk (cf. Chapter 10). The pigeon takes much longer to earn each food delivery than it had when the key remained white. In fact, even though this pigeon maintained an adequate diet at 240 pecks per food delivery when the key was always white, we'd now be wise to watch carefully to be sure its daily food intake doesn't decrease drastically. And we produced this change just by adding some stimuli (cf. Catania, Yohalem, & Silverman, 1980).

When we added the colors, we broke the 240 pecks down into four distinct units of 40 pecks each. We call these *chained* FR schedules: the separate stimuli correspond to the links of the chain. But the chain breaks the 240-peck sequence up into six FR 40 components instead of holding it together more cohesively. When the key was always white, pecking at the start of the sequence wasn't so very different from pecking at the end, just before the reinforcer. Once the stimuli were added, however, pecking during the early stimuli became less like pecking later. In blue, for example, pecks never produced food; at best they turned the key green, but pecks never produced food during green either. With pecking reduced early in the chain, the time to complete each 240 pecks increased. In chained schedules, a stimulus supports less responding the further it is from the end of the sequence. Even severe food deprivation may not counteract this effect.

These differences occur with various schedules. So far our example has involved tandem and chained FR schedules, but Figure 16-5 illustrates similar effects with cumulative records from schedules with FI components. Comparable effects of chained schedules occur with other organisms besides the pigeon, with different kinds of responses and different schedules, and with different kinds and orderings of stimuli (Kelleher & Gollub, 1962). Because these effects have such

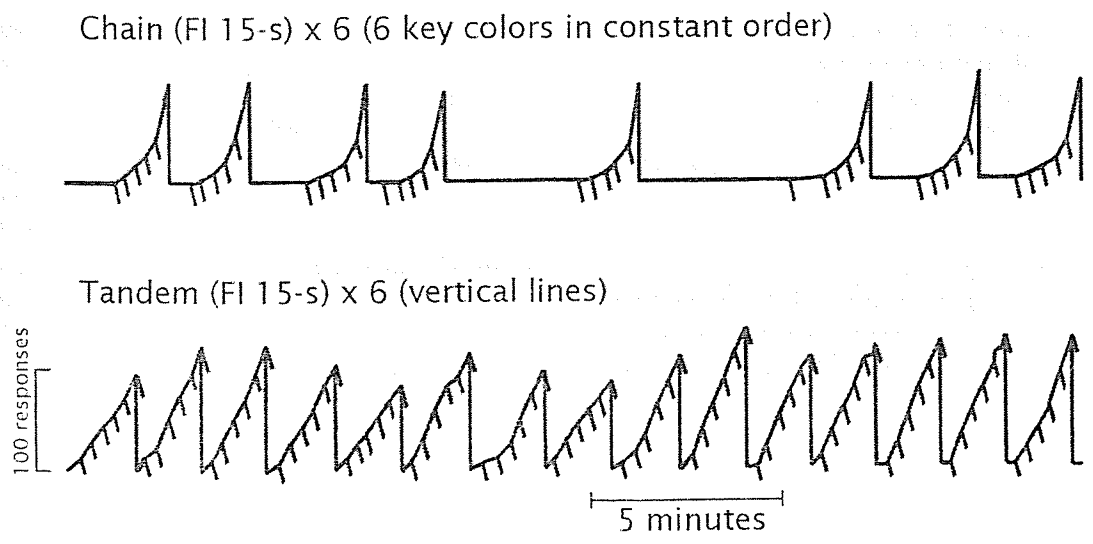

Chain (FI 15-s) x 6 (6 key colors in constant order)

Tandem (FI 15-s) x 6 (vertical lines)

Figure 16-5 Sample pigeon cumulative records from a 6-component chained FI 15-s schedule in which the sequence of key colors remained constant over successive reinforcers and from a comparable 6-component tandem schedule with vertical lines as the stimulus. The records reset to baseline at each reinforcer and successive components are shown by pips on the records. Note the long pauses in the early chain components and the high rates in the last component, compared to the relatively continuous responding during the tandem schedule.

generality, we may wonder whether they are relevant to human behavior. The things that people do, of course, depend on consequences more complex than food deliveries. Nevertheless, so much of what we do involves sequential behavior that we might wonder whether we sometimes operate under the strain of too many links in our chains. Did you ever have trouble writing that first page of a paper that was due? That first page is certainly the beginning of a chain. The ability to formulate and achieve long-range objectives is supposed to be a unique characteristic of the human species. But if adding even a single link to a chain can so devastate the pigeon's behavior, perhaps we should be alert for similar effects in our own behavior.

These effects depend on a constant ordering of the chained stimuli. The long pauses decrease markedly if the stimulus order changes from one reinforcer to the next (Gollub, 1977). Analogous but inverse effects also hold for schedules of punishment; relative to tandem schedules, chained schedules of punishment reduce responding mostly in the later components of the chain (Silverman, 1971). An implication is that punishment after a deed is done probably has its greatest effect on the behavior that precedes getting caught and only minimal effects on the much earlier behavior that led up to the misdeed.

Brief Stimuli

The stimuli in chained schedules can become conditional reinforcers, but they combine with discriminative effects in such a way that responding is reduced. Yet this outcome seems inconsistent with the effects of some stimuli that acquire their reinforcing properties. In human behavior, for example, money presumably becomes a reinforcer by virtue of the various commodities for which it can be exchanged (it's sometimes called a *generalized reinforcer*, because it doesn't depend on a specific primary reinforcer; (cf. Ayllon & Azrin, 1968, on token economies).

Early experiments on the effects of conditional reinforcers were conducted during extinction, after a history of consistent pairings with a primary reinforcer (e.g., making tone a conditional reinforcer by following it with food and testing later with tone alone). Although this procedure countered the objection that responding might be maintained directly by the later primary reinforcer rather than by the conditional reinforcer itself, it was also one in which the effectiveness of conditional reinforcers diminished rapidly once the primary reinforcer was removed. Convincing demonstrations of conditional reinforcers came only when reinforcement schedules were applied to their analysis: Schedules were arranged not only for the production of conditional reinforcers by responses but also for the contingent relation between conditional and primary reinforcers (Zimmerman, Hanford, & Brown, 1967). For example, a tone may function as a conditional reinforcer even if it's followed by food only 1 time in 20.

In second-order schedules, the completion of one schedule is a behavioral unit that is reinforced according to another schedule, as when the second-order schedule FR 10 (DRL 5-sec) arranges a reinforcer for every tenth interresponse time that is longer than 5 s. This arrangement, with a brief stimulus at the completion of each first-order schedule, is one of several types of second-order schedules. Consider a schedule arranged for a pigeon's pecks in which the peck that completes each 60-s fixed interval produces a brief green light on the key and in which every tenth such interval is followed also by food; the notation for such schedules may include the brief stimulus: FR 10 (FI 60-sec: green). Such schedules typically maintain FI scalloping within intervals even though most intervals end without food.

In contrast to chained schedules, second-order schedules with brief stimuli can greatly amplify reinforced responding. For example, when a chimpanzee's pushbutton presses were reinforced with food according to an FR 4000 schedule, postreinforcement pauses ranged from many minutes to hours. But when the light accompanying food delivery came on briefly after every 400 responses, responding increased and typical postreinforcement pauses decreased to 5 min or less. The light converted the simple FR 4000 schedule to a second-order schedule, FR 10 (FR 400: light), that amplified the amount of behavior maintained by the food reinforcers (Findley & Brady, 1965). The cumulative records in Figure 16-6 illustrate these schedule effects. The three FR 4000 sequences (unshaded) are each preceded by pauses of half an hour or more; the three FR 400 x 10 sequences with brief stimuli (shaded) include some early pauses, but once they are under way they are completed in less than half an hour.

Variables such as the relation between the brief stimuli and primary reinforcers determine the effectiveness of second-order schedules (Gollub, 1977). Both chained schedules and second-order schedules with brief stimuli involve conditional reinforcers, but **where chained schedules may attenuate responding, second-order schedules may amplify it.** Their opposite effects illustrate how critically schedule effects depend on the detailed relations among stimuli, responses and consequences (Malone, 1990, pp. 294-296; Morse & Kelleher, 1977).

Second-order scheduling can also include other kinds of operants, as when correct responses in matching-to-sample (cf. Chapter 9) are reinforced according to various schedules. In such contexts, lower levels of accuracy may be correlated with those times at which responding is least likely to be reinforced. For example, errors are more likely early in FR runs or in the early portions of FI scallops than in responding that occurs later and therefore closer to reinforcer deliveries (e.g., Boren & Gollub, 1972; Thomas, 1979). Classes defined sequentially may also imply higher-order classes. For example, individual pecks are functional units, but within FR performance the entire ratio may function as a unit. The FR performance has a property that is consistent with higher-order classes: As long as the higher-order class is reinforced, the subclasses within it may also be maintained even though they are no longer reinforced. In other words, the first peck of the fixed ratio does not extinguish, even though by itself it never produces the reinforcer.

CONCURRENT SCHEDULES: MATCHING AND MAXIMIZING

Any reinforced response is likely to occur in a context of other behavior maintained by other consequences. We must therefore examine the effect on

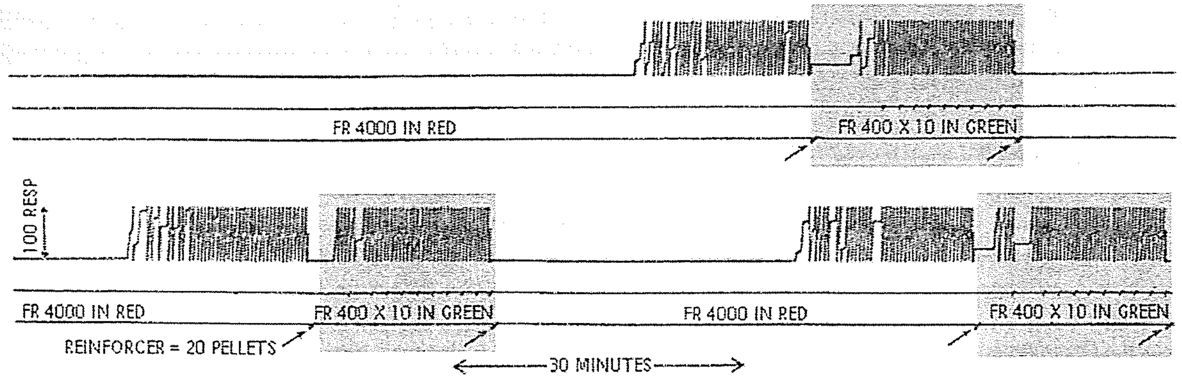

Figure 16-6 Sample chimpanzee cumulative records from an FR 4000 schedule during red and a second-order FR 400 x 10 schedule during green. During red the FR 4000 performance was preceded by very long pauses; responding on the second-order schedule began relatively promptly during green. (Adapted from Findley & Brady, 1965).

one response of reinforcement schedules operating for other responses. Concurrent schedules are schedules arranged simultaneously for two or more responses. Consider an FR 25 schedule of food reinforcement for a pigeon's pecks on one key and an FR 50 schedule operating concurrently for pecks on a second key. Alone, either schedule maintains responding, but when they operate concurrently, responding is likely to be maintained exclusively on the key with the FR 25 schedule. The outcome isn't surprising. A reinforcer requires only 25 pecks on the first key but 50 on the second.

Now consider concurrent interval schedules, such as VI 30-s reinforcement of pecks on one key and VI 60-s reinforcement of pecks on the other. In this case, the pigeon produces 120 reinforcers/hr by pecking only the first key or 60/hr by pecking only the second. By pecking both, however, it produces the reinforcers of both schedules, or 180/hr. In this case, responding is likely to be maintained on both keys. Variables with small effects in single-response schedules often have large effects in concurrent schedules, which are therefore useful for studying effects of reinforcement variables (e.g., reinforcer duration: Catania, 1963a; response force: Chung, 1965). As with multiple schedules, increases in the reinforcement of one response reduces the response rate of others (but cf. Killeen, 1972). This interaction depends on the reinforcers and not on competition between responses for available time (Catania, 1963b; but cf. Henton & Iversen, 1978). Perhaps the most general feature of concurrent performances is that increases in the reinforcement of one response reduce the rate of other responses (Catania, Sagvolden, & Keller, 1988; Rachlin & Baum, 1972). If the response rate generated by a given rate of VI reinforcement is independent of how these reinforcers are distributed to the two keys, it follows that increasing the reinforcement of one response will reduce the rate of the other.

These interactions are illustrated in the sample cumulative records of Figure 16-7. On the left are shown segments from a VI 180-s schedule for a pigeon's pecks on a green key (G) that operated concurrently with the multiple-schedule alternation on another key of VI 180-s during red (R) and extinction during yellow (Y). Despite the different response rates on the multiple-schedule key, rates on the green VI key were roughly constant. In a different experiment, pecks on one key were maintained by a VI 180-s schedule operating concurrently either with extinction (EXT) on a red key or with a signalled VI 180-s schedule on the other key that was mostly dark and therefore on which the pigeon rarely responded. In this signalled VI schedule, the key was mostly dark, but a discriminative stimulus appeared whenever a reinforcer was set up according to its schedule (the S^D was a lit key); typically the pigeon pecked there and then quickly returned to the yellow key after the reinforcer. Few responses occurred on the other key to compete with yellow-key pecking in either condition, but during the signalled VI those rates decreased relative to the rates with EXT on the other key. In these examples, the response rates varied independently of each other and the rate interactions depended on reinforcers (Catania, 1966).

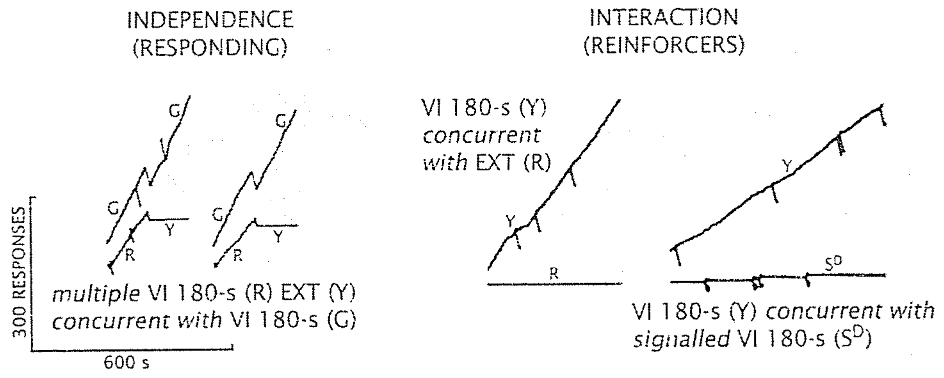

Figure 16-7 Sample cumulative records of pecking maintained by concurrent schedules. On the left, VI 180-s on one key (G = green) operated concurrently with multiple VI 180-s (R = red) EXT (Y = yellow) on a second key. On the right, VI 180-s on one key (Y) operated concurrently with either EXT on a second key (R = red) or signalled VI 180-s on that key (S^D = signalled); the key was ordinarily dark but lit whenever a reinforcer had been set up (pips on the record). The VI (Y) rate was lower when reinforcers were delivered on the other key, even though little responding occurred there. The left records illustrate the independence of response rates from different rates on another key; the right records show that one response rate is reduced by the reinforcement of another response. (Adapted from Catania, 1966, Figures 11 and 12)

Although pigeons distribute their pecks to both keys with concurrent VI VI schedules, there is a complication. If pecks on one key are immediately followed by a reinforced peck on the other, the reinforcer may act on both responses, so that the pecks on one key are maintained partly by reinforcers scheduled for the other (cf. Chapter 10 on delayed reinforcers). For this reason, concurrent VI procedures have often incorporated a *changeover delay*, which prevents either response from being reinforced immediately after a changeover from the other. With a changeover delay, the pigeon distributes its responses to concurrent VI VI schedules roughly in proportion to the distribution of reinforcers they arrange (Herrnstein, 1961); in the example, the pigeon pecks a VI 30-s key about twice as often as a VI 60-s key.

This phenomenon is sufficiently general that it's been proposed as a general law of behavior, called the *matching law* (Davison & McCarthy, 1988; Herrnstein, 1970). It states that relative responding matches the relative reinforcement produced by that responding. The law even holds for concurrent ratio schedules, because exclusive responding on one schedule means that all the reinforcers will be delivered according to that schedule. Herrnstein's account has also been applied to the responding maintained by a single reinforcement schedule, on the assumption that other events besides the reinforcers arranged by the experimenter may have reinforcing effects even though we can't identify them.

The matching law summarizes performances in a variety of schedules, but its status as a convenient description or as a fundamental property of behavior rests on whether it can be derived from simpler processes (Catania, 1981; Hineline, 2001; Rachlin, 1971). For example, consider how concurrent VI VI schedules operate when arranged for a pigeon's pecks on two keys. As the pigeon pecks one key, time passes during which the VI schedule for the other key may set up a reinforcer. A time will come when the reinforcement probability for changing over to the other key exceeds that for continuing to peck the same key. If the pigeon emits the response with the higher current reinforcement probability and this shifts from one key to the other as time passes, the pigeon will distribute its responses to both keys in concurrent VI VI schedules (Hinson & Staddon, 1981; Shimp, 1966). This has been called *maximizing*; with several responses available, maximizing means emitting the response with the maximum reinforcement probability. With unequal concurrent ratio schedules, this is always the one with the smaller ratio, but with concurrent VI VI schedules the response with the maximum reinforcement probability changes from moment to moment; concurrent VI VI performance has therefore been called *momentary maximizing*. Thus, momentary maximizing at the molecular level may lead to matching at the molar level. Matching and maximizing may seem contradictory alternatives, but they are measured in different ways (Hineline, 2001).

We can't speak about matching without some sample of responses and reinforcers from which to estimate relative frequencies, but we can speak about maximizing with a single response, just by noting whether it was the one with the maximum reinforcement probability. To some extent, the issues involve the level of detail at which performances are analyzed. Furthermore, matching and maximizing don't exhaust the possibilities. For example, other analyses have examined whether concurrent performances can be described as *opti-*

mization (the organism produces the highest possible overall reinforcement rate), *satisficing* (it meets some minimal requirement, such as a given food intake) or *melioration* (it balances performance so as to produce equal reinforcement rates under different conditions); this is not an exhaustive list, and in their quantitative details these and other treatments are beyond the scope of the present account (e.g., Killeen, 1994; Mazur, 1991; Nevin & Grace, 2000).

The matching law converts separate response rates to ratios. Sometimes the ratios must be adjusted because the two responses differ in some way that biases the responding organism to one or the other. Such considerations have led to the generalized matching law (Baum, 1974), which takes such bias and other factors into account and which transforms the data to logarithmic rather than linear coordinates. Figure 16-8 illustrates some features of the matching law using data from two experiments on concurrent schedules. In Catania (1963b), key-B pecks were maintained by a VI schedule that provided 20 food reinforcers per hour while the rate of key-A food reinforcement varied. In Hollard and Davison (1971), key-B pecks were maintained by a schedule that provided 60 reinforcers per hour of intracranial stimulation while the rate of key-A food reinforcement varied; the reinforcer difference created a bias that could be corrected by weighting one food reinforcer to be about equal to six intracranial stimulations.

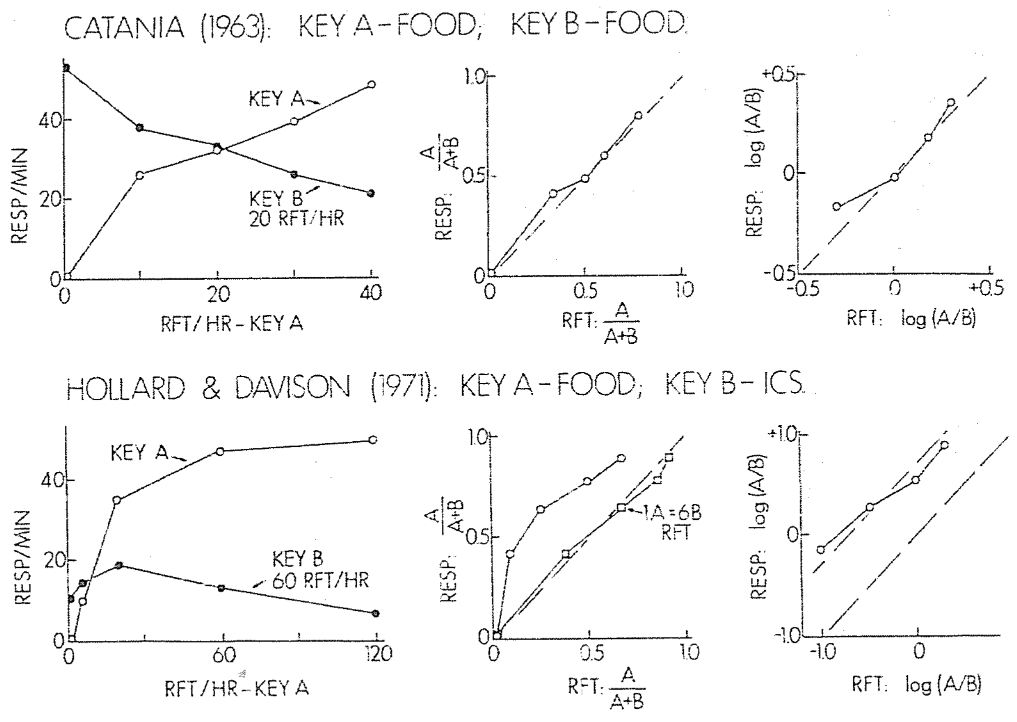

Figure 16-8 Data from two concurrent-schedule experiments converted from response rates (left) to relative rates (middle) and then to the logarithmic coordinates of the generalized matching law (right). When converted to relative rates, both provided roughly linear functions, though the latter data had to be converted for a bias. When converted to generalized matching ratios, the data are roughly linear in these logarithmic coordinates but some features of the response-rate data are lost. (Adapted from Catania, 1981, Figure 2)

The figure shows that converting from response rates to generalized derivatives reveals some similarities but also loses some features of the data. In this case, 10 data points were reduced to 4 (a value of zero cannot be plotted in the logarithmic coordinates on the right). More important, the characteristic decrease of one response with increases in the reinforcement of another shown in the top data set does not occur with the second data set, where both responses first increase together before their rates begin to change in opposite directions, but these differences cannot even be displayed in the data transformations called for by the matching law.

Whether generalized matching is a fundamental process that in some way dictates the details of schedule performances or is a derivative of the moment-to-moment responding generated by reinforcement schedules has long been a source of controversy. Is behavior built from the top down or the bottom up? Those who argue for the former call theirs a molar approach (Baum, 2002), contrasting it with molecular analyses. Biology has gone through similar controversies in its history, but it now appears that the molar biologists were wrong to give up on the microbiologists. "In evolutionary parlance, Evo Devo reveals that macroevolution is the product of microevolution writ large" (Carroll, 2005, p. 291). Why should we expect behavior to be any different?

CONCURRENT-CHAIN SCHEDULES: PREFERENCE

When two or more responses are available at the same time they are sometimes called choices and engaging in one or the other is sometimes called choosing. But is choice something I do just before responding or is it my responding itself? If it is something I do just before responding, then it too is behavior and I should be able to analyze it. For example, I may look over a menu before placing an order; I may read reviews before going to a movie; I may list alternatives and their pros and cons before making an important decision. But none of these is a choice; they are all preliminaries to a choice, which is simply what I end up doing (cf. Skinner, 1950). In any case, a dominance of one alternate over others in a sequence of choices is typically called a preference, and concurrent-chain procedures are particularly well-suited for the analysis of preferences.

Concurrent-chain schedules (Herrnstein, 1964b) were introduced in Chapter 14 as a procedure for studying preference for free choice versus forced choice. Two equal schedules operate in *initial links*; the consequence arranged for each is another schedule, a *terminal link*. Suppose I'm standing in front of two doors behind each of which is a room where some activity is available and I must knock on a door to gain access to the room behind it. Then knocking on the door is my initial link and what I do after entering the room is my terminal link.

This procedure separates the reinforcing effectiveness of the terminal link from the contingencies that maintain responding in that link. For example, the response rates maintained by concurrent VR DRL schedules wouldn't tell us whether a pigeon prefers VR to DRL schedules; we should not take the high VR and low DRL rates maintained by these contingencies as implying a VR preference. Given a choice between slower DRL responding and faster VR responding, the pigeon might prefer the DRL. Concurrent chains allow this kind of choice when they arrange the VR and DRL schedules as terminal links (e.g., schedules A and B in Figure 14-1), because the initial-link pecks that produce the VR or DRL schedules don't share in either the VR or the DRL contingencies. **We judge preferences among situations not by how much behavior they produce but by the relative likelihoods with which an organism enters them.** We must distinguish response rates and resistance to change when we study schedule performances; we must similarly distinguish between response rates and choices when we study preferences (Nevin *et al.*, 2001a).

Concurrent chains have shown that reinforcement rate is a more important determinant of preference than the number of responses per reinforcer (e.g., Neuringer & Schneider, 1968) and that variable schedules are preferred to fixed sched-

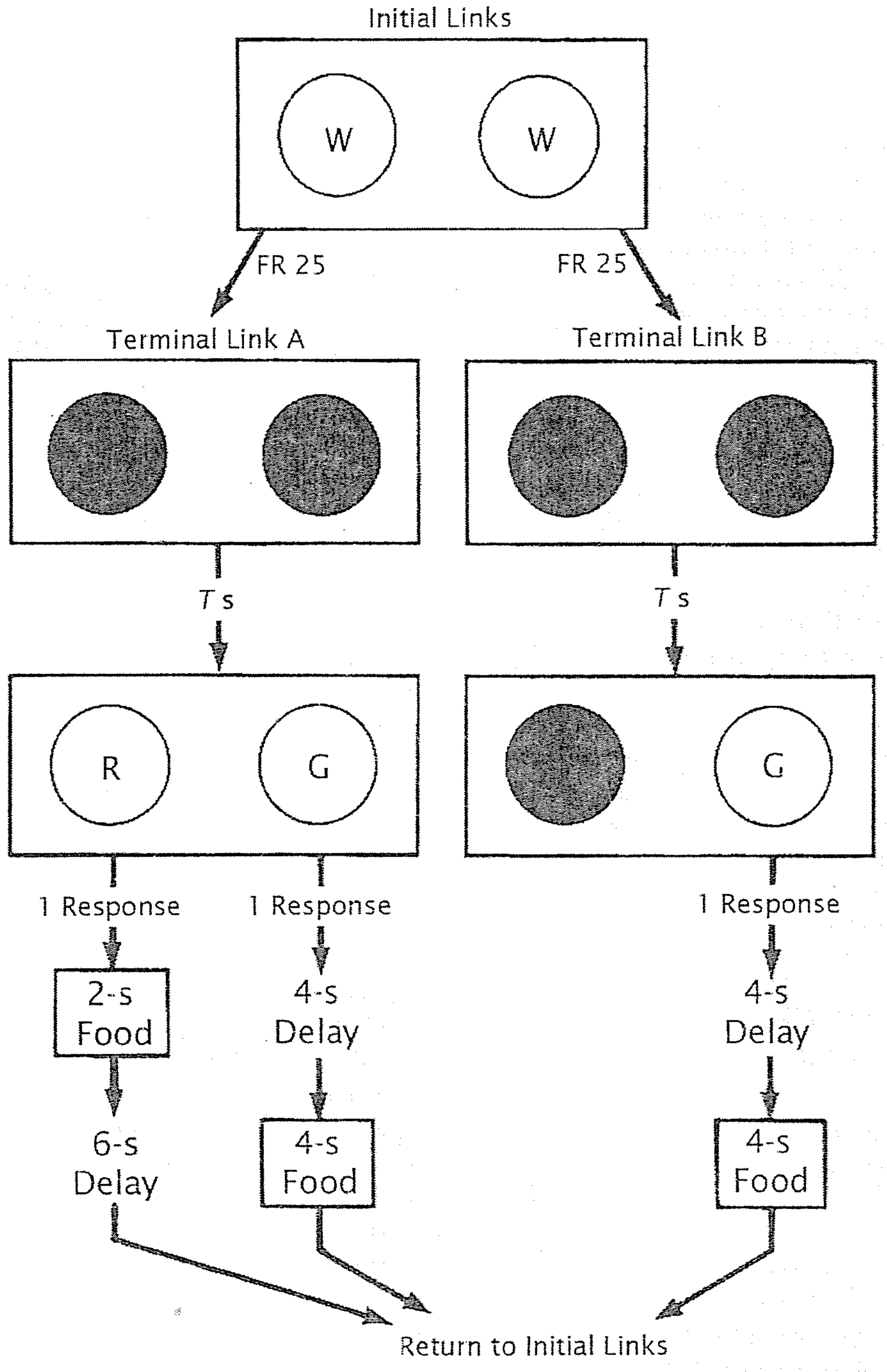

Figure 16-9 A concurrent-chain procedure that synthesizes some properties of impulsiveness, commitment and self-control. According to FR 25 schedules, pecks on white (W) initial-link keys are followed after *T* s by terminal links. In terminal-link A, red (R) and green (G) keys respectively make an immediate small reinforcer or a delayed large one available. In terminal-link B, the green key alone makes only the delayed large reinforcer available. (Adapted from Rachlin & Green, 1972)

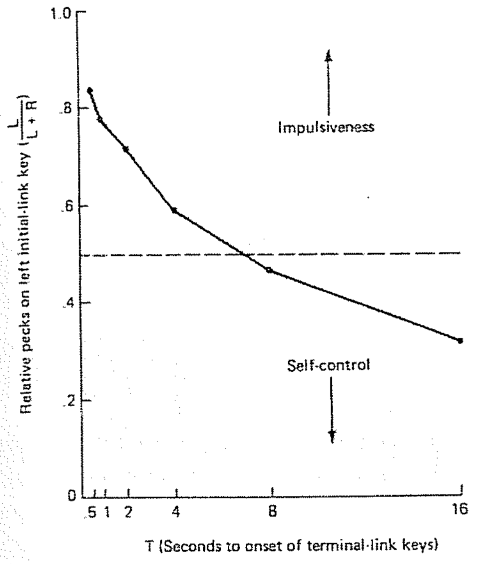

Figure 16-10 Relative initial-link pecks producing terminal-link A (left initial-link pecks divided by total initial-link pecks) as a function of *T*, the time to the onset of the terminal-link keys. (Adapted from Rachlin & Green, 1972, Table 1)

ules (e.g., Herrnstein, 1964a). Studies of preferences among various parameters of reinforcement schedules can be technically complex, because they must control for differences in time or responses per reinforcer in terminal links and for occasional biases toward particular colors or sides.

SELF-CONTROL

We usually speak of self-control when we forgo a relatively immediate consequence in favor of a later larger one. Both the immediate and the deferred consequences may be reinforcing or aversive (cf. Rachlin, 1974; Skinner, 1953). A small reinforcer now may be declined in favor of a large one later, as in skipping a small purchase now so as to save up for a larger one later. A small reinforcer now may be declined in favor of avoiding a large aversive event later, as when an alcoholic refuses a drink now and avoids a hangover later. A small aversive event now may be accepted in favor of a large reinforcer later, as when strenuous exercise prepares an athlete for a major event later. A small aversive event now may be accepted in favor of avoiding a large one later, as when a maintenance dental visit now avoids gum disease or toothache later. In each of these cases we may say that an individual has engaged in self-control or with self-restraint. If these kinds of arrangements characterize the situations we discuss in terms of self-control, we should be able to synthesize them in laboratory settings.

An example of behavior synthesis with concurrent-chain schedules is provided by the procedure in Figure 16-9 (Rachlin & Green, 1972). Initial links consisted of concurrent FR 25 FR 25 schedules; unlike the usual VI initial links, which equalize the pigeon's exposures to each terminal link by making them available equally often, the pigeon could enter one terminal link more often than the other with these schedules. In terminal-link A, the keys were dark for *T* s and were then lit red and green respectively. A peck on red immediately produced 2 s of food; a peck on green produced 4 s of food after a 4-s delay. In terminal-link B, the keys also were dark for *T* s, but after that only the green key was lit. As in the other terminal link, a

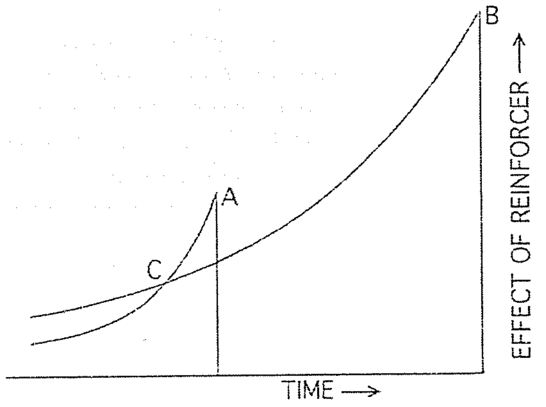

Figure 16-11 Delay-of-reinforcement gradients for an early small reinforcer at A and a later large one at B. Before time C the B gradient is higher than the A gradient, so the B reinforcer should be preferred; after time C the A gradient is higher, so the A reinforcer should be preferred.

peck on green produced the large reinforcer after a 4-s delay.

Confronted with both red and green in terminal-link A, the pigeon almost invariably pecks red, producing the small immediate reinforcer and not the large delayed one; this has been called *impulsiveness*. Confronted with only green in terminal-link B, the pigeon necessarily produces the large delayed reinforcer. But what about the pigeon's preference for A versus B, in its initial-link responding? The answer depends on *T*, the time until the terminal-link keys are lit. When it is short, say 1 s, the pigeon usually produces terminal-link A and then pecks red. When it is longer, the pigeon is more likely to produce terminal-link B, in which only green is available. Figure 16-10 shows the relative rate of pecking the left initial-link key (left initial-link pecks divided by total initial-link pecks) as a function of *T*. As *T* increased from 0.5 s to 16 s, the proportion of pecks producing terminal-link A decreased; the pigeon became more and more likely to enter terminal-link B and produce the larger reinforcer. Translating to a more colloquial vocabulary, when the opportunity to choose was coming up very soon, the pigeon was likely to be impulsive; when it came up later, the pigeon was likely to show self-control.

During initial links, the time to food equals *T* for the small reinforcer but *T* plus the 4-s delay for the large one. When *T* is short, this difference is relatively large (e.g., with a *T* of 1 s, the respective delays are 1 and 5 s) and the short delay outweighs the difference in reinforcer magnitudes. When *T* is long, however, the difference becomes relatively small (e.g., with a T of 10 s, the respective delays are 10 and 14 s) and the difference in reinforcer magnitudes becomes effective. By producing terminal-link B when *T* is long, the pigeon commits itself to the large reinforcer even though it wouldn't do so at the onset of green if red were also present. For this reason, pecks that produce terminal-link B have been called *commitment responses*: They guarantee the large delayed reinforcer by making the small immediate one unavailable. Other syntheses can be created with other temporal arrangements (e.g., see Mazur, 1996, on procrastinaton by pigeons).

Figure 16-11 illustrates these relations in terms of two hypothetical delay-of-reinforcement gradients produced by a relative small reinforcer at A and a later large reinforcer at B. If at any time the pigeon's peck is determined by which gradient has the higher value, then where the A gradient is higher as the pigeon approaches time C, it will commit to reinforcer B, showing self-control. At time C the gradients cross over so the A gradient is now higher; between times C and A it will choose reinforcer A, showing impulsiveness.

With the concurrent-chain procedure we can examine impulsiveness and commitment with immediate and delayed reinforcers or with immediate and delayed aversive stimuli (e.g., Deluty, 1978). We can also examine whether commitment is modifiable and whether the components of this synthesis of self-control involve variables comparable to those that operate in human behavior (e.g., Grosch & Neuringer, 1981). We humans are far more likely than pigeons to forego small and currently available reinforcers in favor of large delayed ones. Human instances of self-control presumably also involve verbal behavior (cf. Chapter 23). But if more complex processes operate in human self-control, they can only be identified by examining whether human cases are consistent with analyses in terms of reinforcer delays. For this reason, these behavior syntheses provide an essential reference performance for the analysis of self-control and illustrate the relevance of reinforcement schedules to human behavior.

BEHAVIOR SYNTHESIS: NATURAL FORAGING

Concurrent-chains have been broadly applied to the synthesis of complex behavior. If the interpretation of complex behavior in a natural habitat suggests that it consists of several simpler components, the interpretation can be tested by trying to assemble those components in a laboratory setting. A successful synthesis supports the interpretation; an unsuccessful one may reveal inadequacies in the assumptions about what was going on in the natural setting.

In the field of *behavioral ecology*, this strategy is illustrated by studies of natural foraging (e.g., Fantino & Abarca, 1985; Kamil, Yoerg, & Clements, 1988). In their foraging, animals in the wild travel from one patch of food to another, staying or moving on to new ones depending on what they find. For example, a bird might fly to a bush in which the eggs of an edible insect have just hatched. As it eats, it gradually depletes its prey, and when it moves on depends on such factors as how much is left, how far it must go to find another bush and what the chances are of finding other food there (cf. Wanchison, Tatham, & Hineline, 1988).

Because analyses of natural foraging involve switching among food sources as food availability or other conditions progressively change, as in the last example, progressive schedules have been useful tools (Hackenberg, 1992; Neuman, Ahearn, & Hineline, 2000). In a progressive schedule, some parameter of the schedule is varied over successive reinforcers or blocks of reinforcers. By analogy to the cost of switching over from one food source to another, some procedures provide an alternative response that allows the organism to reset the schedule back to its starting value. For example, a progressive FI schedule arranged for a pigeon's pecks on one key might start with a 30-s FI and add 15-s to the FI value with each reinforcer, but completion of an FR on a second key might reset the progressive FI back to 30-s. The FI value at which the pigeon resets the progressive schedule will vary with the size of the required ratio, so this schedule can be used to study how the pigeon trades off time (FI) and number (FR).

Some factors within natural ecological settings can be simulated within concurrent chains. For example, varying the schedules that operate in initial links is analogous to varying the time and effort involved in traveling from one bush to another, and varying the schedules in terminal links is analogous to varying the availability or depletion of different food sources at different sites. Concurrent-chain schedules in the laboratory that simulate those in natural habitats have revealed some properties of foraging. For example, organisms are less selective in the food they accept if they spend more travel time (more time in initial links) between potential food sources; and if one food is preferred over another, the availability of the preferred food (the schedule that operates during the terminal link in which that food is the reinforcer) is a primary determinant of the choice of food patches (as shown by initial-link responding). In other words, natural foraging may be treated in terms of concurrent-chain schedules, and properties of natural foraging, in turn, may suggest variables that are important in concurrent-chain performances.

A SCHEDULE TAXONOMY

We've seen that reinforcement schedules are tools that can be applied to the study of a variety of behavioral phenomena relevant to human concerns. Our examples have ranged from causal relations between behavior and environment to self-control and freedom of choice. Yet even so-called simple schedules aren't simple; the complexity of schedule effects has made schedule analysis highly technical. We examined properties of ratio and interval schedules in the last chapter; in this one we've briefly sampled other topics, including multiple, chained, second-order, concurrent and concurrent-chain schedules. We've often omitted procedural details. This is perhaps inevitable. Unlike most other areas in this text, reinforcement schedules didn't even exist as a systematic subject matter until relatively recently (Ferster & Skinner, 1957; Skinner, 1956). One concern of this subject matter has been maintenance of behavior in the steady state, and yet the change in behavior that accompanies any transition from one schedule to another is an instance of learning.

We examined multiple and mixed, chained and tandem and second-order, and concurrent and concurrent-chain schedules, but we haven't exhausted the possibilities for combining the basic schedules. For example, reinforcers can be arranged for completing *either* of two schedule

Table 16-1 Schedule combinations.
For convenience, each case is defined in terms of just two arbitrary component schedules, A and B, but combinations can include any number of components. Stimuli are designated by S with a superscript that identifies the schedule each one accompanies.

Schedule	*Definition*	*Example (with Abbreviation)*
Multiple	A during S^A alternates with B during S^B.	(A) VI during red alternates with (B) EXT during green (mult VI EXT)
Mixed	A alternates with B (like multiple but without different stimuli).	(A) DRL alternates with (B) FI, with no correlated stimuli (mix DRL FI)
Chained	During S^A, completion of A produces S^B; during S^B, completion of B produces reinforcer.	Completing (A) VR during blue produces yellow; completing (B) FR during yellow produces food (chain VR FR)
Tandem	Completing A produces B; completing B produces reinforcer (like chained but without different stimuli).	Completing (A) VR produces (B) DRH, and completing DRH produces food, with no correlated stimuli (tand VR DRH)
Concurrent	A operates for one response; B operates for another response.	(A) One VI operates for pecks on a left key and (B) another VI operates for pecks on a right key (conc VI VI)
Conjoint	A and B operate independently for a single response (like concurrent but without different responses).	(A) VI and (B) avoidance operate simultaneously for presses on a single lever (conjt VI avoidance)
Second-order	Completing A is reinforced according to B; reinforcing second order schedule according to C creates third order schedule, etc.	(A) Successive FRs treated as response units reinforced according to (B) FI schedule (FI [FR])
Progressive	Some schedule parameter changes systematically over successive reinforcers or blocks of reinforcers.	After every *n*th reinforcer, *t* s is added to the value of an FI (progressive FI); may include reset contingencies
Interlocking	Reinforcer depends on completing some combined function of A and B.	Responding is reinforced when the sum of (A) responses plus (B) elapsed seconds reaches some constant value (interl FR FT)
Alternative	Reinforcer depends on completing either A or B.	Responding is reinforced on satisfying either (A) VR or (B) VI contingencies, whichever occurs first (altern VR VI)
Conjunctive	Reinforcer depends on completing both A and B.	Responding is reinforced on satisfying both (A) FR and (B) FI contingencies (conjunc VR VI)

requirements (*alternative* schedules) or *both* of them (*conjunctive* schedules). In an alternative FI 30-s FR 50 schedule, either the first response after 30 s or the 50th response is reinforced, whichever comes first. In a conjunctive FI 100-s FR 20 schedule, a response isn't reinforced until both 100 s pass and 19 responses have already been emitted. In a *progressive* schedule, some parameter of a schedule changes systematically over successive reinforcers or blocks of reinforcers. For example, a ratio might increase by 10 responses after every 5th reinforcer (sometimes a second response is available that resets the progression to some starting value according to some schedule). *Adjusting* schedules vary as a function of some property of performance, as when a ratio varies in proportion to the last postreinforcement pause, or as when shock delivery changes the RS interval of an avoidance schedule. A schedule in which time and number requirements interact is an *interlocking* schedule. For example, an interlocking FR FI schedule might shorten the interval as a function of number of responses, or lengthen the ratio as a function of time. An example is winding a grandfather clock, in which the reinforcer is the tension of the fully wound spring; until the spring is completely run down, the number of turns required to wind it increases as time passes.

For convenience of reference, the major schedule combinations are summarized in Table 16-1. The table provides definitions, examples and standard abbreviations. The formal designations may make some relations among the schedules more obvious. For example, multiple and mixed schedules both involve the alternation of component schedules, and chained and tandem schedules both involve sequences in which completion of one schedule produces another. In both pairs, the only distinction is whether the component schedules each operate during different stimuli or during a single stimulus. Similarly, the only distinction between concurrent and conjoint schedules is whether the two simultaneous schedules operate for different responses or for a single response.

These schedule combinations are our tools. Once we've used schedules to analyze complex behavior we may be in a position to use these tools to put the parts back together in a synthesis. For example, we can test our interpretation of behavior in a natural habitat by trying to assemble its components in a laboratory setting. We can't create a behavior synthesis without making explicit our assumptions about the properties of the behavior we're trying to synthesize. For that reason, when we attempt synthesis we probably profit more from our failures than from our successes. In fact, it may be a general principle of scientific research that we learn the most when our experiments produce data we didn't expect. After all, what's the point of doing experiments if we always know exactly how they're going to turn out?

ADDENDUM 16A: BEHAVIORAL ECONOMICS

Behavioral economics began when a few behavior analysts recognized that some performances and concepts derived from schedules of reinforcement were relevant to properties of the behavior studied by economists (e.g., Rachlin *et al.*, 1976). One important point of contact was that behavior analysts and economists had a mutual interest in choice. Some economists have since adopted behavioral economics as their own, but economic behavior should share properties with other varieties of behavior, so it is likely to be profitable to study how this behavior arises out of simpler behavior. Schedules of reinforcement provide useful tools for such an endeavor. Here it is feasible only to briefly review a few sample concepts from behavioral economics and to briefly suggest how studies of them can be realized experimentally (cf. Bickel, Green, & Vuchinich, 1995; Green & Freed, 1998; Hursh, 1980, 1991).

One familiar feature of economic theory, of course, is supply and demand. It is obvious enough that an experimenter can arrange a supply of some reinforcer, such as food, and that the demand for food by an experimental organism can be altered by establishing operations such as deprivation and satiation. The experimenter can manipulate the cost of food as a commodity by changing response requirements. For example, a reinforcer produced by a fixed ratio of 100 responses can be thought of as twice as expensive as one produced by a fixed ratio of 50.

Not so obvious is the difference between a situation in which the only food reinforcers available to the organism are those it earns within an experimental session and one in which what it earns is supplemented by additional feeding outside of the session. The former is analogous to what economists called a closed economy and the latter to what they call an open economy (according to this distinction, the data presented in Figure 15-3 were obtained under an open economy). The results of experimental comparisons depend on a variety of experimental parameters, such as schedule parameters and type of schedule, but some experiments that produced increases in response rate with increases of reinforcement rate in open economies produced opposite results with closed economies.

Another economic concept is elasticity. When the demand for a commodity changes substantially with changes in price, its demand is said to be elastic; to the extent that demand does not change flexibly with changes in price, it is said to be inelastic. For example, the demand for leisure items that one can easily do without is highly elastic relative to the demand for essentials such as electricity or fuel. In experiments arranged as operant analogues of elasticity, the behavior maintained by reinforcers can vary differently given changes in their cost if the alternatives differ substantially in elasticity. For example, if an organism is food deprived and has also become dependent on opiates, a doubling of a fixed ratio requirement arranged for food reinforcers will not have the same effect as the doubling of an equal fixed ratio requirement arranged for drug reinforcers.

One determinant of elasticity is substitutability, i.e., whether some reinforcers will substitute for others. For example, if the cost of coffee rises, the effect on demand for coffee will be different in a setting where some other source of caffeine such as tea is easily available than in a setting with no other source of caffeine. Substitutability is relevant in experiments that involve different responses maintained by different reinforcers, especially when the reinforcers share properties (e.g., different foods containing some common ingredients).

Substantial attention within behavioral economics has also been given to discounting (Green & Myerson, 2004). If you have a choice between getting $100 now and getting $100 next year, you will undoubtedly choose the immediate $100. But what about $100 now versus $120 next year? If that is not enough, how large would the delayed amount have to be for you to prefer it? This example illustrates that we typically discount the value of delayed consequences relative to immediate ones. If $200 in one year is worth just $100 right now, the delayed $200 has been discounted by 50 percent. Studies of human discounting typically examine verbal judgments, but such procedures have analogs in nonhuman studies that pit small immediate consequences against larger delayed ones (Kahneman, 2011).

Studies of discounting and other economic analogs often summarize findings in terms of changes in the value of consequences: $100 delayed for a year has a far lower value than $100 available immediately. A problem with interpretations of economic behavior in terms of changes in value is that descriptions in such terms may leave out the behavior. Though that may make it easier to communicate with economists and those in other disciplines, the advantages of behavioral treatments and interpretations may be lost or unrecognized when their main strengths are obscured by such verbal transformations.

Furthermore, verbal estimates differ from choices that are products of actual contingencies. This means that what people say they will choose when hypothetical alternatives are described to them may not agree with what they choose when confronted with real alternatives. Whether verbal mediation is involved or not, we cannot assume that choices such as these or economic behavior in general are based on rational decisions. Instead, we must examine the contingencies. For example, rationality is not involved when small immediate reinforcers are selected in preference to large delayed ones. Words matter, of course, but in the analysis of behavior contingencies are the bottom line.

ADDENDUM 16B: SCHEDULES AND ATTENTION DEFICIT HYPERACTIVITY DISORDER

Schedules of reinforcement have lent themselves to a many applications (e.g., screening for effects of environmental pollutants on behavior, or arranging educational contingencies in the classroom for children with developmental disabilities). Some of the most interesting applications are those in which a variety of different schedule phenomena are brought together to clarify complex behavior. One example (Sagvolden *et al.*, 2005) is in the analysis and interpretation of the components of attention-deficit hyperactivity disorder (ADHD).

Children diagnosed with ADHD show one or more of a spectrum of behavioral symptoms. They may be hyperactive in school or home settings; they may be inattentive to instructions or school materials; and they may behave impulsively in situations where other children exhibit self-control. Some children may show all of these symptoms, but others may show only one or another. Each may be understood in terms of delay-of-reinforcement gradients that decrease more steeply than ordinarily gradients, so that reinforcers that do not follow responses closely have less impact on responding than is ordinarily the case (perhaps drugs that meliorate ADHD symptoms allow delayed reinforcers to be effective over longer time spans).

Delay-of-reinforcement gradients describe how the effects of reinforcers vary as a function of the time separating them from the responses that preceded them. Suppose a standard gradient in a given context declines slowly enough over time that reinforcers are still somewhat effective even if they follow 10 or 15 seconds after a response, but that ADHD gradients decrease so rapidly that reinforcers are ineffective unless they follow responses by no more than 5 or so seconds.

Consider first a response sequence like crossing a room to find a toy. A child might do this either quickly or slowly. Starting up the sequence would participate in the effects of the reinforcer either way given a standard gradient. With the ADHD gradient, however, the toy as a reinforcing consequence would work on the entire behavior sequence only if the child found it quickly. This is differential reinforcement: Given rapid responding, the reinforcer is effective with respect to the entire response sequence, but given slow responding it is effective only with respect to the end of the sequence. Thus, faster sequences will be selected by their consequences. The differential strengthening of rapid responding takes time, which may also explain why hyperactivity often takes a while to develop and develops separately in different environments.

In any case, hyperactivity is an appropriate name for behavior in which rapid sequences have displaced more leisurely ones. Of course, if an ADHD delay gradient declines steeply enough that only a single response can ever be within reach of the reinforcer, rapid sequences cannot be differentially reinforced, so hyperactivity will not emerge. The point is that gradients that differ in steepness can produce different degrees of hyperactivity.

The onset of a stimulus that sets the occasion for responding may be followed by a reinforced response after a shorter or a longer delay. If reinforcers are delivered in its presence, the stimulus becomes a conditional reinforcer, but its potency depends on the delay. Observing responses are maintained by such conditional reinforcers. But given a steep ADHD gradient, the onset of stimuli in the presence of which responses produce reinforcers only some time later may not be effective, so these stimuli will be less likely to be looked at or listened to. A child who does not look or listen appropriately might be described as having an attention deficit. The magnitude of this too will depend on the steepness of the gradient, so different degrees of hyperactivity and of attention deficit might be evident in children with ADHD gradients of different steepness.

The child for whom the delay gradient is unusually steep will be less able to deal with or tolerate longer delays, which is perhaps why the attention of such children is so easily captured by computer games, which typically provide very rapid feedback. It is useful to know that in reinforcing the behavior of such children the reinforcers must be delivered promptly. Consider also the inverse relation between impulsivity and self-control. Forgoing a small immediate reinforcer for a later larger one is described as self-control, but whether this happens depends on the steepness of the delay gradient. The steeper the gradient the less potent the larger reinforcers that follow longer delays, so the greater the likelihood of impulsive behavior.

In other words, these three different manifestations of ADHD, hyperactivity, attention deficit and impulsivity, may all be derivatives of a single variable, the steepness of the delay gradient. Individual differences in the steepness of the delay gradient could account for variations in the dominance of one or another of these symptoms over others in different children (e.g., Catania, 2005a, 2005b). The case for an interpretation of ADHD in terms of delay gradients in Sagvolden et al. (2005) is also strengthened by comparisons of the schedule-maintained behavior of Wistar Kyoto rats (WKY) and a closely related strain of spontaneously hyperactive rats (SHR). The SHR rats are a useful model for ADHD in children, not only because their behavior is affected by the same drugs used to treat hyperactivity in children but also because schedules can be used to assess the differences between the delay gradients of the WKY rats and the SHR rats. This example makes it evident that beyond the study of particular schedules of reinforcement lies the possibility of coordinating data drawn from a variety of types of schedules to enrich our understanding of complex human as well as nonhuman behavior.

Chapter 17
Respondent Behavior: Conditioning

Reflex is derived from the Latin re-, *back, plus* flectere, *to bend. A reflexive response was thought of, in effect, as a reflection of the stimulus. It is debatable whether* flectere *and the Latin* plicare, *to fold, share a common Indo-European root. If they do,* reflex *is closely related to* reply, complex *and* multiple.

Responses, depending on their relations to eliciting stimuli, consequences, establishing operations and discriminative stimuli, are variously said to be elicited, emitted, evoked *or* occasioned. *All four terms have Latin roots. The first three share a prefix abbreviated from* ex-, *out:* elicit, *derived from* laqueus, *lure or snare, is related to* delight *and* latch; emit, *derived from* mittere, *to send or let go, is related to* omit *and* intermittent; evoke, *derived from* vocare, *to call, is related to* vocal *and* invoke. Occasion, *derived from the prefix* ob-, *against, and* cadere, *to fall, is related to* case, accident, chance *and* coincidence.

Historically, respondent conditioning has so dominated the psychology of learning that it has provided the opening chapters of many learning texts, and it is still typically treated before operant learning in many introductory textbooks. It has also been variably called *classical* or *Pavlovian* conditioning. The language of conditioned reflexes has to some extent entered everyday vocabulary, although in its popular usage it is often confused with instances of operant learning. Many still speak indiscriminately about both operant conditioning and respondent conditioning; a case for some common underlying processes can be made (Donahoe & Palmer, 1994), but the persistence of that usage runs the risk of confusing two very different procedures.

The term *conditioned* came from the Russian in the phrase for conditioned reflexes, условный рефлекс (*uslovnyi refleks*). But the term should have been translated as *conditional*, because it was

KEY TERMS: Conditional or Conditioned Reflex; Conditional or Unconditional Stimuli and Responses; Contingency and Contiguity; Simultaneous, Delay, Trace, Temporal, Backward and Differential Conditioning; Pseudoconditioning; Autoshaping.

applied to reflexes conditional upon relations among environmental stimuli. As with operant phenomena, vocabulary has consequences: in the language of conditioning, it is too easy to talk about the organism as being conditioned while neglecting the stimulus-stimulus contingencies that led to the changes in its behavior.

Respondent conditioning is an instance of stimulus control applied to stimulus presentations rather than to consequential contingencies. In other words, instead of signalling the consequences of responding, the stimulus simply signals that some other stimulus will be presented. Pavlov's conditioned salivary reflexes are the prototype example: when a bell repeatedly signalled food to a hungry dog, salivation began to be elicited by the signaling stimulus as well as by the food itself.

Ironically, Pavlov may never have used a bell in his experiments; his rare mention of bells occurs only in later work, and there probably refers to electrically operated devices. The Russian word used by Pavlov for bell, ЗВОНОК (*zvonok*), also has buzzer and doorbell as its synonyms; Pavlov did not use another Russian word for bell, КОЛОКОЛ (kolokol), which does not have these synonyms. In addition to this ambiguity of translation, the ubiquitous references to Pavlov's bell may have originated with the common use of the dinner bell and salivation as examples in the writings of John B. Watson and others. In 1993 I had the good fortune to visit Pavlov's apartments in St. Petersburg in conjunction with a meeting of the Language Origins Society. His apartments have been converted to a Pavlov museum. The only bell there was one on his desk, presumably for summoning a servant or an assistant.

Individual responses are never exactly repeated, so we speak of operant classes of responses instead of individual instances. Similar issues arise for elicited behavior. For example, successive elicitations of salivation by food may differ in latency, quantity, viscosity and other properties. It is therefore appropriate to extend the language of classes to responses defined by the stimuli that produce them. These classes, called *respondents*, correspond to the behavior we earlier called elicited or reflexive. Thus, salivation produced by food in the mouth is one respondent class; it must be distinguished from salivation produced by acid in the mouth, which is another respondent class, and from spontaneous salivation, which isn't a respondent class at all because it has no eliciting stimulus (spontaneous salivation is emitted rather than elicited; if we could identify an eliciting stimulus, we wouldn't call it spontaneous).

CONDITIONAL REFLEXES

Stimuli produce respondent behavior. Change the stimuli and the behavior changes. For example, different acidic concentrations on the tongue (e.g., dilutions of vinegar) elicit different quantities of saliva. This means that we are limited in modifying respondent behavior. The properties of respondents are so highly determined by their eliciting stimuli that no procedure equivalent to the creation of new operants by shaping exists for them. But we can alter the eliciting effects of stimuli. Let's examine Pavlov's (1927) procedure in more detail.

We begin with a dog in a harness with one of its salivary ducts connected to a system that records salivation. We use two stimuli: the sound of a buzzer and food in a form that can be delivered directly into the dog's mouth. First we examine the effects of each stimulus separately. When we first sound the buzzer, the dog pricks up its ears and turns its head toward it. This has been called an *orienting* response. It diminishes with successive soundings of the buzzer, perhaps even becoming undetectable; it can be reinstated by waiting a while before sounding the buzzer again (cf. habituation, Chapter 4). When food is placed in its mouth, the dog chews and salivates. These responses may diminish a little over food presentations, but their magnitude remains substantial over the course of a session.

Now let the buzzer signal food by sounding 5 s before each food delivery. After a number of trials salivation sometimes begins in the 5-s period between buzzer and food, and sometimes the buzzer is followed by salivation even on an occasional trial when the food is omitted. In neither case can we say that the food elicited the salivation: in the first the salivation began before the food was presented and in the second the food wasn't presented at all. Now that the buzzer elicits salivation, we say that we've created a new respondent class, salivation elicited by the sound of the buzzer. We call the relation between buzzer and salivation a conditional reflex because it is conditional on the relation between buzzer and food.

The sequence of events is illustrated in Figure 17-1. The buzzer at first elicits orienting responses, but these disappear with repeated presentations; at this point, the buzzer is a *neutral stimulus* (NS). In an unconditional reflex, food elicits salivation; in this relation, the food is an *unconditional stimulus* or *US* and the salivation is an *unconditional response* or *UR*. Conditioning begins when the buzzer reliably precedes food; at this point, the buzzer still has no effect on salivation and can still be regarded as a neutral stimulus. After a period of conditioning, a *conditional reflex* has been created; the buzzer elicits salivation before food is presented (a), or the buzzer elicits salivation even when food is omitted on an occasional trial (b). The buzzer is now a *conditional stimulus* or *CS* and salivation elicited by the buzzer is a *conditional response* or *CR*.

The difference between a conditional stimulus and an unconditional stimulus isn't simply which comes first. If we reversed their order, the eliciting effects of food followed by buzzer wouldn't be much different from those of food alone. In fact, we can often predict the relative effectiveness of stimuli as CSs and USs from the probabilities with which they elicit their respective responses (cf. Chapter 6 on the relativity of reinforcement). For example, consider salivation elicited by food in a dog's mouth and leg flexion produced by shock to its leg. A mild shock may become a conditional stimulus eliciting salivation if reliably followed by food, but a strong shock followed by food isn't likely to do so. On the other hand, food may

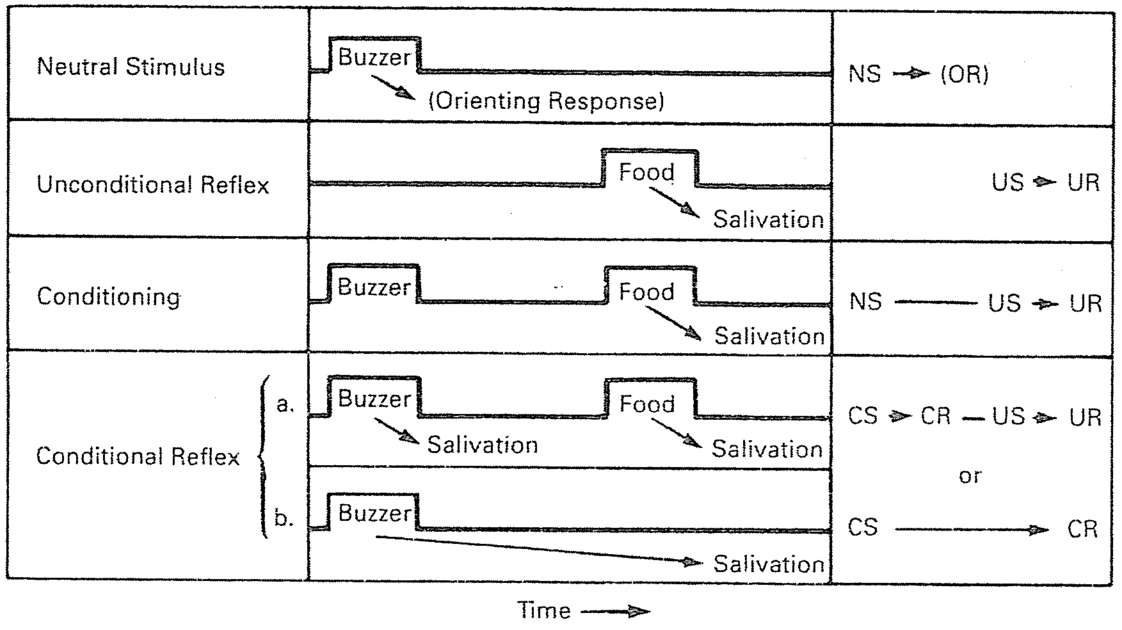

Figure 17-1 Relations between stimuli and responses in respondent conditioning. An initially neutral stimulus (BUZZER:NS) is followed by an unconditional stimulus (FOOD:US) that elicits salivation. Once the neutral stimulus begins to elicit responding like that elicited by the unconditional stimulus, the neutral stimulus is called a conditional stimulus (CS). OR = orienting response; UR = unconditional response; CR = conditional response.

become a conditional stimulus eliciting leg flexion if reliably followed by a strong shock, but food followed by a weak shock isn't likely to do so. The differences are consistent with the relative probabilities with which food may elicit salivation and mild or strong electric shock may elicit leg flexion.

The CR and the UR differ in more than temporal order. For example, the form or topography of a conditional leg flexion typically differs from the unconditional flexion elicited by electric shock. In general, a CR isn't merely a UR now elicited by a new stimulus; in other words, **respondent conditioning cannot be interpreted as stimulus substitution.** In Pavlov's classical case, for example, the buzzer doesn't substitute for food (the dog doesn't try to eat the buzzer). A CS can affect a range of responses, and not just those resembling the ones elicited by the US.

Many different conditional reflexes have been created through respondent procedures (e.g., Hull, 1934). Pavlov's salivary conditioning is probably the most familiar, but other studies demonstrated conditioning based on such unconditional reflex relations as the knee jerk elicited by a tap on the patellar tendon (Twitmyer, 1974) and limb withdrawal elicited by electric shock (Bechterev, 1933). Conditioning has worked with eyeblinks elicited by puffs of air to the eye (Gormezano, 1972). For many years conditioning of pupillary constriction elicited by light to the eye (Stanners *et al.*, 1979) had been unsuccessful (Young, 1958), but more recent research has reported conditioning of pupillary dilation (Reinhard & Lachnit, 2002). The Russian literature includes varied demonstrations of conditioning ((e.g., Bykov, 1957). For example, for several days a dog was placed in a waiting area at neutral temperature before being moved into a heated room, and its metabolism and oxygen consumption began to decrease in the waiting area as well as in the heated room. Inversely, after time in the waiting area was followed by being moved into a cool room, the dog's metabolism and oxygen consumption began to increase in the waiting area and not just in the cool room (Bykov, 1957, pp. 183-210). Conditioning effects have been demonstrated with a range of physiological responses, such as reactions of the immune system: (Ader & Cohen, 1985).

Consider another example. The release of insulin by the pancreas is a UR produced by the US of sugar in the gut (Deutsch, 1974). This US is reliably preceded by the taste of sugar, which makes it likely that the taste will become a CS that releases insulin. Suppose now that you've just switched from consistently drinking soft drinks with sugar in them to their sugar-free diet versions. As a CS, the sweet taste of the drink elicits the release of insulin, which is ordinarily used up as you digest the sugar. But now your drink has no sugar in it, so you can expect to feel weak or heady as the insulin produces a substantial drop in your blood sugar level (hypoglycemia). This effect was especially common when sugar-free drinks were first introduced and people were familiar only with the standard varieties. At the time students reliably reported it when I asked about it in class. The effect is hardly noticeable nowadays, however, because just an occasional substitution of a sugar-free for a standard drink can weaken it considerably.

Respondent conditioning with some types of USs leads to CSs that elicit compensatory responses, responses that counteract the effects of the US, rather than responses similar to those elicited by the US. Such cases provide another kind of evidence that respondent conditioning isn't merely stimulus substitution. One example occurs in respondent conditioning with opiates such as morphine or heroin (Siegel, 1977). Among the effects of these drugs is analgesia, an elevated pain threshold. With continued doses, events that lead up to using the drug (e.g., preparing the needle) are likely to become CSs that elicit a CR. But this CR doesn't enhance the analgesia and other drug effects; instead, it counteracts them (for example, it produces hyperalgesia, a lowered pain threshold or, in other words, the opposite of analgesia). As a result, larger and larger doses are needed to produce the original drug effect (these effects may contribute to the development of drug tolerance). To summarize: The US is the drug in the bloodstream and one component of the UR is analgesia; the CS is any event reliably preceding the drug administration and the CR is a physiological response that counteracts the analgesia and other effects of the drug.

Heroin addicts often take their drugs in the same place in the same company, using a consistent drug ritual. Now consider the addict who for some reason takes the drug somewhere else and in different company. The dose is large but many of the CSs that usually precede it are absent, so a CR much smaller than usual is elicited. The drug effect occurs, but this time it isn't counteracted by the usual compensatory response. Under such circumstances, a drug dose that is ordinarily tolerated can instead be fatal (Siegel *et al.*, 1982); hospital admissions and/or deaths from heroin overdose are especially likely when addicts take their drugs under unusual or unfamiliar circumstances.

The Role of Elicited Responses

Most accounts of the crucial determinants of respondent conditioning appeal to the relation between the CS and the US, without much regard to the role of the UR. This seems to make sense, first because the experimenter arranges the timing of only the CS and the US, and second because the timing of the UR is determined by the timing of the US that elicits it. But what if the temporal relations between the CS and the US and between the CS and the UR could be separated so that each could be varied independently? Of the CS-US interval and the CS-UR interval, can we test which is more important in respondent conditioning?

The key to conducting such an experiment is identifying a reflex preparation in which some US very reliably produces a UR but with a long enough latency that the CS can consistently be presented after the US but before the UR. With such a preparation, various degrees of temporal overlap between CS and UR can also be arranged. The discovery of a preparation that allowed such timing (Donahoe & Vegas, 2004) set the stage for asking whether conditional responding depends on the CS-US relation, as traditionally assumed, or instead depends on the CS-UR relation. This reflex preparation used water to elicit a throat movement in pigeons. A brief squirt of water into the pigeon's mouth was reliably followed some 300 or 400 milliseconds later by the throat movement, which is a component of swallowing.

We may think of our own swallowing as operant behavior, but only the initial portion is. To swallow we must move what is in our mouth toward the throat, where it makes contact with the back of the palate. The contact elicits the swallow, a sequence of responses that transports material to the esophagus. That is why you cannot swallow with a dry mouth if you have nothing in it; without saliva, you have nothing else to send back there. Human swallowing is particularly complex because eating and breathing and talking all share some of the same anatomy, and swallowing therefore involves such elicited components as movements of the vocal chords and other nearby structures that protect the airway from the aspiration of fluids or food particles.

The throat movement parts of the pigeon's swallowing reflex were recorded by positioning a very light counterweighted rod on the bird's neck and amplifying the slight movements of the rod produced by each throat movement with an optical system. A green light near the pigeon's left eye served as a CS, the squirt of water as a US, and the elicited throat movements as either CR or UR. With this system in place, initial procedures first compared standard respondent conditioning trials, in which the CS was consistently followed by the US, with trials in which either the CS alone or the US alone was presented. Throat movement responses were obtained as CRs with CS-US intervals of either 1 second or 1.5 seconds but conditioning did not occur when the CS and the US each occurred alone in separate trials. It was important to show that the pigeon's swallowing reflex was an appropriate preparation for the study of respondent conditioning and these preliminary procedures did so.

The temporal arrangements in the procedures that followed included timing the CS so that it came after the US but before the UR, timing it so that it came after both the US and the UR, and timing it so that it overlapped with the UR. Conditional responding that depends on the experimental CS and US arrangements may be indistinguishable from unconditional responding that is elicited by the US when the CS is presented at or near the same time as the US or the UR, so the effectiveness of the CS was assessed by presenting it alone in occasional probe trials.

The CS elicited salivation even when it followed the US, as long as it preceded or overlapped with the UR. Conditioning was also obtained with a puff of air in the mouth substituted for water as the US. This stimulus too elicited swallowing and allowed conditioning and therefore showed that the results were not an artifact of traces of water left in the mouth by the US. Overall, the implications were that the CS-UR relation matters more than the CS-US relation or, in other words, that the environment is significant mainly in the context of the behavior that occurs within it.

Questions remain about extending these findings to other respondent systems, especially when those systems involve CRs that differ from URs (as in drug dependence and overdose), but divergences in the outcomes of conditioning with different systems may also depend on their temporal properties. For example, different degrees of temporal overlap between CSs and URs are possible when USs elicit long-lasting emotional responses, as in conditioned suppression, than when they elicit brief motor responses, as in the conditioning of eye blinks. Other differences may arise when such factors interact with individual differences in responsiveness to various classes of stimuli, as when the effects of a traumatic experience are more severe and/or more long-lasting in some individuals than in others (e.g., Mineka & Zinbarg, 2006). The implications for the analysis of emotions and for clinical practice may be substantial (cf. Friman, Hayes, & Wilson, 1998).

Types of Conditioning

Experiments ordinarily specify only the temporal relations between the CS and the US, and these can be arranged in varied ways. Situations with the onset of the CS preceding that of the US by no more than 5 s are sometimes arbitrarily grouped together as instances of *simultaneous conditioning*. This usage originated because the optimal interval between CS and US is roughly half a second (Kimble, 1947), short intervals like half a second provide too narrow a window for observing conditional responding, and omissions of the US to allow such observation usually reduce the reliability of conditional responding. Thus, lengthening the interval between stimulus onsets became favored over occasionally omitting the US, so distinctions within this range of delays were overlooked and all such cases were called simultaneous.

The effect of occasionally omitting the US had theoretical ramifications. In a usage that has become rare, presentations of the US in respondent conditioning were called reinforcement and the procedure in which this stimulus was omitted on occasional trials was called partial reinforcement. Distinctions between operant and respondent conditioning were then argued on the basis of the so-called partial reinforcement effect or *PRE*: Relative to reinforcement of every response, partial reinforcement generated substantial quantities of responding in operant behavior, as we saw in Chapter 15, whereas the analogous procedure in respondent conditioning usually reduce responding (but see Gibbon *et al.*, 1980). Now that the term reinforcement has become more restricted in its scope, the comparison no longer seems relevant. In its time, however, it provided one of the most persuasive grounds for distinguishing between the two types of conditioning.

Some arrangements of CSs and USs are contrasted with simultaneous conditioning in Figure 17-2. In both *trace conditioning* and *delay conditioning* a relatively long time elapses between the onset of the CS and that of the US; they are distinguished by whether the CS turns off or remains present during that time. (The time from CS onset to US onset can vary independently of the temporal overlap of CS and US. For example, in delay conditioning the CS might end with US onset or end at the same time as the US; this feature of

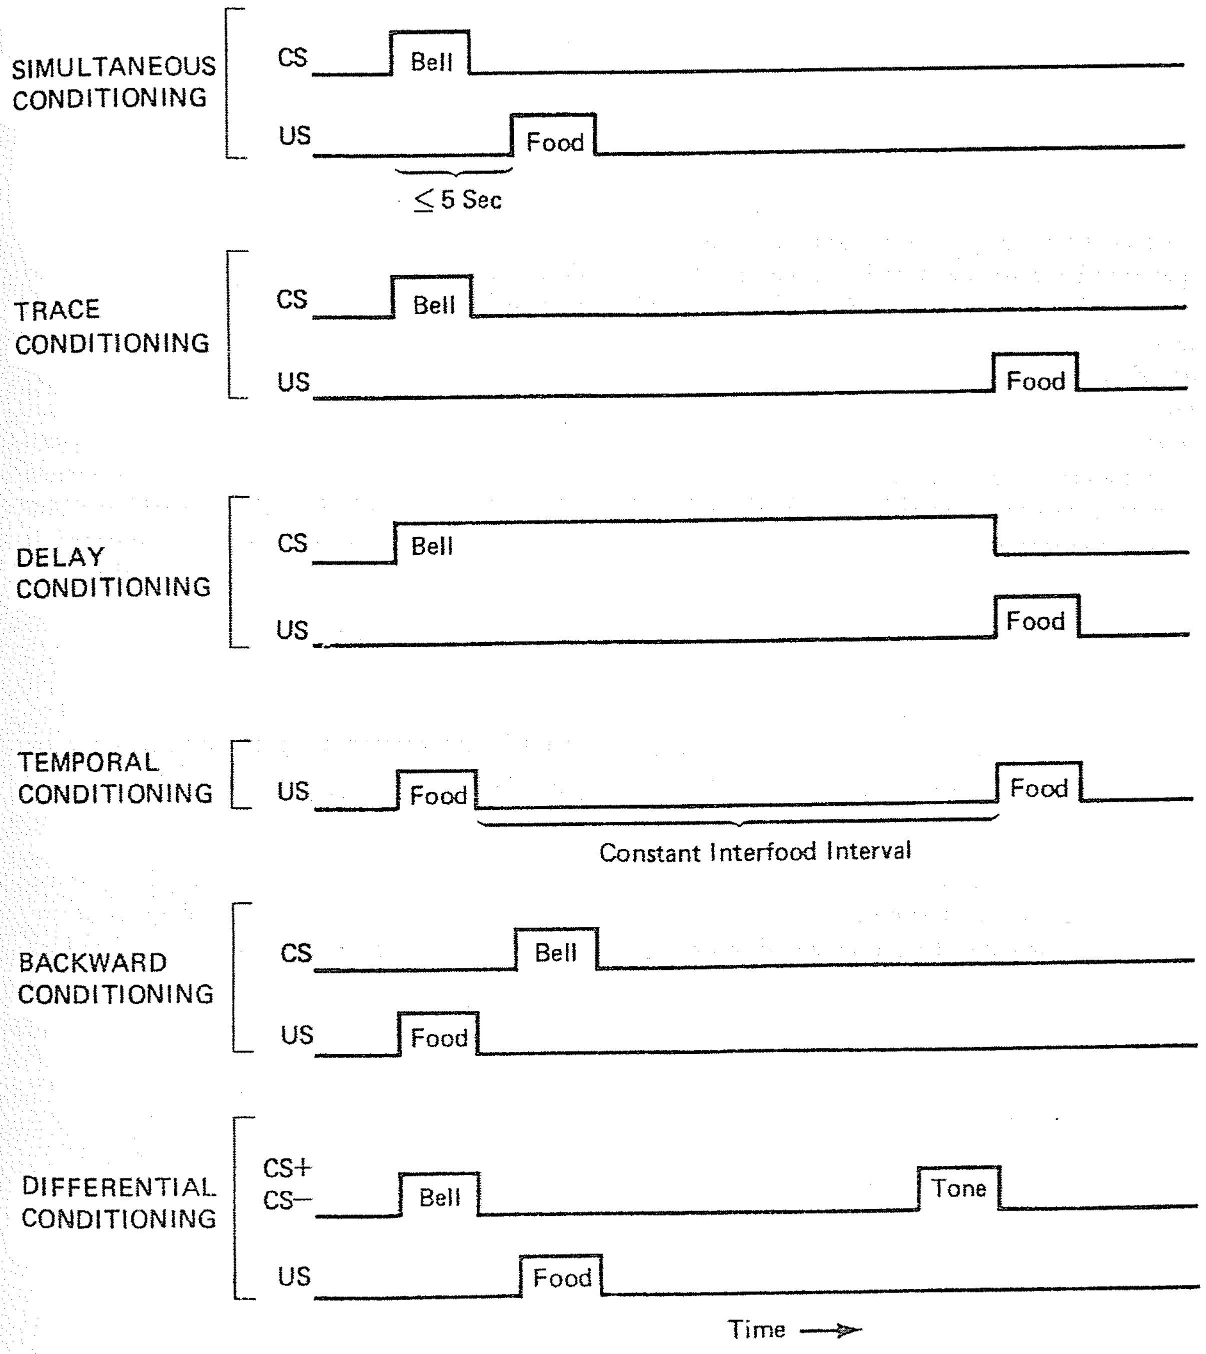

Figure 17-2 Schematic presentation of various respondent procedures, showing a buzzer as CS and food as US. In simultaneous conditioning, the buzzer is followed within less than 5 s by food. Different time relations are illustrated in trace, delay and temporal conditioning, and the stimulus order is reversed in backward conditioning. In differential conditioning, the buzzer is followed by food but a tone isn't. CS = conditional stimulus; US = unconditional stimulus.

the timing of CS and US isn't relevant to the procedural distinctions in Figure 17-2.) In both trace and delay conditioning, conditional responding at first occurs shortly after the onset of the CS, but with successive trials it gradually moves closer to the time at which the US is to be delivered. Trace conditioning acquired its name from the assumption that the CS had to leave some trace in the organism's nervous system to be effective.

Successive presentations of just the US at regular intervals (e.g., every half hour) also produce conditional responding; this procedure is called *temporal conditioning*, and responding has sometimes been spoken of as conditioned to time as a stimulus (as the repeated presentation of a stimulus, temporal conditioning is equivalent to some of the stimulus-presentation procedures in Chapter 4).

Reversing the order of the two stimuli is called *backward conditioning*. It is often less effective than other conditioning procedures but it has occasionally been demonstrated, particularly with aversive CSs:

> common sense would lead one to expect animals to have the ability to respond defensively to a novel stimulus detected after a sudden aversive event. An animal that sighted an unfamiliar predator following an abortive attack surely would not submit to another attack (i.e., a forward pairing of the predator and pain) before reacting defensively. (Spetch, Wilkie, & Pinel, 1981, p. 163; cf. species-specific defense reactions in Chapter 8)

Finally, a procedure with two stimuli in which one (the positive stimulus or CS+) becomes a CS through its relation to the US while a second (the negative stimulus or CS-) does not because it never precedes the US is called *differential conditioning*.

In all of these cases, the CS must produce the CR because of its relation to the US and/or UR, and not for other reasons. For example, if a visual stimulus and a traumatic shock occur together, a later startle to the visual stimulus might not mean it had become a CS. A startle response might be elicited by a variety of innocuous stimuli after a traumatic shock, even though these stimuli never occurred near the time of the shock. These are cases of *pseudoconditioning* (Ison & Hoffman, 1983) and *sensitization* (cf. Chapter 4).

Conditioning and Contiguity

Some of the attention historically given to respondent conditioning may have depended on how easily it could be related to the concept of *association*, a principle of learning with substantial precedent in philosophy and psychology. Learning had been said to take place through the association of ideas, and conditional reflexes seemed to represent a primitive example of the formation of such associations. If ideas were associated, it was argued, then one could lead to another. In a kind of mental chemistry, ideas were supposed to become associated through such properties as having common elements or occurring together in time. It remained then to suggest that ideas could be interpreted as responses generated by environmental events, so that remembering one event in the past would call up others it had occurred with.

We needn't dwell on the details. Even those parts of the contemporary psychology of learning that are still associationistic have evolved considerably from earlier formulations. We've already noted that respondent conditioning can't be interpreted as simply making one stimulus a substitute for another. For now, the point is that respondent conditioning was regarded as a process at the root of all learning, and it was assumed to take place merely through the *temporal contiguity* of events, their occurrence together in time. Theoretical debates then revolved around the primacy of respondent conditioning and other processes, and particular attention was given to finding ways of interpreting instrumental or operant behavior as an instance of behavior generated by respondent principles (for various sides of the argument see Guthrie & Horton, 1946; Hull, 1943; Konorski, 1948; Mowrer, 1960; Schlosberg, 1937; Skinner, 1935b; Smith, 1954).

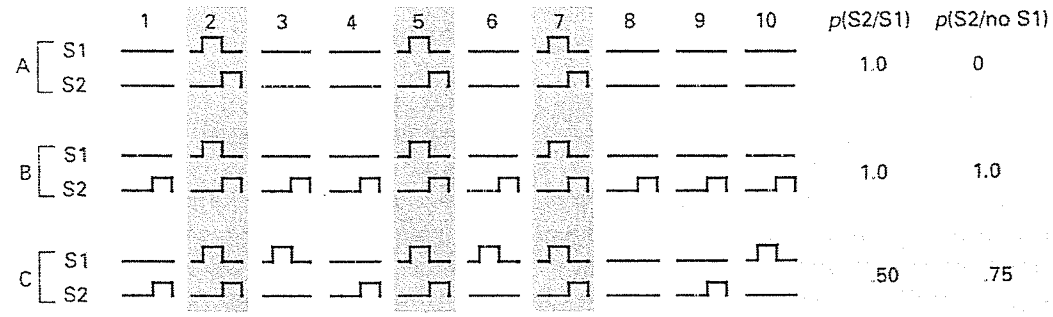

Figure 17-3 Conditional relations between two stimuli, S1 and S2. Rows represent samples of 10 trials from three conditioning procedures. Each involves exactly three pairings of S1 and S2 (in trials 2, 5 and 7), but S1 predicts S2 only in the top procedure. In the middle one, S2 is equally likely given S1 and given no S1, and in the bottom one it is less likely given S1 than given no S1. Probabilities of S2 given S1 and given no S1 are shown at the right and are respectively plotted as A, B and C in Figure 17-4. (cf. Rescorla, 1967)

Part of the problem was that *contiguities* among stimuli were not adequately distinguished from stimulus-stimulus *contingencies*. Contiguity is defined by stimulus pairings, the number of times stimuli occur together. But the number of stimuli can remain constant while contingency relations between CSs and USs vary. For example, assume that buzzer (S1) and food (S2) are arranged within trials and that we can ignore the stimuli that demarcate the trials. The rows in Figure 17-3 show samples of 10 trials from three conditioning procedures. In each, S1 is paired with S2 in trials 2, 5 and 7. In the top row, only those trials include S2, which is therefore perfectly correlated with S1: In other words, the probability of S2 is 1.0 given S1 but zero without it, so S1 perfectly predicts S2. In the middle row, S2 occurs in every trial and so S1 is irrelevant to whether S2 occurs: The probability of S2 is 1.0 with or without S1. In the bottom row, S1 occurs in 6 trials, but in only half of these is it followed by S2, whereas S2 occurs in three-quarters of the trials in which S1 isn't presented: Despite the pairings, the probability of S2 is lower with S1 (.50) than without it (.75).

Figure 17-4 shows these three conditions within a contingency space for stimulus-stimulus relations. Only in the first procedure is *S1* likely to become an effective CS; in the last procedure, *S1* might even reduce the likelihood of conditional responding elicited by the trial stimulus. The conditional relation between the two stimuli (*contingency*) rather than the number of pairings (*contiguity*) is the appropriate basis for classifying conditioning procedures (Rescorla, 1967, 1988).

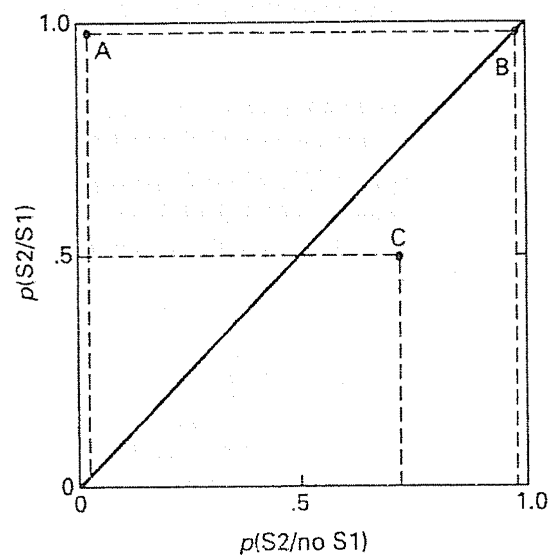

Figure 17-4 A stimulus-stimulus contingency space. The unit square shows conditional probabilities of stimulus S2 given stimulus S1 and given no stimulus S1. The three points, A, B and C, correspond to the three procedures of Figure 17-3. (Cf. Figures 4-2 and 5-9)

Contiguity and Consequences

Instances of conditioning based on aversive stimuli such as electric shock were called *defensive conditioning*, on the assumption that the responses elicited by such stimuli occurred because they had some natural defensive function (cf. Chapter 8 on species-specific defense reactions). One frequently cited example is an experiment with an infant boy named Albert (Watson & Rayner, 1920). The aversive stimulus was the sound, just behind little Albert, of a suspended steel bar struck with a hammer. This sound produced crying or withdrawal or startle responses. When the sound followed presentations of a white rat, these responses began to occur in the presence of the rat and of other stimuli having properties in common with it, such as cotton wool. Watson and Rayner called these responses *conditioned emotional reactions*. Yet the details of the experimental procedure show that the initial clangs of the steel bar weren't independent of behavior:

> 1. White rat suddenly taken from the basket and presented to Albert. He began to reach for the rat with left hand. Just as his hand touched the animal the bar was struck immediately behind his head. The infant jumped violently and fell forward, burying his face in the mattress. He did not cry, however.
> 2. Just as the right hand touched the rat the bar was again struck. Again the infant jumped violently, fell forward and began to whimper. In order not to disturb the child too seriously no further tests were given for one week. (Watson & Rayner, 1920, p. 4)

At least at the outset, therefore, Watson and Rayner's experiment used punishment of reaching toward the rat and not merely response-independent stimulus presentations. This was presumably important in getting Albert to attend to the white rat, but it means that we cannot unambiguously attribute little Albert's responses to conditioning. We can't even rule out the consequences of Albert's responses to the struck bar; after all, these responses led the experimenters to terminate the procedure for a week.

The problem, however, wasn't restricted to Watson and Rayner. Once we look for response consequences in supposed conditioning procedures, we often find them. For example, early demonstrations of limb withdrawal elicited by shock were indifferent to the method of attaching the electrodes. Yet if both electrodes are attached to a dog's leg, a flexion can't prevent shock delivery, whereas with either or both electrodes attached to the floor on which the dog's paws rests, a flexion prevents or terminates shock delivery by breaking the electrical circuit. In fact, the classic defensive reflex of Bechterev (1933) typically had both electrodes on a surface the organism touched, so that a response prevented or terminated the shock; Bechterev was therefore probably studying avoidance and escape behavior rather than respondent conditioning. Recognizing the different implications of the two methods of electrode placement was an important step in the evolution of the distinction between operant and respondent behavior (Kimmel, 1976; Schlosberg, 1937; Skinner, 1935).

Once consequences became implicated in a few cases of presumed respondent conditioning, it was tempting to seek them in all. For example, can't the flexion have consequences even with both electrodes attached to a dog's leg in defensive conditioning? Suppose a buzzer reliably precedes shock. How can we tell whether the shock is as aversive passing through a flexed limb as through an unflexed one? In defensive conditioning, conditional flexions are ordinarily slower and different in magnitude than unconditional flexions. Perhaps this is because a dog whose leg is already flexed doesn't have to adjust its posture as much when shock is delivered as when it stands on all fours and must shift its weight to the remaining three legs once flexion is elicited (Wagner, Thomas, & Norton, 1967). Clearly salivation too has consequences; it affects taste and swallowing in the case of dry food and dilution in the case of acid on the tongue (e.g., (Hebb, 1956)Hebb, 1956).

The place of respondent conditioning in learning theory began with attempts to reduce all instances of operant learning to special cases of respondent conditioning, but the new arguments had turned the situation around. The case was made that all instances of respondent conditioning could be interpreted in terms of consequences that were earlier unnoticed. The next step was to observe that autonomic responses such as salivation and constriction or dilation of blood vessels were often accompanied by somatic responses (e.g., muscle contractions producing skeletal movement). The position could then be taken that autonomic responses in respondent conditioning were artifacts, incidental accompaniments of the behavior generated by instrumental processes (Smith, 1954). The status of respondent conditioning therefore came to depend on demonstrations of conditioning that couldn't be interpreted in terms of response consequences.

One approach was to see whether conditional responding could be modified by explicitly arranged consequences. If such consequences were ineffective, then the argument that the new reflex relation depended on other unidentified consequences would no longer be convincing. Sheffield (1965) therefore added consequences to the conditional salivation generated by the classical Pavlovian situation. Specifically, a tone CS preceded food but food was omitted if the dog salivated on that trial. (The procedure is an example of negative punishment sometimes referred to as *omission training*.) In other words, this arrangement converted the standard Pavlovian procedure to one in which the consequence of salivating was no food and the consequence of not salivating was food.

Salivation wasn't modified by its consequences in this procedure. At the beginning of training, the dog Vicki's conditional salivation hadn't yet begun, so the tone was consistently followed by food. This contingency produced conditional salivation, but once Vicki began salivating in a trial, food was omitted, so conditional salivation decreased on later trials. Once conditional salivation decreased, the tone was again consistently followed by food, so conditional salivation reappeared. Vicki repeated this cycle of conditional salivation, omitted food, decreased salivation, reinstated food, and return to conditional salivation many times over 40 days (800 trials). Although she could have received food on every trial by not salivating during the tone, she didn't learn to do so and received food on only some trials each day.

It would be premature to conclude that this settled the issue. A consequence effective as a reinforcer for one response may not be effective for another (Chapter 6). Food elicits salivation, so it isn't surprising that salivating is ineffectively reinforced by food. The reduction of salivation by the omission of a reinforcer has been demonstrated with a water reinforcer, which doesn't itself elicit salivation (Miller & Carmona, 1967). Salivation, elicited in some circumstances, can be modified by its consequences in others. The issue is no longer that of reducing operant learning to respondent conditioning or vice versa, because too many lines of evidence distinguish between them. For example, operant cases require responses but respondent conditioning can occur without responses, as when stimulus-stimulus contingencies arranged during paralysis by curare affect behavior after recovery from the paralysis (Holland & Ross, 1981; Rizley & Rescorla, 1972; Solomon & Turner, 1962). Instead, the crucial issue in dealing with operant and respondent cases is to be able to tell which is which.

Autoshaping and Automaintenance

The cases of respondent conditioning so far have included both autonomic responses (e.g., salivation) and somatic or skeletal responses (e.g., leg flexion). We've already considered how these two classes contributed to theoretical distinctions between operant and respondent behavior. The demonstration that autonomic responses such as salivation could be modified by their consequences had substantial impact on these theories. That impact was paralleled by the demonstration that somatic or skeletal responses could be affected by respondent procedures. Both demonstrations implied that operant and respondent processes couldn't be distinguished by physiological criteria that defined kinds of responses; the critical difference was instead in the respective response-stimulus and stimulus-stimulus contingencies.

We discussed earlier the ambiguity of experiments on conditional leg flexions to shock. The problem was that procedures couldn't be designed to guarantee that leg flexions would have no consequences. Perhaps partly for this reason, the demonstration of respondent conditioning of another skeletal response, the pigeon's key peck, received special attention in a procedure called *autoshaping* (Brown & Jenkins, 1968). Because the key peck is a common response in studies of consequential responding, it was important to determine the extent to which respondent processes enter into these performances.

Autoshaping originated as a convenient alternative to shaping the key peck by successive approximations (Chapter 9). It begins with a pigeon that eats reliably from the feeder but hasn't yet pecked the key. At intervals the key is lit and a few seconds later the feeder is operated independently of the pigeon's behavior. Thus, the lit key becomes a stimulus that signals food. The food occasions eating, which in the pigeon includes pecking. We therefore might say that food is a US and that pecking food is a UR. After a few presentations of lit key and then feeder, the pigeon begins to face and move toward the key when it lights. Within perhaps fewer than 10 and rarely more than 100 trials, the pigeon pecks the key whenever it is lit. After pecking is generated by autoshaping, the continuation of the procedure is called *automaintenance*. *Autoshaping* and *automaintenance* simply distinguish the changes in behavior leading up to the first peck from the maintained behavior following this peck.

In autoshaping and automaintenance, food deliveries occur independently of behavior. It is therefore difficult to attribute autoshaped pecking to its consequences. Nevertheless, once such pecking begins it is often followed by food. An argument based solely on the observation that no consequences of pecking are obvious may not be persuasive. Autoshaped pecks were therefore studied in omission procedures analogous to Sheffield's experiment with salivation: Food was delivered after the key light only on trials in which the pigeon didn't peck the key (Williams & Williams, 1969). As with salivation, pecks often occurred in a substantial proportion of trials even though they caused the omission of food. The pecking presumably stabilized at a level at which enough trials occurred without pecks (and therefore with food) to maintain pecking in the other trials.

When food is repeatedly presented to a hungry pigeon, pecking becomes a dominant component of its behavior between food presentations (cf. Chapter 4). The pigeon's autoshaped key pecks may therefore be interpreted as behavior generated by repeated food presentations. They occur mainly during the key light and come to be directed toward it so strongly that they strike the key. Once autoshaped pecking begins, it may be maintained indefinitely by repeated presentations of key light and food even though it has no obvious consequences (in fact, relative to keeping its head in or near the feeder, the pigeon may delay its access to food by pecking the key). The production of key pecking in autoshaping has the critical features that define respondent conditioning, so the respondent vocabulary is appropriate. The key light is a CS. It acquires its capacity to elicit a CR, key pecking, through its correlation with food. Food is the US and it elicits a UR, pecking. As in other cases, the contingent relation between key light and food, not their pairing, determines whether autoshaped pecking occurs.

A distinctive feature of autoshaping is the directed nature of the CR: Pecks generated by the key light could occur anywhere in the chamber (in the air, at the wall, around the food hopper); instead, they're directed at the key. The inverse relation also holds: Pigeons tend to move away from stimuli correlated with the absence of food (Wasserman, Franklin, & Hearst, 1974). For example, if a green key precedes food but a red key doesn't, autoshaped pecking occurs during green but the pigeon moves to the opposite side of the chamber during red. The directed nature of autoshaped pecking has been called *sign tracking* (Hearst & Jenkins, 1974). A response must be emitted before it can be reinforced, and it may sometimes be emitted because it was once elicited. Perhaps then the

behavioral relations that occur in autoshaping are prototypes of the processes from which operant behavior evolves (c.f. Chapters 11 and 16 on attention).

Autoshaped key pecks often differ in duration and topography from pecks maintained by operant contingencies (e.g., Keller, 1974; Schwartz & Williams, 1972), and the topography of autoshaped pecks is more likely to be affected by the nature of the US than by contingencies (Jenkins & Moore, 1973). For example, a pigeon's peck at grain is briefer than and differs in form from its drinking. Autoshaped pecks produced by key light and food resemble eating pecks (when autoshaped pecks are based upon food, the pigeon looks as if it is eating the key), whereas those produced by key light and water resemble drinking (when they're based upon water, the pigeon looks as if it is drinking the key). On the other hand, autoshaping can occur even when the US elicits behavior unrelated to pecking. For example, pigeons will come to peck a lighted key if the key light is reliably followed not by food but rather by access to a social area (Peele & Ferster, 1982).

ADDENDUM 17A: STIMULUS COMBINATIONS IN CONDITIONING

Whether a stimulus becomes effective as a CS depends on the context of stimuli within which it appears (Kamin, 1956). An unusual stimulus or a familiar stimulus in an unusual setting is more likely to become effective than a familiar stimulus in a familiar setting. Sometimes the context itself can become effective as a CS. For example, a rat may learn that a CS is followed by a US when it is in the experimental chamber but not when it is in its home chamber. The experimental chamber is, in effect, a stimulus in the presence of which the CS-US contingency operates (as a result, the rat may not respond to a CS presented in its home chamber). In the context of respondent conditioning, such stimuli have been called *occasion-setters*, in the sense that they set the occasions on which stimulus-stimulus contingencies operate (e.g., Rescorla, 1988). Essentially, they are discriminative stimuli. Such stimuli don't elicit responding; they modify the eliciting effects of other stimuli. Occasion-setting is one of many possible stimulus functions that may be produced by stimulus combinations in respondent conditioning. Let's consider examples of a few others.

Overshadowing, Blocking, and Inhibitory Stimuli

Suppose we present loud tone and dim light together followed by some US, such as a shock that elicits leg flexion. After conditional responding develops to these stimuli, we'd probably find tone a much more effective CS than light if we presented each stimulus separately. Attention is relevant to respondent conditioning as well as operant discrimination: We might say that the organism was attending more to tone than to light, or that tone was more salient than light (cf. Rescorla & Wagner, 1972).

The tone-plus-light example assumes that the organism has no history of conditioning with either stimulus of the compound. In such cases, this dominance of one stimulus over another is called *overshadowing*. When the stimuli of a compound don't become equally effective CSs, the more effective stimulus is said to *overshadow* the less effective one. But attention can be affected by history. For example, suppose dim light is a CS by itself before it is presented together with tone. If simultaneous light and tone are then followed by shock until conditional flexions are observed, we might find when presenting the stimuli separately that light but not tone is an effective CS, even though tone preceded the shock as reliably as did light. When a stimulus fails to become effective as a CS because it is presented with some other already effective stimulus, we say that the stimulus with the prior history *blocked* conditioning to the new stimulus. The organism continued to attend only to the original CS. In such cases, the effect is called *blocking*.

In overshadowing, neither of two stimuli of a compound (S1 and S2) has a prior conditioning history but only one becomes effective as a CS (or one becomes more effective than the other). In blocking, one stimulus is made effective as a CS, and this history prevents the other from becoming effective when the two are presented together as a compound (or the first reduces the effectiveness of the second).

Stimuli that signal the omissions of other stimuli (cf. C in Figure 17-4) sometimes acquire the capacity to reduce the effectiveness of other CSs; they are then described as *inhibitory*. Here is an example involving food-elicited salivation in a dog. First, a buzzer is made a CS. Once the buzzer reliably elicits salivation, a tone is presented either alone or with a light on irregularly alternating trials. When tone is presented alone, it is followed by food. When it is presented with the light, food is omitted. Eventually tone alone elicits conditional salivation but the tone and light together do not. Later, when the buzzer is presented with the light the buzzer elicits less conditional salivation than it would have if presented by itself, and the light is said to *inhibit* the effect of the buzzer. In this instance, the dog is doing more than simply discriminating tone alone from the tone-plus-light combination. Note that this usage of the language of inhibition is different from its usage in cases of schedule interaction (e.g., Chapter 16).

Sensory Preconditioning and Second-Order Conditioning

We've considered respondent conditioning based on USs that serve in other situations as reinforcers (e.g., food) or as punishers (e.g., shock). But does conditioning occur when relatively neutral stimuli such as lights or sounds serve as USs? The problem is that conditioning is difficult to assess in the absence of elicited responses. Two procedures concerned with effects of such stimuli, *sensory preconditioning* and *second-order conditioning*, are illustrated for conditional leg flexions in dogs in Figure 17-5.

First, consider sensory preconditioning (Brogden, 1939). In the first phase, preconditioning, one stimulus signals a second stimulus. In this example, a buzzer is consistently followed by tone. In the second phase, a conditional reflex is created in which the second stimulus becomes a CS. In the figure, tone is followed by shock. Once the conditional reflex is created, so that tone elicits leg flexion, the eliciting effects of the buzzer are tested. Leg flexion to the buzzer is then taken to mean that the buzzer became a CS relative to the tone during preconditioning. (For convenience, control groups used to counterbalance stimuli and to rule out sensitization and other effects have been omitted; cf. Chapter 4.)

In second-order conditioning, the order of phases is reversed (cf. Rescorla, 1980). In the example in Figure 17-5, a conditional reflex in which tone elicits leg flexion is created first, by presenting tone and then shock. Later, a buzzer is followed by the tone. In this instance, the question is whether the CS created in the first phase can function as a US for another stimulus in the second phase. The difficulty is that the tone loses its effectiveness as a CS as it is repeatedly presented without shock in the second phase, but during just this time the buzzer must acquire its conditional properties. An alternative procedure, with the buzzer followed by the tone and then shock on all trials within a single phase, would be ambiguous; we wouldn't know whether leg flexions elicited by the buzzer occurred because of the relation between buzzer and tone or the relation between buzzer and shock.

If the CS is presented alone after sensory preconditioning, thereby extinguishing the conditional reflex, the preconditioning stimulus also will no longer elicit a CR. In other words, in the example of Figure 17-5, presenting tone alone after the conditioning phase until the tone no longer elicits flexions will also make the buzzer lose its effectiveness as a CS. But the comparable procedure after second-order conditioning doesn't always extinguish the second-order conditional reflex (cf. Holland & Ross, 1981; Rizley & Rescorla, 1972). In the example of Figure 17-5, presenting tone alone until it no longer elicits flexions may not eliminate the conditional flexions to the buzzer created during the second-order phase.

This outcome is paradoxical. Consider an analogous human case history. A young man sees blood in painful circumstances and the sight of blood becomes a CS that elicits those emotional responses we call fear. At a later time (equivalent to the second-order phase), he sees blood in an elevator and thereby acquires a fear of elevators. Then he goes into hospital work and in that context gradually gets over his fear of the sight of blood. According to Rizley and Rescorla's findings, this change won't reduce his fear of elevators even though the CR to blood had been the basis for the fear. This case is hypothetical and we must generalize cautiously from experimental to real-life situations. But the analysis of such respondent contingencies is relevant to behavior therapies that are assumed to eliminate fears or phobias by extinguishing acquired responses to conditional aversive stimuli. Examples include *flooding*, which involves exposure to intense versions of the aversive stimulus, usually for an extended time, and *systematic desensitization*, which involves the gradual fading in of the aversive stimulus (Wolpe, 1958, 1990).

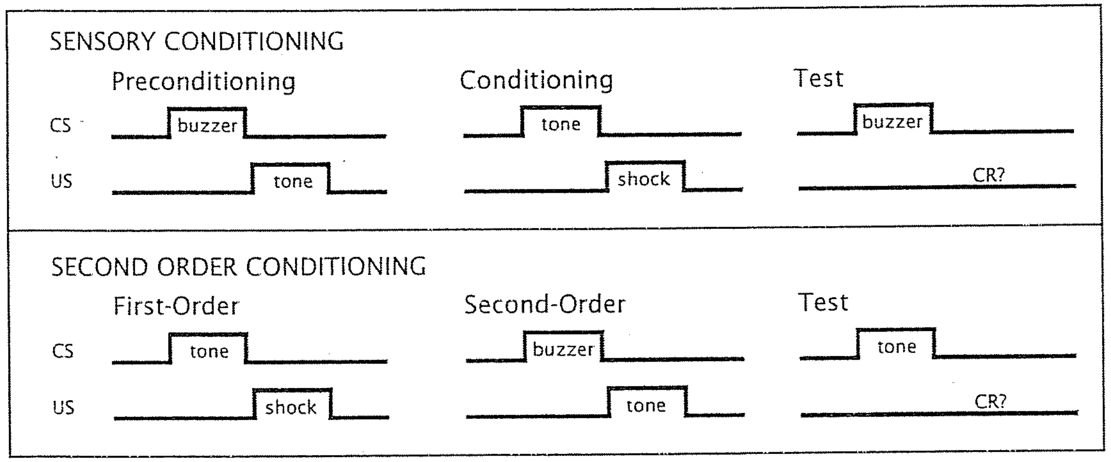

Figure 17-5 A schematic diagram of the phases of sensory preconditioning and second-order conditioning procedures, using buzzer and tone as CSs and shock-elicited leg flexions in a dog as the UR.

Chapter 18
Operant-Respondent Interactions: Emotion

The language of emotion includes several etymological clusters. For example, eager; anger *and* anxiety *have common roots, as do* choleric, melancholy, glad *and* glee. *Another cluster includes* wrath *and* worry, *and still another* sad *and* satisfy *(the last pair is also related to* satiate*).* Fear, *from the Indo-European* per-, *to try, to risk, to press forward or lead, has an extensive group of relatives that includes* experiment *and* apparatus, probability *and* opportunity, prepare *and* repertory *(*repertoire*),* approach *and* deprivation, *and two contemporary synonyms for behavior,* comportment *and* performance.

- Conditioning and Emotion
- The Language of Emotion
- Preaversive and Preappetitive Stimuli
- Recapitulation
  - *Kinds of Contingencies and Contingent Stimuli*
- Addendum A: Biological Constraints on Learning
  - *Sensory Constraints*
  - *Motor Constraints*
  - *Constraints on Consequences*
  - *Preparedness*

Some birds develop pecking orders, rankings in which each bird will yield to attacks by birds above it in the ranking but can attack with impunity any bird below it. Skinner once designed a demonstration apparatus in which to create and reverse pecking orders in pairs of pigeons (Reynolds, Catania, & Skinner, 1963; Reynolds & Skinner, 1962; Skinner, 1959a). A feeder could be operated either by a vertical bar on one side or by a treadle on the other. One pigeon had been trained to use the bar when a lamp was lit one color and the other to use the treadle when the lamp was lit another color, so one or the other color could be used to reinforce the behavior of just one of the pigeons, even though they both occupied the same chamber. Given a blue ceiling light, one pigeon's pecks at the other were reinforced with food, and given a green ceiling light these contingencies were reversed, so that the pigeon that had been the attacker became the target. A white ceiling light signaled an intermission during which neither pigeon's pecks were reinforced.

Over time under these conditions, the pecking of the attacker pigeon often changed from a topography that resembled pecks on a standard pigeon key to one that included some of the typical properties of aggressive behavior, such as a fluffing up of the breast feathers, while the target bird tucked its wings down and crouched defensively. When the ceiling light changed from green to blue or vice versa, the attacker bird quickly collapsed its feathers and assumed the defensive posture while the other bird turned and puffed itself up and became the aggressor. When the white ceiling light came on the aggressive behavior ended and both pigeons effectively moved to their neutral corners.

KEY TERMS: Operant-Respondent Interaction, Emotion; Preappetitive and Preaversive Stimuli; Anxiety, Conditioned Suppression; Constraints on Learning, Preparedness.

We might think of aggression and submission as patterns of behavior that include emotional components, but here those patterns were quickly turned on or off or were reversed simply by changing the discriminative stimuli correlated with a reinforcement contingency. We may think of emotional behavior as elicited by various events: "I couldn't help it," shouts the red-faced parent, "What you did made me angry." But if this is how we think it works, the outcome of this demonstration should be counterintuitive. It seems to imply that some features of this behavior are operant rather than respondent. On the other hand, have you never seen or perhaps participated yourself in a shouting match that is unexpectedly intruded upon by a respected third party? What seemed a battle royal abruptly changes to sweetness and light once this person comes on the scene. This is a case where an environmental event turns off the emotions as surely as does the white ceiling light in the pigeon apparatus. It makes sense, then, to look carefully at the operant as well as the respondent features of the behavior we call emotional.

When we get to verbal behavior, especially in Chapter 22, we'll discuss some of the origins of the language of emotions, but some comments are appropriate now. We tend to speak of emotions as something inside us, ready to come out under certain circumstances and sometimes even hard to keep in. But we cannot find any key pecks inside the pigeon when it is not pecking and we cannot find any reaches and touches inside ourselves when we are not reaching and touching, so maybe we should start thinking about emotional behavior in the same way.

Darwin got a lot of other things right, but maybe he erred in his book, *The Expression of the Emotions in Man and Animals* (Darwin, 1872). To say that we express our emotions is certainly to imply that we are showing something that would otherwise remain unseen. But what if emotions have evolved not as inner conditions but rather because of the effects they have on others? An angry male gorilla surely influences the behavior of those around him; his gestures and sounds are exceedingly effective expressions of his emotions. If confronted by him we would probably worry more about his public behavior than about anything inside him.

It has been argued that a function that drove language evolution was that it allowed us to express our emotions. But evolutionary contingencies don't ordinarily create systems that duplicate functions already in place, and the angry gorilla tells us that nonverbal primates can already express their emotions pretty well. Why then should they need language as a substitute? It is unlikely that language evolved for that reason. On the other hand, if language is a way to change the behavior of others, a consequence of describing emotions might sometimes be to get treated differently.

Our language of emotion has other problems embedded in it. If I see a man red-faced and shouting and gesticulating, I may well call him angry. But if you come to me and ask why he is behaving that way, I hope you'd object if I replied that he is behaving that way because he is angry. Yet that is what we often do with the language of emotions. We've learned to label various patterns of behavior, as when we say that people are happy or sad or angry or in love. But too often when we're asked why they do what they do we turn things around and offer the emotions as the explanations for their behavior. This sort of circular reasoning is pervasive. It even sneaks into technical usages, as when certain patterns of social behavior are summarized as attitudes and later the attitudes serve as explanations when those social patterns are observed elsewhere.

Another problem with the language of emotion is that it cannot easily be converted from nouns to verbs. We can speak of attending instead of attention or of thinking instead of thoughts or of remembering instead of memory. In each of these cases, the verbs make us less likely to treat these categories as things rather than as behavior. We often say that anger is directed, as when we are angry at someone or something, but we don't say that someone angered at someone else. The language of emotion is highly variable (we'll consider some reasons why this is so in Chapter 22). *Love* is sometimes a noun and sometimes a verb (we can love or we can be in it), but *hate*, its counterpart, functions differently. We speak of joy and fear and sadness more as possessions than as actions. Yet if these are classes of behavior that have antecedents and consequences we are at a disadvantage if we lack suitable verbs (cf. the discussion of hunger in Chapter 14).

Given these concerns, we must proceed cautiously, as usual favoring procedures that show us how behavior works over accounts in terms of theoretical entities. Operant and respondent processes can interact when respondent procedures are combined with operant procedures. For example, a stimulus that reliably precedes or signals shock may not just elicit leg flexions; it may also interfere with behavior that has been maintained by its consequences, such as lever pressing maintained by food reinforcement. We sometimes describe comparable behavior in humans in terms of fear or anxiety; thus, we should find procedures such as these relevant to the experimental analysis of emotion.

CONDITIONING AND EMOTION

Stimuli that signal the presentation of other stimuli can be superimposed on baselines of ongoing operant behavior. For example, suppose a rat's lever presses are maintained by food reinforcement; from time to time a tone is presented and the tone ends with the delivery of a shock. In such circumstances, the tone typically reduces lever pressing, especially as the time of shock delivery approaches. This phenomenon, originally demonstrated by Estes and Skinner (1941), has been variously called *anxiety*, *conditioned suppression* and *conditioned emotional response* or *CER*. It is illustrated in Figure 18-1, which shows the development of and recovery from suppression (Geller, 1960). The rat's lever presses were maintained by a VI 2-min schedule of food reinforcement; a brief shock followed 3-min presentations of tone. After the tone came to suppress responding, shock was discontinued and pressing during the tone recovered to earlier levels.

The procedure is an instance of respondent conditioning: One stimulus, tone, signals another stimulus, shock. The shock is presumed to be aversive, so we call the tone a *preaversive* stimulus. This is another case in which behavior produced by the CS differs from that produced by the US. The tone suppresses reinforced lever pressing, but pressing begins again promptly after the shock.

If we looked more closely at the rat's behavior during the tone, we'd find that these contingencies affected many other classes of responses besides lever presses, such as heart rate and respiration (cf. Blackman, 1977; Rescorla & Solomon, 1967). We are most likely to invoke the language of emotion when an event affects a broad array of different response classes, so we may be tempted to speak of the rat's fear or anxiety. If we do so, we must recognize that for the reasons we have discussed above such terms don't explain the rat's behavior. It wouldn't do to say later that the rat stopped pressing during the tone because it was afraid; the effect of the tone on the rat's pressing led us to speak in terms of the rat's fear in the first place.

Our language of emotions is complicated. We speak of emotions in others and in ourselves on the basis of both situations and the behavior that occurs in those situations (cf. Ortony & Turner, 1990). For example, we might speak of the behavior produced by preaversive stimuli in terms of fear or anxiety, but if we also observed aggressive behavior we'd be more likely to speak of anger. Either way, we must be clear that these names for behavioral effects don't explain them. Consider, for example, the question of why someone is acting in some way. We might answer that the person is sad or depressed. Once again we must beware of circular reasoning. It would be more useful to know that the person is acting that way because of some specific event, such as the loss of a job or the breakup of a love affair.

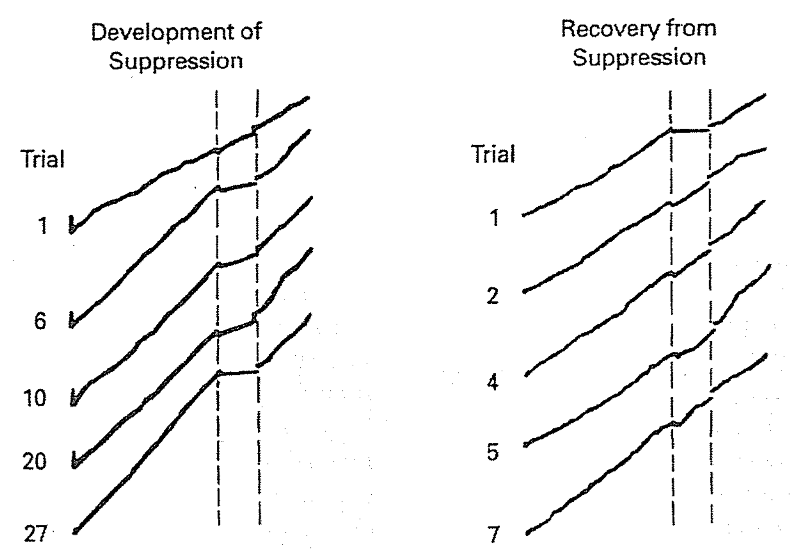

Figure 18-1 Development of and recovery from suppression during a preaversive stimulus. A rat's food-reinforced lever presses were maintained by a VI 2-min schedule (the cumulative records do not show reinforcers). Left: the effects of superimposing a 3-min tone followed by electric shock on baseline lever pressing. The downward displacements of the record between the dashed vertical lines mark off periods of tone. By trial 27, pressing was almost completely suppressed by the tone. Right: recovery from suppression when the tone was no longer followed by shock. (Adapted from Geller, 1960, Figure 3)

A practical application of our understanding of preaversive stimuli is provided by the treatment of children under intensive care, as in hospital burn units (Derrickson, Neef, & Cataldo, 1993). These children are subjected to unpredictable and uncontrollable aversive events at any time of the day or night: injections, changes of dressings, intravenous feedings and so on. A typical outcome of their long-term care is that they become lethargic and withdrawn; they don't react to the events around them (cf. learned helplessness, Chapter 12). For these children, the hospital setting has become one massive preaversive stimulus. Some anxiety disorders may be related at least in part to aversive events that are uncontrollable rather than controllable (Mineka & Zinbarg, 2006).

We can help by giving the children at least some control over part of the environment some of the time (e.g., in social interactions with visitors or ward staff, in choices of meals whenever possible, and so on), but constraints on care delivery may limit the feasibility of this approach. Another alternative is suggested by our analysis in terms of preaversive stimuli. If we turn on a red light over a child's bed at least 10 minutes before we start any aversive procedure, the red light becomes a preaversive stimulus. Its onset will become aversive too, but while it does so its absence also becomes a *safety signal*, a time when the child is safe from aversive medical procedures. The safety may be relative, in that the child may still be in pain some of the time and emergencies may sometimes not allow time for use of the red light, but relative safety is better than none at all. During the safe times, while the red light is absent, the child's physiological reactions to the conditions that signal aversive events relax, and behavior with reinforcing events as a consequence will be more likely to be maintained; this may speed the child's recovery as well as reduce the child's lethargy and withdrawal.

THE LANGUAGE OF EMOTION

For many years, my behavior was affected without my knowing it by a sensitivity to a substance called *tyramine*. Tyramine is a derivative of fermentation found in a variety of foods and drinks, including beer, aged cheeses, processed meats, soy sauce and red wines, along with coffee and colas and chocolate and most other sources of caffeine, to mention some of the major culprits though not necessarily in their order of severity in producing symptoms (the tyramine content of some foods is highly variable). Many who are sensitive to tyramine get migraine headaches as their primary symptom, but mine were different and except for a facial rash were mainly behavioral. Although I didn't know it at the time, too much tyramine made me somewhat paranoid, gave me bad dreams and lowered my anger threshold. These symptoms were no doubt harder on my family than on me. I won't go into details, but it is fair to say that I was sometimes hard to get along with.

I had been somewhat short-tempered and had suffered occasionally from facial rashes early in my academic career. I'm fairly sure that my father had a similar problem. In the spring of 1980 I had a sabbatical semester and spent three weeks visiting my colleague, Terje Sagvolden, at the University of Oslo in Norway, and then attended a professional meeting in London. In Norway my symptoms vanished, but rather than attributing that to a different diet I concluded that I had escaped from the stress of work and family life. When I moved on to London, however, where a pub lunch was likely to include aged cheeses and a pint, my symptoms returned.

By the time I was back home I was thinking about what I was eating and drinking. I had often attended an afternoon faculty reception in which wine and cheese were served, so on some occasions I abstained from the wine and on others I abstained from the cheese. Neither made much difference; my symptoms persisted. It hadn't occurred to me that I should also try abstaining from both.

Both of my two sons had occasionally been seen for possible allergies, and after little success with other lines of inquiry I consulted their pediatrician. The role of tyramine in migraines was a fairly new development at the time, and he gave me a list of foods high in tyramine and suggested that I should see what happened if I removed them from my diet. Along the way he also pointed out that allergies, which involve the immune system, should not be confused with food sensitivities. He dealt with the distinction with the same care I would have taken if I were explaining the difference between operant and respondent behavior to one of my students.

I went cold turkey in removing everything on the tyramine list from my diet and within a short time my symptoms disappeared, but they were soon replaced by agitation and headaches and deep muscle pains in my legs. By the end of week my wife pointed out that I had removed not only my sources of tyramine but also all of my sources of caffeine. She made a cup of coffee and within seconds of my first sip all of my new symptoms disappeared. I had been suffering from caffeine withdrawal. Caffeine is addictive but does not share many of the problems of narcotic substances. Many of us do not go without caffeine long enough to encounter its withdrawal symptoms. It has since been recognized that some after-effects of surgery are not produced by the surgery itself but rather by the caffeine withdrawal produced by restricted presurgery diets (Fennelly, Galletly, & Purdie, 1991; Hampl *et al.*, 1995). I now mostly take my caffeine fixes in the form of tea, which contains less tyramine than coffee or colas. My sensitivity to tyramine seems also to have diminished with age. I must concede that I do still get angry from time to time, but it is a far different sort of behavior than once it was.

The relevance of all this to emotional behavior is of course that this behavior, like other responses to environmental events, has multiple causes. One way in which chemical substances such as those in foods can influence it is by changing the thresholds at which stimuli may trigger it. Clearly my anger threshold was lower given tyramine than otherwise. In coping with behavioral problems, we often look to particular environmental stimuli that might be functioning as antecedents or as consequences, but we should not neglect variables that might have their effects only by changing the thresholds for other behavior.

Learning about tyramine made a difference. On some occasions my symptoms reappeared and on checking my recent diet I discovered that I had unwittingly consumed something that contained tyramine. More important, from then on when I found myself on the verge of shouting about something I could ask myself whether it was because of something my wife or one of my sons had done or because of something I'd eaten. That was often good enough, and it also illustrates the potency of verbal behavior (see Chapter 23). The consequences of asking about the contingencies that maintain or shape emotional behavior differ from those of asserting that people act as they do simply because their emotions make them do it.

PREAVERSIVE AND PREAPPETITIVE STIMULI

The language of emotion is important in our interactions with other people, but we each learn it in idiosyncratic ways, so it is inevitably variable and inconsistent (cf. Chapter 22). For that reason it has sometimes stood in the way of behavioral analyses of the emotional effects of stimuli, such as those of preaversive stimuli. The interactions between respondent conditioning and operant behavior, as when preaversive or preappetitive stimuli are superimposed on reinforced responding, can sometimes be seen more clearly in the context of experimental parameters, such as baseline reinforcement schedule, baseline response rate and so on.

The finding that positively reinforced responding can be suppressed by preaversive stimuli, as in Figure 18-1, was later supplemented by the finding that avoidance, which is negatively reinforced responding (Chapter 8), can be enhanced by such stimuli (cf. Sidman, Herrnstein, & Conrad, 1957). In other words, a rat whose lever presses avoid shock may increase rather than decrease its pressing during a stimulus preceding an unavoidable or inevitable shock. This enhanced responding has been called *conditional facilitation* or *conditional acceleration*. Once such enhanced responding develops during negatively reinforced responding, it may continue with positively reinforced responding. For example, rhesus monkeys' lever pressing maintained by orange juice as a reinforcer was originally suppressed during a clicking noise that preceded shock, but after they acquired a history of pressing that avoided shock and were returned to the initial procedure, their pressing during the clicking noise was enhanced rather than suppressed (Herrnstein & Sidman, 1958). This is another case where a history with controllable rather than uncontrollable events makes a big difference.

Preappetitive as well as preaversive stimuli can be superimposed on baseline operant behavior. For example, a key light that precedes response-independent food deliveries increases pigeons' key pecking when it is superimposed on pecking maintained by DRL reinforcement (Herrnstein & Morse, 1957). By analogy to the labeling of suppression during preaversive stimuli as *anxiety*, it was tempting to speak of such enhancing effects of preappetitive stimuli in terms of *joy*. The enhancement of positively reinforced responding and suppression of negatively reinforced responding by preappetitive stimuli seemed to parallel the suppression of positively reinforced responding and enhancement of negatively reinforced responding by preaversive stimuli (e.g., Azrin & Hake, 1969; Leitenberg, 1966). But parameters can make a difference (Blackman, 1977).

As an example, Figure 18-2 shows that shock level and baseline response rate jointly determine whether preaversive stimuli suppress or enhance a rat's lever pressing reinforced with food (Blackman, 1968). During red light and noise, lever presses were reinforced according to a DRL 15-s schedule with a limited hold of 5 s (i.e., a press was reinforced only if emitted within 15 to 20 s of the last press);

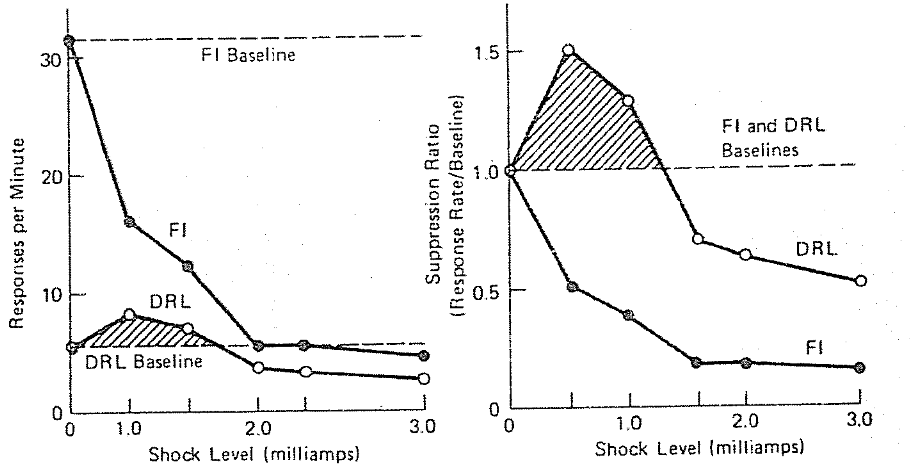

Figure 18-2 Response rates (left) and suppression ratios (right) during a stimulus that preceded shock. Effects of the preaversive stimulus depended jointly on shock level and the schedule that maintained responding. Data are from a rat's lever pressing maintained by multiple DRL FI schedules of food reinforcement. With FI responding, the preaversive stimulus suppressed responding at all shock levels; with DRL responding, it enhanced responding at low levels and suppressed it only at higher levels. (Adapted from Blackman, 1968, Figure 2)

during white light and no noise, an FI 20-s schedule operated with a 5-s limited hold. In these multiple DRL FI schedules, DRL components maintained lower response rates than did FI components. Later, occasional 1-min tones that preceded brief shock were added and shock level was varied to determine the relation between shock magnitude and degree of suppression. Figure 18-2 (left) shows response rate during tone (the preaversive stimulus) as a function of shock level. In the FI component, response rate consistently decreased with increasing shock level. In the DRL component, however, response rate increased at low shock levels and decreased only at higher levels. Figure 18-2 (right) shows the same data converted to a suppression ratio, the change in response rate expressed relative to baseline response rate.

In other words, **you can't say whether one of these procedures will produce enhanced or suppressed responding unless you know what was maintaining the behavior.** The different FI and DRL performances show how behavioral effects can vary as a function of baseline conditions. Superimposing preaversive stimuli on reinforced responding is a respondent procedure that has opposite effects depending on the baseline performance upon which it's superimposed. Similar types of interactions are critical to psychopharmacology, the analysis of the effects of drugs on behavior; there again, the effect of a drug may vary considerably depending on baseline performance.

The schedule of stimulus presentation can also modify the effects of preaversive stimuli. For example, suppression varies jointly with the duration of a preaversive stimulus and its frequency of presentation. It also depends in part on how much the reduced response rate affects reinforcement rate: Less suppression occurs when reduced responding substantially reduces reinforcer deliveries than when it only slightly affects them (Smith, 1974).

As with other respondent cases, **the effects of preaversive and preappetitive stimuli on operant behavior depend on stimulus-stimulus contingencies rather than stimulus-stimulus pairings.** The point is illustrated in Figure 18-3, which shows how various combinations of shock probabilities in the presence and absence of preaversive stimuli suppress a rat's reinforced lever pressing (Rescorla, 1968). For example, if 40% of the preaversive stimuli are followed by shock (shock probability given CS = .40), a range of effects from complete suppression to no suppression at all can be obtained depending on shock probability when the preaversive stimulus is absent. (We can speak of the different contingencies in terms of *predictive value*: The preaversive stimulus is said to have predictive value when shock probability in its presence differs from that in its absence; it is said to have no predictive value when those probabilities are equal; cf. Figure 17-4.)

Responding in the presence of a signaling stimulus is affected by its relation to the stimulus it signals. In the classical Pavlovian case, it seemed at first as though one stimulus substituted for the other, but with preaversive and preappetitive stimuli we've again seen that **respondent phenomena cannot be treated as stimulus substitution.** As usual, it is appropriate to recall that stimuli have multiple functions, and it was probably inevitable that the combination of operant and respondent procedures would raise such issues (cf. Hoffman & Fleshler, 1962).

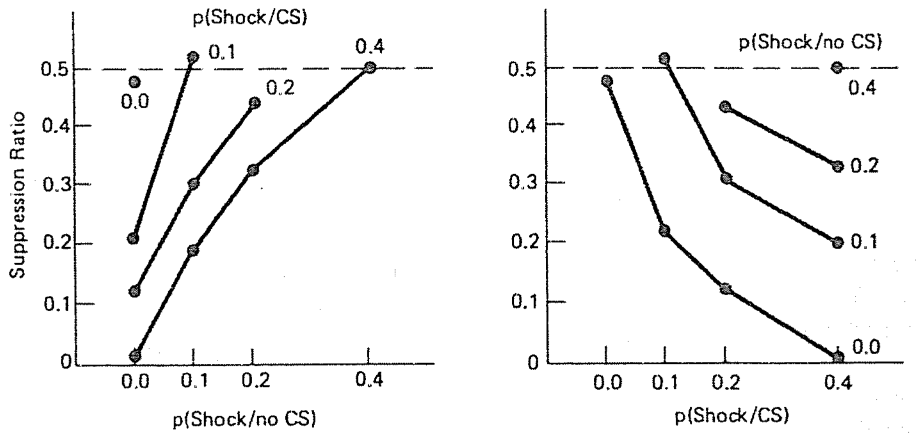

Figure 18-3 Suppression of a rat's lever pressing during a preaversive stimulus (CS) as a function of different shock probabilities during its presence and absence. For example, with a shock probability during the preaversive stimulus of .4, or *p*(SHOCK/CS) = .4, suppression depended on shock probability in its absence, or *p*(SHOCK/no CS). Effects ranged from complete suppression when the latter probability was zero to none when it equalled *p*(SHOCK/CS). The same data are plotted in both halves of the figure: on the left the parameter is *p*(SHOCK/CS) and on the right it is *p*(SHOCK/no CS). In this suppression ratio, baseline equals .5. (Adapted from Rescorla, 1968, Figure 3)

RECAPITULATION

In this text so far, we've considered various behavioral procedures and findings: phenomena of elicitation, reinforcement, discrimination and conditioning, among others. Before we apply these concepts to more complex types of behavior such as language, a review seems appropriate. We'll do so in the context of examples of social behavior involving parents and a child.

Beyond mere observation, the simplest of our procedures was stimulus presentation (Chapter 4): No signal precedes the stimulus and no response needs to occur before it is presented. Suppose a nursing mother starts out by feeding her newborn infant independently of his behavior (let's assume the child is male, not for sexist reasons but rather so we can easily distinguish the mother and the infant as *she* and *he*). Her presentation of the nipple is an instance of stimulus presentation, and it may affect the infant's behavior. He is likely to turn toward her breast and begin to suckle. From the nursing mother's point of view, the suckling is also a stimulus, and it elicits the letting down of her milk. This glandular response moves the milk toward her nipples, where it becomes available to the suckling infant.

Assume now that the mother switches to demand feeding, and so feeds the infant only when he begins to cry. The feeding becomes a consequence of a response, crying. This relation is, of course, a response-stimulus contingency; we translate it as the effect of a response on the probability of a stimulus. In this instance, the infant isn't fed unless he cries; without the crying the probability of a feeding is zero. This is, of course, just one example of a contingency. Responses can raise or lower stimulus probability (Chapters 5-8); they can turn things off as well as turn things on (they can also change other contingencies, but our examples here don't require such levels of complexity).

The contingency between crying and feeding is likely to affect the infant's behavior. We can expect an increase in crying, but only after some time has passed since his last feeding, when milk has again become reinforcing through deprivation as an establishing operation.

Let's now add stimulus control: We superimpose a discriminative stimulus on either of the other operations. First consider the mother's behavior. When she feeds the infant, the suckling produces the letting down of milk into her breasts. Once she begins feeding the infant whenever he cries, the crying becomes a reliable antecedent of the feeding and the mother discovers that she begins to let down her milk as soon as he begins to cry. The stimulus, crying, is followed by another stimulus, the infant's suckling at her breast. This relation should be familiar as an example of Pavlov's respondent conditioning (Chapter 17). When Pavlov presented dogs with light and then food, the dogs began to salivate after light as well as after food, just as the mother lets down her milk during the crying as well as during the feeding.

But a discriminative stimulus can also be superimposed upon the consequential operations of reinforcement or punishment (Chapters 11-12). Now the infant has grown a bit and sleeps through the night. The mother begins the practice of feeding him when he begins to cry during the day but not when he does so at night. These times are correlated with daylight and darkness, and soon the infant begins to discriminate between them. Nighttime crying decreases relative to daytime crying, and later subtler discriminations allow the mother to begin to shape other kinds of behavior to replace the crying. Until then, during daylight the infant is fed when he cries but when it is dark he isn't; in other words, light and dark become discriminative stimuli. During the light, crying raises the likelihood of a feeding, but during the dark it doesn't.

When we examine behavioral situations, it's often a useful exercise to describe them in terms of the appropriate operations; sometimes a proper description requires a combination of operations. For example, suppose that the mother has learned that when her infant begins to make fussy noises near bedtime, he's likely to fall asleep quickly if she picks him up and rocks him. His fussy noises set the occasion for rocking him and his subsequent falling asleep sets the occasion for tucking him into his crib. With regard to the mother's behavior, both parts of this sequence involve stimulus-control operations superimposed on consequential operations. The rocking produces a consequence, the sleeping infant, that is in turn a discriminative stimulus for the tucking in; a stimulus that's contingent on responding in one part of the situation serves as a discriminative stimulus in another (cf. chains in Chapter 10).

These examples have involved interactions between the infant's behavior and the mother's. In the context of such interactions, the infant learns to give special attention to social stimuli. These later become significant in many kinds of social behavior, such as attending to what others say, taking turns in conversation and saying things that affect the behavior of others. Maybe you've noticed that many of the human examples we've used to illustrate basic processes have involved children with little or no verbal behavior, such as young infants or children with autism or other developmental disabilities. This was no accident, because the basic contingencies are most effective when they're uncontaminated by talk. As we'll see later, verbal behavior is a very special kind of social behavior.

Kinds of Contingencies and Contingent Stimuli

Let's call the stimuli involved in response-stimulus relations *contingent stimuli*. The preceding account used the mother's milk and the infant's suckling as examples of contingent stimuli. But stimuli are of various sorts. We could replace those of the preceding examples with others that are aversive. The infant might be less fortunate, and the shouting of an abusive parent might occur independently of the child's behavior; or it might occur only after some response, such as crying; or it might occur only in the presence of some other stimulus, as when the child learns that his father shouts only when his mother is around; or it might be that if he cries his mother always quiets his shouting father, so that her presence is an occasion when he can avoid or escape from his father's shouting by crying. These are all examples of contingencies involving aversive stimuli.

It's sometimes convenient to distinguish among different types of contingent stimuli. Organisms work to produce or remain in the presence of some, called *appetitive*, *rewarding* or *reinforcing*: food, entertaining company, money, among many others. They work to remove or stay away from others, sometimes called *aversive*, *noxious* or *punishing*: noise, dull company, extremes of heat or cold, among many others. Some of these stimuli have obvious biological significance while others acquire their significance during the organism's lifetime. It is difficult to classify any stimulus unambiguously as a reinforcer or a punisher, especially given the relativity of reinforcement and punishment (Chapters 5-8). And, having admitted appetitive and aversive stimuli as classes of contingent stimuli, we must also recognize that relatively neutral or insignificant stimuli can enter into contingencies. If the infant reaches out and touches his mother, for example, his movement is a response and contact is its consequence. Obviously, no stimulus is likely to be completely without significance; these classes represent points or regions in a continuous range of stimulus types rather than three discrete categories, and the designation of particular stimuli can change as a result of establishing or motivational operations.

A response may lower as well as raise the probability with which events occur. Suppose that the father finds the infant's cries aversive. If the infant is especially likely to start crying if his diaper hasn't been changed for some time, then the father can avoid the infant's crying by changing the diaper. The probability of crying increases as time passes without a diaper change, and decreases whenever a diaper change occurs. In other words, the diaper change is the father's avoidance response. A response-stimulus contingency is defined as the effect of a response on the probability of a stimulus, and we distinguish among contingencies in terms of whether the effect is an increase or a decrease in probability. In this last case, in which the diaper change avoided the infant's crying, the stimulus was aversive and its probability was reduced by the father's response.

Now suppose that the infant has fallen asleep and while watching television the father accidentally awakens him by turning up the sound too loud. The infant begins to cry. The contingent stimulus, the crying, is again aversive, but this time a response, turning up the sound on the television, has raised its probability. The crying may punish the turning up of the sound, in that the father may from now on be less likely to do so while the infant is sleeping. Changes in probability are, of course, not limited to all-or-none cases. For example, the infant might not be fed every time he cries, and he might not awaken and cry every time someone turns up the television sound (Chapters 15–16 on schedules).

In classifying contingent stimuli, we recognized relatively neutral or insignificant stimuli as well as appetitive or aversive stimuli. Contingencies also range from those in which responses raise stimulus probability to those in which they lower it, and within this range is the special case in which a response has no effect on stimulus probability. This special case is equivalent to the stimulus-presentation operation. Stimulus presentations may involve the response-independent delivery of appetitive stimuli or aversive stimuli or relatively neutral or insignificant stimuli.

Kinds of contingencies and contingent stimuli are summarized in Table 18-1, in which they are shown combined with a discriminative stimulus and with some names that have been applied to them. The procedures aren't exhaustive, so the list is incomplete. For any procedure it may be instructive to locate it or its various stages among the classes shown in Table 18-1.

Consider an example. The infant is older, and his mother now allows him to crawl around and explore some of the rooms in his home. After he's done this a few times over several days, the mother takes him to a room in the corner of which is a new toy chest. She opens the chest and gives him a favorite toy. The next day she puts him down in another room and he immediately sets out for the room with the toy chest, getting there quickly and without making any wrong turns. His exploration of the room on the previous days involved behavior that produced relatively neutral consequences. But when something more significant was introduced, the new chest with some of his toys in it, the child proved that he had learned the layout of the rooms. The example is analogous to an experiment on latent learning.

Consider one more example. The mother has often taken the infant to the pediatrician's office for routine exams. Time in the waiting room has reliably been followed by seeing the pediatrician in the examining room. One evening the infant becomes ill and the mother takes him to a hospital urgent care unit rather than to the pediatrician's office. There the pediatrician checks the infant's symptoms and then gives him an injection. The injection makes the infant cry. A few days later the mother takes the infant to the pediatrician's office for a follow-up exam. Even though the aversive injection hadn't been given to him in the examining room, the infant begins to cry as soon as he and his mother enter the waiting room. When the pediatrician's significance was changed by the injection, the crying in the waiting room proved that the infant had learned the contingent relation between the waiting room and seeing the pediatrician. The waiting room and seeing the pediatrician are analogous to the initially neutral stimuli of a sensory preconditioning experiment (cf. Addendum 17A). Constructing other examples that correspond to the various cells in Table 18-1 is a worthwhile exercise.

One way to judge the relative significance of events is to compare the probabilities of the responses they occasion. If we were interested in the child's playing with toys and eating, we might see which he did when both the toys and food were freely available. When the child becomes older, an opportunity to play with friends might function to reinforce eating if he's reluctant to finish a meal, but if he's just encountered some cousins he's never met before at a family get-together and is reluctant to play with them, the opportunity to eat a favorite food may function to reinforce playing.

Table 18.1 Kinds of Response-Stimulus Contingencies and Contingent Stimuli

Type of Contingent Stimulus	*Response Raises Probability of Stimulus*	*Response Doesn't Affect Probability of Stimulus*	*Response Lowers Probability of Stimulus*
Appetitive, rewarding or reinforcing	Positive reinforcement (operant discrimination)	Stimulus presentation (respondent conditioning)	Negative punishment (omission training)
Relatively neutral or insignificant	Sensory consequences (latent learning)	Stimulus presentation (sensory preconditioning)	Sensory consequences (latent learning)
Aversive, noxious, or punishing	Positive punishment (discriminated punishment)	Stimulus presentation (defensive conditioning)	Negative reinforcement (discriminated avoidence)

*Entries are representative classes of experimental procedures; those in parentheses illustrate cases in which a discriminative stimulus is superimposed on the contingency.

In other words, the effect of a contingency may depend on the relation between the responses that produce and are occasioned by contingent stimuli. We are again reminded that establishing or motivating operations are concerned with the conditions determining the effectiveness of contingent stimuli as reinforcers or as punishers.

These classifications don't guarantee that any stimulus or response will have just a single function; a stimulus in a contingent relation with one response may be in a discriminative relation with another, and a response elicited by one stimulus may be involved in contingencies with other stimuli. For example, the mother's presence may be a contingent stimulus when she comes at the infant's cry, and a discriminative stimulus when the infant learns that things happen when she's there that don't happen when she's not; the infant's cry is sometimes elicited by events, such as painful stimuli, while at other times it occurs because such events as the mother's presence are its consequence. We come to understand behavioral situations by separating the various stimulus and response relations that enter into them. That is the business of an experimental analysis.

On close examination, some distinctions implied by our taxonomy seem to diminish in importance. In the analysis of behavior, classifications often have fuzzy boundaries and distinctions can become arbitrary. We noted such a case with respect to presenting or removing stimuli (e.g., is water effective as a reinforcer by virtue of its presentation or because it terminates dry mouth or other events correlated with thirst?). In the final analysis, we were able to discard the distinction between presenting and removing stimuli in favor of a behavioral account that considered the relation between the responding that produced the reinforcer and the responding produced by that reinforcer. That account recognized that all consequences of responding can be characterized as environmental changes.

We might then note that every procedure takes place in some environment, so we could next get rid of discriminative stimuli, observing as we did so that we can deal with them by incorporating them into our response definitions. Thus, if a child's requests for candy are granted when his grandparents are present but not when they're absent, the response of this contingency can be defined to include only those requests he makes in the presence of his grandparents. If later analysis shows us that we can't even define response classes independently of contingencies, we might ultimately decide to dispense with that distinction too. But distinctions like these have been useful along the way. We can, so to speak, throw away the ladder after we've climbed it.

ADDENDUM 18A: BIOLOGICAL CONSTRAINTS ON LEARNING

Chapter 3 examined the joint phylogenic and ontogenic sources of behavior and learning. Both sources may impose constraints not only on the stimuli and responses that enter into operant and respondent contingencies but also on the relations that can be established among stimuli and responses. This section considers a few examples.

Sensory Constraints

Some of the most obvious constraints on learning depend on the organism's sensory systems. For example, a pigeon is more likely than a bat to respond to visual stimuli, whereas a bat is more likely than a pigeon to respond to auditory stimuli. If an organism's sensory capacities aren't taken into account, experiments can yield misleading results. For example, the sound frequencies at which sensitivity is maximal are much higher for rat ears than for human ears. The experimenter who uses auditory stimuli that are easily heard by rats may be unable to tell whether the stimuli are on or off, but the one who uses auditory stimuli that are easily heard by humans may be giving the rat sounds that are hard for it to hear and may incorrectly conclude that the rat learns slowly and with difficulty.

Experimenters must be alert to the possibility that stimuli to which they're insensitive are important discriminative stimuli for the organism they're studying. For example, the results of early studies of maze learning must be interpreted with caution because rats have keen olfactory sensitivity. If a maze isn't thoroughly cleaned between subjects, a rat's performance may depend on odor trails left by other rats instead of what it learned on its own earlier runs through the maze. Similarly, if the odor of a food US reaches a dog during the presentation of the CS in a Pavlovian procedure, the salivation that follows the CS may depend on the odor instead of the CS-US contingency.

Constraints can involve stimulus configurations as well as single dimensions of stimuli. For example, in contagious yawning one person's yawn elicits yawns in others. The effectiveness of the yawn as an eliciting stimulus is determined by a complex combination of facial features that includes movements of the eyes as well as the mouth (Provine, 1989). In humans, the properties of faces that are involved in yawns and smiles and frowns have become important over a long phylogenic history of social behavior (Provine & Fischer, 1989). If they need to be learned at all, they are more easily learned than arbitrary geometric configurations.

Motor Constraints

Anatomical constraints on responding pose no problem. We don't expect flight to be the same in pigeons and bats and bees (and we don't even ask about the possibility of flight in rats). Species differences in motor capacity are more likely to raise questions when they don't have a clear anatomical basis.

In a study of leg movements in infancy, Thelen and Fisher (1983) recorded the timing and topography with which 3-month-old infants kicked at a mobile or hanging toy. The visual consequences of kicking varied: Some infants saw the mobile move when they kicked at it and others didn't. These consequences affected the rate and vigor of kicking but not the temporal coordinations among the flexion and extension phases of the kick. In other words, some features of the kicks were modifiable but others weren't. Similarly, contingencies may affect the direction in which you walk but not your detailed coordinations of muscles and joints as you do so.

Locomotion has both phylogenic and ontogenic components. The details of motor coordination don't arise from contingencies between responses and stimuli. Coordinations in walking involve relations among the muscles within a given limb and among the limbs and other parts of the body (e.g., the relaxation of one muscle as an opposing muscle contracts). Sometimes we can recognize people at a distance simply on the basis of their walks. Many aspects of these coordinations operate independently of the environment (cf. Gallistel, 1980; Gray, 1953); they are often called *motor programs*. The horse trainer doesn't have to shape the details of stepping or the order of leg movements as a horse speeds up from a walk to a trot to a canter to a gallop. A show horse might be taught special steps, such as the rack, but even in these cases the new topography modulates existing patterns.

Another example is flight in birds. How do the wings come to beat in synchrony? Must the bird fly to discover that it can't stay in the air by flapping with just one wing or by bringing one wing down while raising the other? Hatchling chicks were deprived of flapping and flight experience by wing restraint or other means and then their wing coordinations were tested at various later stages (e.g., Provine, 1981). Flapping was synchronous from the start, showing that this aspect of flight doesn't depend on environmental contingencies. Many aspects of flight coordination are *prewired*; they're built into bird behavior. Nevertheless, the environment remains important. The evolution of flight in birds depended on the aerodynamic environments of their ancestors. And whatever the details of its flight coordination, when a bird flies and where it goes are determined by the current environment.

Species differ in many ways, and constraints on the topography of responses mustn't be confused with constraints on their functions. We can illustrate this point by contrasting a cat stalking a mouse in a natural environment with a cow that's taught to stalk:

> given a prey which bears the same relation to a cow, in speed and mutual stimulation, as a mouse does to a cat, it should not be too difficult to set up contingencies under which a cow will "stalk"—that is, approach slowly when at a distance in order not to alert the prey and then move quickly to capture. The prey would have to be something like an animated bundle of corn. (Skinner, 1977b, p. 1011)

Skinner then points out that this "stalking" by the cow, in speed and other characteristics, would look quite different from the stalking by a cat. Nevertheless, the functional properties of the behavior of the cow and the cat would be similar, even if they differed considerably in structural details.

Constraints on Consequences

We may extend our examples to the capacities of various stimuli to reinforce or to serve as USs. Like sensory and motor capacities, these also differ across species. It hardly needs saying that the effectiveness of reinforcers has a phylogenic basis. An organism for which neither food nor water was ever effective as a reinforcer would hardly be likely to live long enough to pass its genes on to another generation. But more subtle properties of environments may also be important, such as the sensory consequences that maintain exploratory behavior or the novel consequences that may make an organism cautious in sampling an unknown food or familiar food in an unknown place, in the phenomenon called *neophobia* (e.g., Mitchell, Scott, & Williams, 1973).

Before the relativity of reinforcement was recognized, findings in which standard reinforcers for a species failed to have their characteristic effects were hard to deal with. Breland and Breland (1961) used several cases to argue against the generality of reinforcement as a behavioral process.

In one demonstration with raccoons, food was delivered when a raccoon picked up coins from the ground and deposited them in a container. After some repetitions of the procedure, the raccoon began persistently to rub the coins together instead of releasing them into the container. The Brelands and others saw this outcome as invalidating the principle of reinforcement. But a relevant aspect of raccoon behavior is that raccoons ordinarily rub and wash their food before eating it. The coins apparently provided a better opportunity for this behavior than the food that was supposed to function as a reinforcer. In other words, food wasn't effective as a reinforcer because rubbing was more probable than eating. It is likely that an opportunity for rubbing would have been effective as a reinforcer for other responses, perhaps including eating.

Given that reinforcement relations are based on relative probabilities of responses, the different behavioral hierarchies of different species inevitably constrain what they can learn. When an experimental procedure is applied, there's no guarantee it will be effective. A procedure that leads to learning with one response or organism may not do so with different responses or organisms, and a response or an organism affected by one procedure may be unaffected by others. An essential part of the analysis of learning is to explore such limits.

Preparedness

Constraints may also involve the relations between the stimuli and responses that enter into operant and respondent contingencies. Such relations were implicit in the examples of constraints on consequences, because we treated them in terms of the relative probabilities of the reinforced responses and the responses occasioned by the reinforcers. We considered other examples in the treatment of species-specific defense reactions in Addendum 8A. For example, the ease with which avoidance responding was acquired and maintained depended on species-specific relations between various avoidance responses and aversive stimuli.

Some relations between discriminative stimuli and responses may be easier to learn than others. For example, does it help if the stimuli and responses share common properties, as in responding on the left to a stimulus on the left and on the right to one on the right, instead of responding left to green and right to red? In the first case, locations are relevant properties of both stimuli and responses; in the second stimulus qualities are correlated with response locations (e.g., Miller & Bowe, 1982). Would a shift from location to a nonspatial dimension be easier or harder to master than a reversal of locations (recall Stratton's adaptation to inverting prisms in Chapter 5)? Such relations may be crucial in the design of human-machine systems and in the mastery of motor skills (e.g., Bauer & Miller, 1982; Glencross, 1977; Mazur, 1986).

We've considered cases involving relations between discriminative stimuli and responses, and between responses and reinforcers. Phylogenic contingencies may have prepared organisms to learn only some of the many possible relations among stimuli and responses in operant and respondent procedures. The concept of *preparedness* grew out of the observation that discrimination learning may be a function of the context of responses and reinforcers within which it occurs (Schwartz, 1974; Seligman, 1970). Its significance was established through the analysis of a phenomenon called *bait shyness* or *taste-aversion learning*, a variety of discrimination learning based upon differential punishment (but see Rozin & Kalat, 1971).

In experiments on taste aversion, a rat stops eating a food if it later becomes sick (Revusky & Garcia, 1970). Nausea or other systemic consequences of eating can punish eating even though they follow eating only after a considerable delay. In a study of the role of discriminative stimuli (Garcia & Koelling, 1966), thirsty rats drank water sweetened by saccharin; their drinking was accompanied by clicks and flashes of light triggered electronically by their licks. In other words, the rats drank water that was bright and noisy as well as sweet. In one group, drinking was followed by X-irradiation that later produced sickness. In a second group, drinking was followed by shock. Later, the drinking of each group was measured when a sweet solution was presented and when noise and light accompanied drinking.

Rats that had been X-irradiated drank less sweetened water, but their drinking was unaffected by noise and light; those that had been shocked drank less when noise and light accompanied drinking, but their drinking was unaffected by whether the water was sweetened. In other words, when the aversive stimulus was the delayed systemic consequence of X-irradiation, the rats learned only its relation to the taste of the water consumed earlier, but when the aversive stimulus was shock the rats learned only its relation to the noise and light that preceded it. The delayed effects of X-irradiation punished the drinking of sweetened water, and the immediate effects of shock punished drinking accompanied by noise and light. It's not enough to say that the rats learned some stimuli or some responses more easily than others; they were predisposed to learn different relations among discriminative stimuli and contingencies in these different situations.

Revusky and Garcia (1970) discuss these findings in the context of a thought experiment in which you find $100 that's been left for you by an insane billionaire who wants to get you to eat more gooseberry pie. Though you haven't been told about it, you received the money because you ate gooseberry pie for the first time two hours ago at lunch. Given all of the things you did between eating the pie and finding the money two hours later, it is unlikely that you'll start eating gooseberry pie more often. But suppose that instead of finding money two hours later you got sick. Despite the two-hour delay, this might put you off eating gooseberry pie again. In other words, you could learn a relation between flavor and sickness given the two-hour delay, but not one between flavor and money. "The fact that infra-humans also can associate over long delays strongly suggests that there is an innate selective association of flavors with physiological aftereffects and, what is more important, a selective failure to associate irrelevant stimuli with toxicosis" (Revusky & Garcia, 1970, p. 21).

Such experiments might of course be complicated by other differences between the two types of consequences. For example, electric shock has a more abrupt onset than do the gastric symptoms of X-irradiation. Furthermore, a case can be made that these experiments involve respondent relations rather than or in addition to operant ones. If respondent contingencies are arranged with taste as a trace CS and nausea as a US, the taste is likely to acquire its aversiveness as a result of its contingent relation to nausea even though the contingency operates over a delay.

In any case, this is another phenomenon with considerable practical significance. When radiation is used in medical procedures, such as some cancer treatments, the contingencies are analogous to those for the X-irradiation groups of the Garcia and Koelling experiment. Thus, the treatment setting itself may begin to elicit nausea and the patient's appetite may decrease as the palatability of foods eaten in the hours before treatment is affected. It may be appropriate to arrange treatments so that these effects are confined to a narrow range of settings and diets.

We've mainly considered phylogenic contributions to constraints. The issues can become even more complicated if we consider ontogenic as well as phylogenic contributions. For example, the adult sexual and maternal behavior of monkeys depends not only on the evolutionary contingencies that selected aspects of the rearing of young and patterns of mating. It also depends on contingencies of early development that involve contact with the mother and interaction with age mates. The deprivation of a young monkey's maternal contact during its first weeks may constrain its behavior throughout the rest of its life (e.g., Harlow & Harlow, 1966); touch may have special significance, but the temperature of a warm body versus a cold one may also be relevant (Harlow & Suomi, 1970). The analysis of behavior is concerned with identifying the origins of particular instances of behavior. With both operant and respondent contingencies, our conclusions about what is learned must take both phylogeny and ontogeny into account.

Chapter 19
Social Learning

The etymologies of the personal pronouns probably extend back to the very earliest of human languages. The histories of I *and its relatives (such as* ego *and* ich) *may be separate from those of* me *and* my *and* mine. *Some forms of* I *have a verblike quality, and pronouns are incorporated into verbs in some languages (e.g., as in the Latin* sum *and* es *for* I am *and* you are). *The distinctions among these and other personal pronouns, such as* you *and* she *and* they *and* he *and* us *and* it, *may be so fundamental as to be irreducible.*

Self, *which is related to* sibling, separate, select *and* ethnic, *carries an implication of possession (cf. the sense of its relative,* solitary, *as on one's own).* Other *has the etymological sense of the other one of two:* al- *(other) plus* ter *(two), as in* alternative. Community, *like* communicate, *implies a having in common.* Common *is a derivative of* con-, *with, and the Indo-European* mei-, *to go or move together with; some relatives are* immune, mutate, migrate *and* mean.

- Kinds of Social Contingencies
  - Learning about Others
  - Learning from Others
    - *Observational Learning*
    - *Imitation*
    - *The Social Origins of Language*
  - Learning about Oneself
    - *Discriminating Properties of One's Own Behavior*
- The Selection of Cultural Contingencies

One variety of selection is the sort that operates on populations of organisms over successive generations. We considered it in the context of Darwin's treatment of evolution in terms of natural selection. Another variety is the sort that operates on populations of responses within the lifetime of the individual organism. Much of our discussion of the effects of consequences, especially in the shaping effects of reinforcers, was concerned with this variety of selection. Organisms persist in doing some things and stop doing others; shaping provided an explicit example of the selection of behavior by its consequences.

Chapter 3 also introduced a third variety of selection. It too operates on behavior, but it involves more than one organism. Any behavior an organism acquires within its own lifetime is eventually lost if it is not passed on to others. Once social learning becomes possible, behavior that has been learned can survive the death of the

KEY TERMS: Social Learning, Social Facilitation; Observational or Vicarious Learning; Modeling and Imitation, Generalized Imitation; Self-Awareness, Self-Control and Self-Regulation.

organism that learned it. It then survives in what others do, perhaps not only in the behavior of descendants but even in the behavior of contemporaries who need not be closely related. This kind of selection, sometimes called *cultural* or *memetic selection*, involves the selection of behavior as it is passed on from one organism to another, as when local languages or ethnic cuisines or religious practices spread within or across cultures.

KINDS OF SOCIAL CONTINGENCIES

Learning from others is especially important in human behavior. A very substantial part of what any of us knows is what we've learned from others, and much of it has been explicitly taught, either in the informal context of interactions between family members and friends or in the formal context of educational institutions. But the earliest types of social learning must have been much simpler. At what point did some organisms begin to learn to do things just by watching what happened as other organisms did them?

We may think of learning by observation as commonplace because it so often enters into human behavior. But it is not clear how much of it goes on even among nonhuman warm-blooded vertebrates (mammals and birds), and there is hardly any evidence for it at all in invertebrates (e.g., insects), though chemical and other signals play a key role in coordinations within colonies of social insects such as ants and bees. Among primates, one widely cited example involved the feeding of sweet potatoes to macaque monkeys on a Japanese reserve adjacent to a beach (1959). The potatoes typically became sandy, but one juvenile monkey eventually discovered that rinsing them in the ocean removed sand from them. The practice then spread to other juveniles of the colony, but it was not taken up by the adults.

Human cultures offer many examples of the social selection of behavior, the selection that occurs as behavior is passed on from one individual to another (Harris, 1977; Skinner, 1981). Certain ways of raising children, getting and preparing food, building shelters and dealing with group members and outsiders survive over successive generations. Some practices can spread to other groups who aren't close genetic relatives of those who began them. In contemporary Western culture, for example, ethnic foods are prepared and eaten by many who aren't members of the ethnic groups that originated them.

Some behavior that is socially transmitted survives because of its consequences. In early human history, the person who learned from someone else how to make stone tools or fire or garments was probably also more likely to survive long enough to pass the behavior on to someone else than a person who couldn't learn in that way. The survival of other patterns of behavior may involve more complex contingencies.

Consider patterns of child rearing. Suppose that children reared according to most patterns are as likely as adults to rear their own children according to one pattern as according to any other. But suppose also that a very few patterns work so that when the children become adults they are likely to rear their own children in the same way as their parents reared them; we might call these self-replicating patterns of child rearing. In a large population, whenever one of the self-replicating patterns happens to be used by some parents, for whatever reasons, it will be used again in the next generation; the other patterns will come and go. Little by little, over many generations, the patterns that are self-replicating will displace the ones that aren't. Once traditional patterns of child rearing originate in this way, they are likely to survive for a very long time, perhaps even in the face of powerful countervailing contingencies such as discriminatory practices.

In earlier discussions of the first two varieties of selection, natural selection and operant selection, we noted that evolution and shaping depended upon variable populations on which selection could operate. Similar constraints exist at the level of cultural selection. For example, cultural practices that favor ethnic diversity may have advantages over ones that don't simply because

they allow such variability. In every variety of selection, we ordinarily look at how selection operates at the level of the survival of individual members or units rather than at the level of the survival of the group, though the need for that constraint has sometimes been questioned (e.g., D. S. Wilson & Sober, 1994). For example, if variation itself can be selected, it is difficult to define variation as a property of an individual as opposed to a property of the population in which that individual is included (cf. Chapters 10 and 13).

Learning about Others

In many situations, the discriminative stimuli provided by other organisms are more important than those provided by inanimate objects and events. For example, parental investment in offspring may be wasted if the parent can't discriminate between its own offspring and the offspring of others; potential mates must be distinguished from potential competitors, and among the potential mates the receptive must be distinguished from the unreceptive; and so on. In many organisms, such properties are correlated with anatomical features (e.g., colorful plumage in birds); often, behavior is the crucial dimension.

The study of animal communication is concerned with the many ways in which organisms produce stimuli that affect the behavior of other organisms. "Examples of communication are numerous: song in birds, frogs, and crickets; tail-wagging and hackle-raising in dogs; 'grinning' in chimpanzees; human gestures and language" (Dawkins, 1976, p. 67; see also Dawkins & Krebs, 1978). It's more appropriate to speak of such stimuli in terms of their behavioral effects than in terms of the information they carry. Just as it may be misleading to speak of genes as carriers of information about phylogenic contingencies (Dawkins, 1982), it may be misleading to speak of social stimuli as carriers of information (cf. Chapter 22).

Releasers and fixed action patterns provide many examples of effects of social stimuli. In some cases, sometimes called *social facilitation*, the behavior that serves as a social stimulus and the behavior produced by that stimulus are topographically similar, as when one bird's takeoff triggers the flight of the other birds in a flock, or as when a few galloping steers set off a stampede. Such cases may superficially look like imitation, but they are limited to a narrow range of response classes (e.g., contagious yawning or laughing: Provine, 1989, 2004, 2012) and so must be distinguished from it (cf. Field *et al.*, 1982).

Discriminating the behavior of other organisms, whether of one's own or other species, has clear selective advantages. Consider, for example, the relation between predator and prey. If one antelope in a herd is limping just a little, the lion that notices the limp may be more likely to make a capture. The antelope that can tell the difference between one lion that hasn't eaten for a while and another that's just finished eating may be more likely to move away in the safest direction. A predator that can distinguish whether it has been noticed by its prey has a distinct advantage over one that can't; an advantage also accrues to a prey that can distinguish whether it has been noticed by its predator. Such discriminations presumably have an extensive phylogenic history. Attention to the behavior of prey must have been one dimension upon which natural selection operated in the evolution of predators, just as attention to the behavior of predators must have been one dimension upon which natural selection operated in the evolution of prey. Following from such selection, discriminations of social behavior may become so important that they override other types of discriminations.

Discriminations of the behavior of others are at the heart of our human concept of *intentionality* (cf. Dennett, 1987): I say I understand someone's intentions when my discriminations of the properties of that person's past and current behavior enable me to act appropriately with respect to what that person may do in the future. In fact, if discriminating my own behavior is a special case of discriminating the behavior of others (e.g., Bem, 1967), it can be argued that this topic encompasses all of the phenomena considered under the rubric of intentionality. Our judgments of the intentions of others are, above all, social judgments, and it takes no special assumptions about the selective contingencies that must have operated on social behavior both within and across species to see that such contingencies could shape well-prepared capacities for social discriminations.

Social discriminations within species have many functions. They may operate within dominance hierarchies or in the defense of territory or as isolating mechanisms that maintain the integrity of a group against intrusions from outsiders or in the distribution of limited resources among group members. Within species that live as social groups, such as most primates, individuals learn what sorts of behavior to expect of others with whom they have extended contact. The cooperative behavior that can emerge in such contexts (de Waal, 1989) requires social discriminations that may be the precursors of the behavior toward others that is called empathy (e.g., Hoffman, 1975; E. O. Wilson, 2012).

Learning from Others

It's one thing to learn about other organisms. It's another to learn something from them (Zentall & Galef, 1988). Sometimes the behavior of one organism allows another to act on the basis of stimuli available only to the first, as when a vocal call from one monkey allows another monkey to escape from a predator it hadn't seen. Warning calls are well-documented in bird behavior (e.g., Kroodsma & Miller, 1982). In monkeys, predator calls can vary with kinds of predators, and the response to the call can depend on who the caller is and who the listener (e.g., Gouzoules, Gouzoules, & Marler, 1984; Seyfarth, Cheney, & Marler, 1980a).

Observational Learning. Learning based on observing the behavior of another organism is called *observational learning* (e.g., Zentall & Levine, 1972); another occasional term is *vicarious learning* (e.g., Bandura, 1986). Sometimes what seems to be observational learning involves simpler processes. For example, food preferences in rats are learned in social contexts; it's hard to get rid of rats by poisoning because rats that don't ingest the poison may avoid it after interacting with other rats that did and became sick or died. When rats come together they sniff and lick each other, so by smell and taste one rat can become familiar with properties of the food the other has recently eaten (Galef & Stein, 1985). If the food is novel and the other rat is healthy, the first rat will later prefer that novel food over other novel foods, but if the other rat is sick, the first rat will avoid that novel food. In other words, this is a socially mediated taste aversion (cf. Chapter 18). The learning won't occur without some contact between the two rats (e.g., mouth to mouth or mouth to fur). One rat has learned about some novel foods from another, but only in the sense that the combination of food stimuli with social stimuli makes some foods more or less effective as reinforcers or as aversive stimuli, and not in the sense that one rat has learned something on the basis of observing what happens to another rat.

Observational learning has been compellingly demonstrated with rhesus monkeys (Mineka *et al.*, 1984). Monkeys in the wild show fear of snakes by screaming and other agitated behavior and by avoidance of the snake. Even if their parents fear snakes, monkeys reared in the laboratory who have had no experience with snakes don't; for example, if food is on the other side of a container with a snake in it, they'll reach across the container for the food. But if the laboratory-reared monkeys then briefly observe one of their parents behave fearfully toward snakes, they too become fearful. Their fear is intense and persistent; if tested three months later it's undiminished. What they've learned about snakes is based only on observing the parent's behavior toward a snake. There may be a phylogenic component, however, because such observational learning is more likely to occur with snakes or snakelike objects than with some other types of stimuli. These kinds of learning may play a critical role in the onset and maintenance of post-traumatic stress disorders (PTSD) and other human anxiety disorders (Mineka & Zinbarg, 2006).

Observational learning is a prerequisite for other varieties of learning and social interaction. For example, joint attention, in which a child looks where a caregiver is looking or pointing, depends on the child first learning to attend to the caregiver's looking and pointing (Dube *et al.*, 2004; Jones & Carr, 2004; Whalen & Schreibman, 1997).

Observational learning is sometimes treated as if it were itself a fundamental kind of learning (e.g., Bandura, 1986), but it is probably better treated as a higher-order variety of behavior. Many different skills have to come together appropriately for observational learning to work, and in the case of human observational learning it is likely that there is also a very large verbal component (Catania, 1995a; cf. Chapters 22 and 23). At the least, observational learning must include subtle discriminations of another organism's actions and their outcomes (cf. Asen & Cook, 2012) and some history with respect to the effects of related actions on the part of the observer. When we see observational learning, we should try an analysis to determine its components instead of trying to use it to explain other, more complex kinds of behavior.

Imitation. The most important difference between observational learning and imitation is that in imitation the observer's behavior corresponds to that of the organism that has been observed. Thus, imitation might be regarded as a special case of observational learning. Imitation does not imply that the imitating organism has learned something about contingencies, so not all imitations are advantageous. A coyote that sees another coyote step into a trap would do well not to imitate that behavior. A hatchling bird that hasn't yet grown its flight feathers would do well not to follow its parents when they fly from its treetop nest.

Following may sometimes be imitative. In one experiment (Neuringer & Neuringer, 1974), food-deprived pigeons learned to eat from the experimenter's hand. When the hand then approached and pecked the key, producing food, the pigeons followed the hand and began to peck the key. The procedure often worked more rapidly than shaping of the key peck. Under natural conditions, young animals may learn to behave like their parents simply by following parental sources of food. But not all following is imitative. For example, when one rat leads another to food, if the proximity of another rat is a reinforcer nearness to the leader rat may reinforce the behavior of the follower. Following will then emerge as behavior shaped by natural contingencies (cf. the duckling's following of an imprinted stimulus; Chapter 4).

We call responding *imitative* when one organism duplicates the behavior modeled by another organism. But one kind of imitation may be limited to the duplication only of specific instances that have been explicitly taught, while another may include correspondences between the behavior of model and observer even in novel instances, when it is called *generalized imitation*; in the latter case, imitative responding is a class of responses that may be differentially reinforced or, in other words, it is a higher-order class of behavior (e.g., Baer, Peterson, & Sherman, 1967).

For example, suppose a child learns to imitate several instances of behavior modeled by a puppet, such as jumping, skipping and clapping. If we arrange consequences to maintain the child's imitation of all but the clapping, the child will ordinarily continue to imitate the clapping along with the others even though this imitation never produces those consequences (Secan, Egel, & Tilley, 1989). To the extent that imitation of clapping doesn't extinguish, we say it is a member of the generalized class. With generalized imitation, the child will also produce novel imitations if the puppet does something it has never modeled before, such as stamping its foot; in fact, responses occasioned by novel stimuli define generalization, as when a pigeon pecks when it first sees yellow or blue after a history of pecking only during green. Because *modeling* of a new response can produce novel behavior only if a child's imitation has generalized, it is an effective supplement to shaping, as in teaching skills to children with autism or other developmental disabilities (e.g., degli Espinosa, 2011; Peterson, 1968).

We don't know the pertinent physical dimensions of imitative behavior. For example, contingencies can be arranged to create classes that include both nonimitative and imitative responses (Garcia, Baer, & Firestone, 1971) or that include imitations only within certain topographical boundaries. Echoic verbal behavior can be viewed as a variety of imitation, but the different time courses of their development suggest that neither is a generalized form of the other (Camões-Costa, Erjavec, & Horne, 2011; Erjavec & Horne, 2008; Poulson, 1984; Poulson & Kymissis, 1988).

These characteristics of imitation are consistent with what we understand about the contingencies that create higher-order classes, especially given that the behavior of model and observer may look the same to us but we cannot be sure that this is also true for the observer. For example, if you touch the top of your head when that is modeled in the game of "Simon says," you see but don't feel the leader's hand and you feel but don't see your own hand. There is no simple correspondence between someone else's seen hand and your own felt hand, so how did you learn to imitate? Some cells in the primate brain, called *mirror neurons*, fire both with limb movement and with the observed movement of that limb by another primate (Arbib, 2005), but it is not yet clear how these neurons are involved in the initiation of movement as opposed to movement discriminations.

One way you can learn correspondences between the seen and felt parts of your body is by behaving in front of a mirror. Humans and some primates seem to learn these correspondences without explicit training (e.g., Gallup, 1979). For example, if a chimpanzee experienced with mirrors has a spot painted on its eyebrow while it is asleep, it will touch the spot the next time it sees its face in a mirror. Some monkeys pass this test and others don't, but the differences may depend not on cognitive capacities but on behavior that interferes with looking in mirrors (cf. Anderson, 1984; cf. Rajala *et al.*, 2010). For example, species of monkeys differ in *gaze aversion*, the extent to which they look away if another monkey looks at them; thus, when some monkeys look into a mirror and see themselves looking back, they will look away before seeing spots on their eyebrows.

Such responses to one's own body have sometimes been spoken of in terms of *self-awareness*. A monkey who passes the mirror test is sometimes said to have a *self-concept*, but it is more appropriate to ask about the histories that can create such behavior than to invoke some inner entity to explain it.

With other organisms, such as the pigeon, these correspondences have to be taught, by training discriminations among stimuli seen in a mirror. For example, first a pigeon's pecks at blue cardboard dots pasted at various places on its body were shaped. Next, discriminations among blue dots reflected in a mirror were taught by presenting the dots behind holes in one wall only when the pigeon faced a mirror on the opposite wall; each dot was gone by the time the pigeon turned around, but only pecks at the hole where it had appeared were reinforced. When another blue dot was then attached to the pigeon's breast while it wore a short bib that allowed it to see the dot in the mirror but not by looking down, the pigeon pecked down toward the dot on its body even though it was seeing the dot only in the mirror (Epstein, Lanza, & Skinner, 1981). Did the monkeys who failed the mirror test simply lack some relevant experience with mirrors?

The pigeon's novel performance doesn't demonstrate a pigeon self-concept or sense of self, but it does illustrate a special case of discriminative control in which stimuli and responses vary together along a dimension in *continuous repertories* (e.g., Wildemann & Holland, 1972), in which continuous changes in one dimension produce corresponding changes in the other. Discriminated behavior with respect to a mirror involves correspondences between continuous movements and the changing seen positions of one's body. Other instances include keeping a moving object framed in a camera, mixing paints to match a sample color, tuning a musical instrument, and steering a car along a winding road. In human behavior, imitation is a particularly important example of this type of stimulus-control relation; for the individual who already has an established imitative repertory, new behavior can often be generated more quickly and effectively this way than by shaping or other means.

The Social Origins of Language Another way to learn from another organism is through verbal behavior: you can be told about contingencies instead of observing them. But verbal behavior can't have originated in that way, because descriptions of contingencies require sentences and the earliest forms of language must have begun with single words. It's a reasonable guess that functional single-word utterances (protolanguage) have a history of a million years or more in our hominid ancestry; perhaps these units sometimes combined in functional ways, but their combination in the grammatical organizations we call language probably originated only as recently as forty or fifty thousand years ago.

Certainly anatomical and genetic changes played roles: "...the reduction in jaw musculature may have facilitated the eventual evolution of finer control of the mandible, as is required for speech" (Carroll, 2005, p. 273; see also pp. 273-275 on the MYH16 and FOXP2 genes). But, as we shall see, human verbal behavior involves the coordination of many different components, so it is best treated as a system that has evolved and is still evolving rather than as some unitary capacity that somehow emerged full-blown in all its complexity and variation. "There is no need to invoke single dramatic mutations as causes of great leaps in form and function or as explanation for the origins of human traits. Nor is there any scientific foundation for doing so" (Carroll, 2005, p. 277). We'll have more to say about verbal behavior in Part IV, but here we'll deal briefly with its possible origins in human social contingencies.

The simplest and most obvious function of verbal behavior is instructional: It is a way in which one organism gets another to do something. By talking, we change each other's behavior. In verbal behavior, we do things via the mediation of another organism. Sometimes what gets done involves nonverbal effects, as when we ask someone to move something or to carry something to us; I might have difficulty lifting you from a sitting to a standing position, but if I asked you to stand and you complied I'd have accomplished the same objective with far less effort. Sometimes it involves verbal effects, as when we change what someone else has to say about something; throughout this book I've indirectly been asking you to talk about behavior in new ways. If the primary function of language is that it is an efficient way in which one individual can change the behavior of another, it follows that this behavior is quintessentially social and can emerge only in organisms whose behavior is already sensitive to social contingencies.

Assume that the calls of a primate leader once determined the behavior of members of its band as reliably as a releaser elicits a fixed action pattern. At first the vocabulary of releasers was limited to just a few calls, not yet qualifying as verbal behavior but with relatively simple effects corresponding to those of words such as *come* or *go* or *stop*. Over many generations, perhaps millennia, a more extensive repertory of more varied calls was differentiated. If the details of these calls were weakly determined by phylogenic contingencies, this rudimentary vocal control could later be supplemented by variations produced by ontogenic contingencies. For example, a dominant speaker might learn to attack a listener who doesn't respond in the characteristic way, thereby punishing disobedience (as we'll see in Chapter 23, many contemporary contingencies continue to maintain the effectiveness of verbal control by reinforcing the following of instructions and by punishing deviations from it).

Once vocal behavior had expanded to an extensive repertory including arbitrary as well as phylogenically determined calls, idiosyncratic repertories developed by particular leaders would ordinarily be lost to later generations unless some way of reproducing this behavior in the leaders' successors had evolved. Thus, the next step in this evolution, perhaps long in coming, was the repetition by the follower of the leader's verbal behavior . Once some individuals began repeating what others said, verbal behavior became a kind of behavior that could survive within the behavior of the group, as a candidate for the third type of selection that we discussed earlier, that of cultural transmission (Catania, 2001b; see Jaynes, 1976, for an alternative scenario). The stage was then set for human verbal memory, for instruction and educational systems, and for the rapid and wide dissemination of cultural practices.

Particular languages also evolved, perhaps from a single source in a language sometimes called Nostratic, from which the Indo-European family of languages would have descended (Watkins, 2011). Darwin (1871) discussed the evolution of languages through selection as analogous to the selection of species in phylogeny. Historical changes in the syntax, semantics and phonology of languages are consistent with selective accounts (Deutscher, 2005, 2010; Vihman, 1996). They are sometimes paralleled by similar changes in language development in ontogeny, but with language as with development in general the assumption that ontogeny recapitulates phylogeny is likely to be misleading (Gould, 1977).

Meanwhile, along the way came other human achievements, such as art and music and religion and science. Human history is only a short piece of this chronology, so we can only guess at the timing of these steps in our evolution. But perhaps we would do well always to focus first on simpler functions of behavior rather than imputing special significance to more esoteric ones. For example, tens of thousands of years ago in the long darknesses of the European ice ages it would have been dangerous to teach the youngsters what they needed to know if they had to be taught while in the midst of hunting bisons and stags and bears. Maybe a few wise elders lit up caves so they could paint pictures on the walls to demonstrate hunting tactics and display special features not only of their prey but also of the big cats and others that might sometimes be their predators rather than their prey. That would explain why so much early cave painting is dedicated to animals and their hunters and so little to vegetation and landscapes. Perhaps we should look at those caves as early human classrooms instead of as art museums or as depictions of the imaginings of Cro-Magnon adolescents (Whitley, 2009). Doing so does not in any way diminish human art and creativity but rather places it in the context of other human achievements that are even more fundamental.

Learning about Oneself

We've discussed imitation in terms of the relation between our own behavior and the behavior of others. Implicit in that discussion was the suggestion that we learn to discriminate properties of our own behavior in the context of learning about others. The case is even more obvious with verbal behavior, because we learn the language with which we describe our own behavior from others. It follows that what we know about ourselves is a social product. We don't see ourselves as others see us; instead, we see ourselves as we see others.

Let's begin with a human example from a verbal learning experiment (Vesonder & Voss, 1985). The experiment included three kinds of participants: learners who talked aloud while learning verbal items and then predicted how well they'd remember the items on the next presentaton; listeners who heard what the learner said and made similar predictions based on what they heard; and observers who made predictions based on how well learners did on past items without hearing what the learners said. The predictions of both the learners and the listeners were substantially better than those of the observers; the crucial point, however, was that the predictions of the learners and the listeners were essentially the same. The public behavior of the learner, to which the listener also had access, was good enough for the predictions; if the learner did know private things to which the listener had no access (e.g., levels of confidence), they didn't make the learner's predictions any better. The findings are consistent with evidence suggesting that, just as we judge others on the basis of observations of their behavior, we judge ourselves on the basis of observations of our own behavior (e.g., Bem, 1967; Nisbett & Wilson, 1977).

Discriminating Properties of Our Own Behavior. The capacity to discriminate properties of our own behavior is important in many types of human behavior. For example, the student who can't tell the difference between a superficial and a thorough reading of a text may stop studying too soon. To the extent that the behavior of interest has environmental effects, it is sometimes difficult to distinguish between control of discriminated responding by the behavior itself and control by the stimuli produced by that behavior. In some other contexts, this might be called the difference between knowing and feeling that you know (e.g., Koriat, 1995). For example, some students might judge their exam performances primarily on the basis of problems they encountered in answering particular questions, whereas others might judge them primarily on the basis of the consequences of their performances (e.g., the grades later posted).

The nonhuman synthesis of such discriminations is again of potential interest. In one procedure, a pigeon's pecks on a center key were followed by the lighting of two side-keys; pecks on the left one were reinforced if the pigeon had emitted 50 or fewer responses on the center key, and pecks on the right were reinforced if the pigeon had emitted more than 50 (Pliskoff & Goldiamond, 1966; cf. Reynolds, 1966; Shimp, Sabulsky, & Childers, 1989). The pigeon's side-key pecks depended on the number of its center-key pecks, so we can regard them as responses under the discriminative control of a property of the pigeon's own behavior. Stimulus control by number of responses and other properties of behavior such as the temporal patterning of responses have been demonstrated in a variety of experiments. Another important discriminable dimension is deprivation (e.g., Capaldi & Davidson, 1979).

But one technical problem in such procedures is that the organism's behavior is usually correlated with environmental stimuli. For example, the pigeon that has spent a longer time pecking has probably also spent more time in a position where it's been looking at the key. Is its discrimination then based on its behavior or on what it has been looking at?

In discriminating our own behavior, stimuli are of course available from our muscles and joints and so on. These stimuli are called *proprioceptive* or *interoceptive*. The effects of *biofeedback* may depend on the ways in which such stimuli are augmented or supplemented (e.g., Hefferline, 1958). For example, if the reading on a meter or the loudness of a tone is made proportional to the electrical activity of a muscle, an individual may learn to control levels of muscle tension and relaxation. Such feedback procedures have been extended to other systems, such as the aura that precedes some epileptic episodes or the muscle and blood vessel events correlated with some migraine headaches (e.g., Sturgis, Tollison, & Adams, 1978). It seems plausible that biofeedback control should be correlated with discriminative capacity, but it remains an experimental question whether biofeedback control of a given physiological system can be shaped based on discriminations of proprioceptive or interoceptive stimuli (e.g., Cott, Pavlovski, & Black, 1981).

Now consider one more nonhuman example. Suppose we give a pigeon a dose of cocaine, pentobarbital or saline (an uncontrolled substance that here serves as a control). We then intermittently reinforce pecks on one of three keys depending on which drug we've administered. The pigeon can learn to discriminate among these three substances in this procedure. Once the pigeon has done so we can study the properties of other dosage levels of these and other substances by administering them and observing which key the pigeon pecks (e.g., Lubinski & Thompson, 1987). For example, we can find out how small a dose of cocaine the pigeon can discriminate from saline, or whether the pigeon can discriminate pentobarbital from other sedative drugs it's never been given before. In effect, we've taught the pigeon to report an internal condition, its own drug state. This example shows how we can use the public correlates of private events to generate reports of those events. Notice that to do this, we don't have to know what the drug state feels like; we only have to know what's circulating in the pigeon's bloodstream.

We'll return to the issue of discriminating private events in Chapter 22. The practical significance of being able to do so should be obvious. For example, the drinker who is a good judge of blood alcohol levels should know when to hand the car keys over to someone else (cf. Critchfield, 1993). Another important feature of the drug example is that the pigeon was taught the drug discrimination by a human experimenter. In human behavior, such discriminations may sometimes be incidentally learned, but the point is that, like the pigeon, we're most likely to learn them when they're taught to us by others. In other words, discriminations of our own behavior very often originate in the context of social behavior.

We've already mentioned other circumstances in which discriminations of our own behavior are important (e.g., as in judging how well we've studied a text). Such discriminations are also critical to a phenomenon once called *self-reinforcement* but now more properly referred to as *self-regulation* (Bandura, 1976, 1995; Catania, 1975b, 1995; Goldiamond, 1976; Mahoney & Bandura, 1972). For example, a student who has made a commitment to watch television only after completing a study assignment might think that this arrangement for watching television after studying will reinforce studying. But any increase in studying that follows can't be attributed to reinforcement; the student will only make the commitment to deal with studying this way if studying has already become important for other reasons. Whatever brings the student to make the commitment to "self-reinforce" studying will probably by itself make studying more likely. It is impossible to pull these variables apart.

In other words, any effect that so-called self-reinforcement might have cannot be attributed to the action of the reinforcers delivered by the organism to itself. The contingencies and establishing events modifying the behavior that is purportedly reinforced cannot be separated from those that modify the behavior of self-reinforcing. This doesn't mean the procedure can have no effects; instead, any effects it does have must depend on how well the individual who appears to "self-reinforce" can discriminate behavior that qualifies for the reinforcer from behavior that does not.

The commitment to reward one's own studying involves setting standards for the discrimination between adequate and inadequate studying, so students who try to deal with their study habits in this way are discriminating properties of their own behavior that have become important to them. The contingencies that generate these discriminations are complex and inevitably involve verbal behavior. The language of self-reinforcement obscures rather than clarifies these phenomena. Talking about this behavior as self-regulation also finesses the problem that the language of self-reinforcement implies reinforcement of the organism rather than of behavior. We considered related issues in the topic of self-control in Chapter 16; we'll treat how discriminations of our own behavior are relevant to language in Chapter 22, when we deal with the class of verbal responses called autoclitics.

THE SELECTION OF CULTURAL CONTINGENCIES

Each of the three varieties of selection involves some kind of variation that provides the source materials upon which it operates, and each has some mechanism or mechanisms according to which surviving variants are selected. We saw in Chapter 3 that the different varieties of selection often complement one another, but it is also possible for selection at one level to oppose selection at another. For example, behavior may be selected through ontogenic processes during the lifetime of an individual even though it reduces the individual's reproductive fitness at the phylogenic level (consider substance abuse or sexual behavior involving partners at risk for AIDS). Phylogeny produced the capacity for behavior to be selected by its consequences in ontogeny, but this does not imply that the individual will always benefit from that capacity, though phylogenic selection may eventually change the sensitivity to such ontogenic selection within a population.

Early in this chapter we considered self-replicating patterns of child rearing, in which we supposed that some patterns might work so that when the children become adults they are likely to rear their own children in the same way as their parents had reared them. Once in place, such patterns of child rearing are likely to survive for a very long time. But now suppose that some patterns of child abuse are more likely to be repeated over generations than other more benign varieties of child rearing, in that an abused child is more likely to become an abusive parent. Such abusive patterns might be less likely than others to be maintained by reinforcement contingencies and might also reduce the probability that children of abusing parents will survive to maturity. But if these patterns are more strongly self-replicating than the others, they may nevertheless become dominant in the population even though their selection is opposed by phylogenic and ontogenic contingencies.

A case in point is the application of techniques of reinforcement relative to those of punishment. The advantages of reinforcement do not make it more likely than punishment to spread through a culture. The problem is that delivering a punisher typically produces more immediate effects than delivering a reinforcer. Whatever else happens over the long term, a parent who strikes a misbehaving child is likely to see an immediate change in the child's behavior. The behavior of concern to the parent will probably stop, at least for the moment, though that may have little to do with whether it will reappear later. If stopping the child's behavior is part of what reinforces the parent's application of punishment, the immediacy of that reinforcer will be an important factor in the maintenance of the use of punishment by the parent. The effects of a reinforcer, on the other hand, may not appear for a while, so, perhaps as this text demonstrates, it is far more difficult to teach reinforcement than to teach punishment. Should we be surprised, then, that punishment is far more common than reinforcement in many human institutions throughout the world?

An analysis of reinforcing effectiveness at the individual level may not be consistent with the effects of cultural selection at the level of the population, as when behavior that is transmitted from one individual to another only with difficulty is displaced by competing behavior less likely to be reinforced but more easily transmitted. Patterns of behavior that depend on the survival of certain types of verbal behavior among the members of a group (perhaps including celibacy, suicide, and modern war) need not be consistent with phylogenic and ontogenic contingencies. This is why the domain of cultural selection in its interaction with phylogenic and ontogenic contingencies may require different kinds of analyses.

PART IV
LEARNING WITH WORDS

Chapter 20
Verbal Function: Formal Classes

Echoic *comes to us from the Latin and the Greek* echo, *to resound, personified as a mountain nymph.* Textual *comes from the Latin* texere, *to weave, later extended to anything written.* Dictation *is derived from the Latin* dicere, *to say or tell; its relatives, via the Indo-European* deik-, *to point out, include* digit, dedicate *and* teach. *From the Latin* trans-, *across, and* scribere, *to write,* transcription *is writing from one place to another;* scribere *and* script *come to us from the Indo-European* sker-, *to cut, perhaps in the sense of carving in wood or stone.* Write, *from Germanic and Old English, similarly has origins in earlier words with the sense of scratching or carving.*

- Correspondences Between Spoken and Written Classes
- Echoic Behavior
    - *The Development of Echoic Behavior*
    - *Categorical Perception of Phonemes*
- Transcription
- Textual Behavior
- Dictation Taking
- Relations Among the Classes
    - *The Replication of Verbal Behavior*
    - *Parallels in Music*

B. F. Skinner regarded *Verbal Behavior* (Skinner, 1957) as his most important work. I was fortunate to get the chance to read it early. During the Spring semester of 1957, along with a half dozen or so other Columbia College undergraduates, I took a weekly seminar on verbal behavior that was jointly conducted by three faculty members: Ralph F. Hefferline, Fred S. Keller, and W. N. Schoenfeld (Catania, 1997a). During the first several weeks we spent our time on a mimeographed document, B. F. Skinner's *William James lectures*, which he had delivered at Harvard University in 1947 and which was a forerunner of the book. Each class began with a thorough summary of the current material by one of the faculty members. We finished with the lectures somewhere midway through the course, when the book based on them was published and finally became available. As we progressed through the book we had the advantage of having read the earlier lectures, with which we compared it.

KEY TERMS: Verbal Behavior, Spoken and Written Verbal Behavior, Vocal Behavior; Formal Verbal Classes; Echoic Behavior, Transcription, Textual Behavior, Dictation-Taking; Replication of Verbal Behavior

I didn't pay much attention to a critical review of the book that appeared not long after (Chomsky, 1959; cf. Knapp, 1992). We can come back to that later. But once I had an academic position I eventually begin to teach from the book. I've since taught both undergraduate and graduate courses based upon it. What is special in it is its thoroughly consistent treatment of verbal behavior as behavior in its own right rather than as a vehicle for something else such as meanings or thoughts. It isn't easy to avoid the temptation to say that our words are filled with ideas, but once we do so we can see some properties of verbal behavior more clearly. One reason is that our language of words is more concerned with their structure than with their function, or what they do. I hope these chapters on verbal behavior will give you some practice in thinking about words in this different way.

Our talking and listening and reading and writing are all behavior, but our everyday vocabulary gets in the way when we talk about them. For example, consider the common term *word.* When we speak of words, we seldom bother to distinguish spoken words from written ones. Yet speaking a word isn't the same as writing it, and speaking and writing usually occur in different circumstances. Even worse, we often speak of *using* words, as if words were things instead of behavior. I don't say I use a reach when I pick something up, but I might say that I use a word when I say something (cf. Skinner, 1957, p. 7).

We also speak of our words as if they were directed toward events or objects. We say that words or sentences refer to, deal with, speak of, call attention to or are about things. The language of reference implicitly includes a direction from verbal behavior to environment. Everyday language doesn't include words emphasizing the other direction. Has our everyday language prejudiced us about how our verbal behavior works? We hardly ever say that we utter nouns in the presence of relevant objects or that sentences are occasioned by relevant events. Instead, we say that words refer to objects or that sentences are about events. There is good reason for these usages; they are appropriate to the equivalences that relate different classes of verbal behavior. But in analyses of the behavior of speakers and listeners or readers and writers they can be misleading.

The language of meaning is another complication. Dictionaries don't contain meanings of defined words; they just contain other words. We speak metaphorically when we say that words *contain* meanings and that we *convey* these meanings to others through language (cf. Chapter 22). The metaphor of words as containers for meanings has been around for a long time, and yet the magnetic patterns corresponding to a taped voice or the patterns of pigment on a page corresponding to a handwritten message have no meaning unless someone listens to the tape or reads the note; the meaning isn't waiting to be released from the tape or the ink. If language transmits anything, it is verbal behavior itself; in listening and reading, our own behavior re-creates some features of the behavior of the speakers and the writers who constitute our verbal community. We share our verbal behavior; it is, above all, social behavior.

Our first task is to classify verbal behavior. But our taxonomy must be functional rather than structural or grammatical. A grammatical classification of words in a sentence doesn't tell us about how it was produced or what consequences it had for the one who produced it. The same holds for a phonological classification of the sounds in someone's speech. Functional accounts of verbal behavior examine what verbal responses do. As with nonverbal behavior, structural and functional accounts of language can complement each other. Unfortunately, verbal behavior has been controversial in the history of psychology, and structural and functional accounts have often been pitted against each other as if they were incompatible instead of complementary (e.g., Catania, 1972; Catania, 1973b; Chomsky, 1959; Skinner, 1957). We'll try to deal consistently with both kinds of approaches.

Verbal responses are distinguished by the occasions on which they occur and the consequences they produce. They can be occasioned by either verbal or nonverbal stimuli, and they can have either verbal or nonverbal consequences. For example, a child might say "apple" in the presence of either an actual apple or the written word; as a consequence of saying "apple," the child might get either the apple or the confirming reply, "Yes, that's an apple."

In this chapter we'll consider some functions of verbal behavior, emphasizing various classes of verbal stimuli and responses. We'll begin with relatively simple relations in which verbal behavior is reproduced in either a spoken or a written mode (as in echoing what someone has said or taking dictation). In a later chapter we'll examine how verbal behavior makes contact with the environment, as when we describe objects or events; in doing so we'll discover something about what must be in place before we can learn to talk about ourselves. Those topics will set the stage for exploring more complex verbal processes, as in assertion and negation. Later we'll move on to some consequences of verbal behavior. In particular, **someone can change someone else's behavior by giving instructions.** In fact, **this may be the single most important function of language**. To see how it works, we have to understand the relations between verbal behavior and nonverbal behavior or, in other words, the relations between saying and doing. Our analyses of these functional properties of verbal behavior will prepare us to treat a few structural properties of verbal behavior.

CORRESPONDENCES BETWEEN SPOKEN AND WRITTEN CLASSES

Our verbal communities shape correspondences between things and their names, between words and their definitions, between what we did and what we said we did, between what we promised and what we accomplished, and so on and so on. The way we learn correspondences and the conditions that maintain them may determine how they work in our verbal behavior. Our first examples consist of some formal verbal relations between spoken and written classes, not because these are the most fundamental verbal classes but because they are familiar and involve correspondences well defined by our verbal community (Skinner, 1957). They will illustrate how to talk about words in terms of antecedents and consequences instead of in terms of meaning and reference. In other words, they will help us to talk about words as behavior.

The term *verbal* is general and applies to language in any modality; we distinguish it from the term *vocal*, which is specific to spoken language. We could easily extend our account to other modalities (e.g., the gestural modality of sign language or the tactile modality of Braille), but for convenience we'll mostly restrict our attention to spoken and written classes of verbal behavior. Also, it is sometimes worth distinguishing *spoken* from *vocal*, because not all voiced sounds qualify as speech and vice versa. But these technical issues will not affect our main account.

Correspondences between verbal stimuli and verbal responses in formal verbal relations are implicit in the colloquial vocabulary: We say that words are the same whether they're heard or spoken, seen or written, or, again in other words, whether they're auditory or visual stimuli or responses. One elementary verbal function is the reproduction of verbal behavior: We repeat what others say or copy what others write. Thus, our cases will include the reproduction of verbal behavior in all four possible combinations of spoken or written stimuli and spoken or written responses: spoken to spoken (echoic behavior), written to written (transcription), written to spoken (textual behavior), and spoken to written (dictation-taking).

It is easy to forget that most of these relations are arbitrary. **A spoken letter has no visual shape and a written letter has no sound.** The relations between the visual and spoken forms come about only after a very long history. For the present purposes, the most important point is that saying the letter sets the occasion for writing it down. Saying the letter cannot possibly look like the letter, but given an appropriate verbal history it can become a recipe for writing it down.

ECHOIC BEHAVIOR

Imitation of some properties of vocal stimuli appears relatively early in human infants' acquisition of speech. We call this class of verbal relations *echoic*. When a parent says "mama" and the child repeats "mama," the child's response is echoic to the extent that (1) it is occasioned by the parent's utterance and (2) the phonemes of the child's utterance have a one-to-one correspondence to those of the parents. Essentially, echoic behavior is generalized vocal imitation (cf. Poulson *et al.*, 1991).

Even though the stimulus and the response have common properties, this verbal relation isn't simple. Hearing a sound is not the same as making the vocal adjustments that produce the sound. The stimulus is a complex sound pattern. The response consists of the coordinated articulations of lungs, vocal chords, tongue, lips and so on. These produce sounds but are not themselves sounds. How then does the child know what to do to produce the sounds heard as "mama" or "dada" by the parents (cf. Chapter 19 on imitation)?

The significant dimensions of the units of speech called phonemes are more easily defined by articulations (positions of the tongue, etc.) than by acoustic properties (Lane, 1965; Liberman, 1982). The interactions of articulation and sound are complex; for example, many English consonants (e.g., *p*, *b*, *d*) cannot be produced unless accompanied by a vowel, and their acoustic properties vary as a function of context (e.g., the sounds of *l* and *k* are different in *lick* than in *kill*). **Echoic behavior is not defined by acoustic correspondence; it is defined by correspondences of phonetic units.**

Voices differ in many respects: An adult voice is deeper than that of a child, a woman's voice differs from that of a man, and people speak with varying regional dialects. If a young boy from a small town in New England repeats what a woman from Atlanta just said, their utterances will differ acoustically in many ways. But differences in vocal quality and regional dialect are irrelevant to whether the boy's behavior is echoic; the criterion for echoic behavior is the correspondence of verbal units such as phonemes and words.

That is why the duplication of human sound patterns by parrots and other birds doesn't qualify as echoic behavior: Their duplications are acoustic rather than phonetic. For example, a parrot would reproduce the *th* sound if a child lisps an *s*, but a nonlisping adult would ordinarily use the unlisped *s* instead. One reason why a voice synthesized by a computer may not sound realistic is that the computer doesn't produce its approximations to speech the way a human vocal tract does so.

Echoic units can vary in size from individual speech sounds to extended phrases or sentences. A variety of verbal phenomena, such as speech errors (e.g., Fromkin, 1971), can help us to decide what these units are. The importance of the echoic production of individual sounds is demonstrated by rhyming and alliteration in poetry (Skinner, 1972; Smith, 1968). For adult speakers, the units of echoic behavior are often whole words or phrases. The echoic production of extended phrases or sentences occurs in dramatics, as when an actor repeats the lines whispered by a prompter, and on ritual occasions, as when a bride and groom repeat the phrases of a marriage vow spoken by a member of the clergy. Echoic units are not defined by their size; they are defined by the correspondences into which they enter.

Echoic behavior doesn't simply accompany the acquisition of language and then vanish; it persists in the behavior of mature speakers. For example, I might repeat the name of someone to whom I've just been introduced. Nevertheless, echoic behavior doesn't imply that the speaker has understood what has been echoed; meaning doesn't enter into the definition of echoic behavior. As we'll see later, we must deal with meaning in verbal behavior in other ways.

Some verbal responses may be only partially echoic, as when I say something that incorporates just a couple of words from a question I've just been asked. Skinner took advantage of partial echoics in his design of a system to tap latent verbal behavior. His *verbal summator* (Skinner, 1936) was supposed to be a test of detecting speech in noise. Listeners wrote the words they heard, but the stimuli were actually repeated inflection patterns (e.g., *oo-ah-ee, oo-ah-ee, oo-ah-ee*) instead of words. The verbal responses Skinner collected echoed these inflection patterns, but they also reflected whatever verbal behavior was currently strong in the listener's verbal repertoire, much as the visual stimuli in Rorschach inkblots or Thematic Apperception Test materials do in personality testing (see also Estes, 1940).

The Development of Echoic Behavior

Echoic behavior depends at least in part on the shaping of articulations by their vocal consequences (cf. Risley, 1977; Skinner, 1957, p. 58). Before their own vocalizations begin to be differentiated, infants learn to discriminate among the many aspects of the speech of those around them (Eimas *et al.*, 1971); for example, infants readily discriminate between sounds of their native language and those from an unfamiliar foreign language, but not between sounds from two unfamiliar foreign languages (Mehler *et al.*, 1988). This means not only that they can also hear the differences when they make the sounds themselves, but also that some of these speech sounds may become reinforcers (Friedlander, 1968).

At first their babbling includes an undifferentiated range of human speech sounds, but over time they retain native speech sounds in their spontaneous vocalizations while other kinds of speech sounds gradually disappear. It was once thought that early infant sounds included all possible human speech sounds, but instead infants go through a progression from intonations and stress patterns to vowels and then to consonants, starting with the easiest ones (e.g., Vihman, 1996). Eventually their babbling evolves to self-repetitions (echolalic speech; e.g., *ma-ma-ma-ma-mà*) and then to repetitions of the speech of others (echoic speech).

Vocalizations can be reinforced (Poulson, 1984), and the vocalizations of infants are engendered and maintained by what they hear themselves saying; without these auditory consequences (as in cases of hearing impairment), the behavior doesn't develop. As we know from recordings, our own voices sound different to us than they do to others, because we hear our own voices via bone conduction as well as from sound that travels through the air.

Rare individuals may, through accidents of anatomy, hear themselves as others hear them. Those who do have the advantage of knowing how well they've actually imitated the voices of others. I first heard how different my own voice sounded to me than to others when a high school friend used a hidden tape recorder, then a novel device, to record a group conversation. After we'd listened to the tape one of our group said he didn't hear much of a difference. He was particularly good at doing vocal impressions. Given that he sounded roughly the same to himself as to others, if he talked so he sounded like Humphrey Bogart to himself then he'd sound like Humphrey Bogart to us too. I've occasionally asked in my classes about how people heard their own voices, but over the years I've encountered fewer than half a dozen who said they heard themselves as others heard them. Each of those, however, was pretty good at doing impressions. And perhaps two people could hear each other somewhat as they each hear themselves by pressing their heads together, so that both get some of the bone-conducted sound that is usually limited only to the speaker.

Perhaps native speech sounds become reinforcing relative to sounds of nonnative languages simply because they often accompany the activities of the infant's caregivers (e.g., DeCasper & Fifer, 1980). An articulation that produces something sounding more or less like what mommy says may be reinforced automatically by this correspondence between the infant's and the parent's utterances. The differentiation of phonemic structure may then be attributed to the various overlapping contingencies that different speech sounds enter into (cf. Chapter 12 on the origins of structure).

Categorical Perception of Phonemes

A child raised exclusively in an English-speaking environment will discriminate between the consonants *b* and *p*, but one raised exclusively in an Arabic-speaking environment, which does not include separate *b* and *p* consonant classes, will not (cf. habituation, Chapter 4). Furthermore, if the *b*-stimulus is gradually changed to a *p*-stimulus for the English speaker, the switch from calling the sound a *b* to calling it a *p* occurs discretely; listeners do not report an intermediate consonant that is a cross between them. This type of transition is called *categorical perception*, in that responses to stimuli varying along some continuum do not themselves vary continuously but instead fall into two or more distinct categories (Harnad, 2003).

Another fact about the acquisition of phonemes is that discriminations of speech sounds that are learned easily at an early age may be hard to learn later (e.g., Werker, 1989). For example, the distinction between spoken *r* and *l* in English doesn't exist in Japanese and is much more easily learned by a Japanese child than by an adult Japanese speaker. Infants master simple articulations before moving on to more complex ones. If the vocal units differentiated during babbling are different from one language to another, perhaps an adult learning a new language finds it difficult to master the new phonemes in part because the prerequisite simpler articulations were never differentiated.

The environment may be able to sculpt language units in various ways. Developing nervous systems divide into functional regions sometimes called compartments (cf. Chapter 3); of particular relevance is how boundaries form between compartments (e.g., Irvine & Rauskolb, 2001; Kiecker & Lumsden, 2009). Perhaps boundaries are also significant in the development of behavioral units, such as phonemes; they may be more important than particular localizations (Kim & Bao, 2008). For example, we know that cortical tactile receptor areas in primates include boundaries between the projection areas for different fingers that can be visualized with appropriate stains. Adjacent cells on one finger are likely to be stimulated at the same time, whereas those on a neighboring finger, though projected close by on the cortex, are likely often to be stimulated separately. These differences between correlated and uncorrelated firing may provide the basis for a boundary. Moreover, the number of compartments appears to be determined during development by the number of discrete units at the periphery (e.g., the separate rays in the nose of the star-nosed mole: K. C. Catania, & Kaas, 1997).

A long-standing puzzle is how exposure to a linguistic environment can lead to discrimination among native phonetic units long before children can produce those phonemes with appropriate vocalizations. Discrimination comes long before differentiation, but what is the discriminated response if corresponding vocalizations have not yet been shaped? If there are any spatial projections from the peripheral auditory system to auditory cortex that correspond to phonetic boundaries, a similar compartmental system may perhaps lead to separate areas for *b* and *p* for a child in an English-speaking environment but none for one in an Arabic-speaking environment. The categorical perception that emerged from such a process would certainly depend on the brain, but if it worked like this it could hardly be understood without including the role of acoustic environments. Language units created in this way could then interact with the environment through other functional contingencies.

Behavior analysts do not ordinarily appeal to physiology. When I was an undergraduate, one of my teachers, W. N. Schoenfeld, often took extreme positions to challenge our assumptions (Catania, 1997b). One was that the primary function of the brain was to cool the blood. When the class tried to counter his claims, he pointed out that none of us had ever seen a brain do anything. He then segued into making a case for the crucial role of behavior in our understanding of anatomy and physiology. He used the example of vision, arguing that we would not be much interested in the differences between retinal rods and cones were it not for their relation to important differences in seeing. In so doing, he engendered a healthy skepticism about explanations that appealed to anatomy but did not take into account the roles of behavior and the environment. But he did not deny the physiology, and I suspect he would have appreciated accounts that started from behavior and showed how physiology might depend on behavior and its environments instead of the other way around.

This has been an account of the natural ontogenic selection of the phonetic units of verbal behavior at an ontogenic level. The selection is based upon correspondences between the language already available in a verbal community and the vocalizations of an infant who is the language learner, so it also provides for the transmission of verbal behavior at the cultural level (in this case from one generation of speakers to the next). There are ways it can fail: motor disorders may constrain articulation; hearing may be impaired; relevant consequences may be absent, perhaps as a result of neglect or abuse by caregivers or because neurological deficits have reduced the effectiveness of social stimuli. But once vocal articulations can be shaped and the child can hear the consequences of those articulations, the minimum conditions for the ontogenic selection of vocal units are in place. In contrast, reinforcement of nonhuman vocalizations is typically limited to rate of vocalization; it is difficult to modify its topography (e.g., Hayes & Hayes, 1951; Lane, 1961). A lot may follow from the shaping of speech: "The human species took a crucial step forward when its vocal musculature came under operant control in the production of speech sounds. Indeed, it is possible that all the distinctive achievements of the species can be traced to that one genetic change" (Skinner, 1986, p. 117).

TRANSCRIPTION

Verbal stimuli and responses can correspond when both are written. In such cases, the behavior is called *transcription*. For example, you might copy a number from a telephone book or an author and title in preparing a bibliography. It doesn't matter whether you produce the copy by writing it or typing it or dragging and dropping it. Just as we distinguish vocal articulations and the sounds they produce in echoic behavior, we also distinguish the movements involved in producing words from the looks of the words in transcription. And just as echoic behavior depends on correspondences of verbal rather than acoustic properties, transcription depends on correspondences of verbal rather than visual properties. A handwritten sentence may look very different from the printed text from which it was transcribed (for example, the script letters run together but the printed ones do not). Nevertheless, writing the sentence qualifies as transcription if the script sentence matches the printed one in spelling, word order and punctuation.

Just as units of echoic behavior can vary from individual phonemes to entire phrases or sentences, units of transcription can vary from individual characters to extended passages, depending on the circumstances in which the behavior occurs. A child learns to copy single letters before learning to copy whole words. In doing so, the child learns the correspondences between arbitrary visual forms, such as the printed and script *a* in upper- and lowercase. There may be no visual property common to all forms of the letter *a* (cf. Gibson, 1965). Consider the several different appearances of each form of the following individual English letters:

A a 𝒶 𝒜   B b ℬ   D d 𝒟 δ   E e ℰ €   F ſ ℱ ƒ f
G g ℊ 𝒢   Q q 𝒬   R r ʌ ℛ   S 𝓈 s 𝓈   T t 𝔗 † 𝒯

These differences within the forms of a single letter are sometimes greater than those between different letters, as in the following sample groups:

h n m   O D Q   𝑒𝓁   P R B
M N V W   I l L T   b d p q   𝓂𝓃𝓊𝓋𝓌

The variations of form explain why transcription differs from copying in the pictorial sense (cf. Shahn, 1972, pp. 49, 256). A skilled Asian calligrapher unfamiliar with the European alphabet in which a text is printed might produce an accurate copy, but the copying would not be verbal. This distinction is based on the behavioral units in the two kinds of copying. The critical features of the calligrapher's copying are geometrical properties of the letters in the text and the marks produced by the calligrapher's strokes; the critical features of transcription are the verbal units (letters, words and phrases) in the original text and its copy. Visually, the calligrapher's copy might look more like the original than a handwritten copy by a speaker of the language, but only the latter counts as transcription. An 𝒜 may be transcribed as an A or an *a* or an a, but a mere picture of the 𝒜 doesn't qualify.

Except for their respective spoken and written modes, echoic behavior and transcription are formally similar. Children ordinarily acquire echoic behavior early, even without specific instruction, but they take some time doing so. They don't ordinarily acquire transcription unless it is explicitly taught. Nevertheless, the relative ease with which one or the other is learned provides no grounds for assuming that one is simpler than the other.

Pure transcription is probably rare. It occurs only when transcription is unaccompanied by other behavior, such as the responding to a text that we call reading for meaning. For example, if a skilled typist accurately transcribes a letter while not responding verbally to it in other ways (e.g., while listening to a conversation elsewhere in the office), the typist afterwards may be unable to describe what was in the letter. That would be a case of relatively pure transcription. Responding to the text in other ways, as in reading or understanding it, is doing something else that may combine with the transcription but does not count as part of it. When we're transcribing, we're usually doing a lot more besides. As with echoic behavior, meaning doesn't enter into our definition; it must be dealt with in other ways.

TEXTUAL BEHAVIOR

When a written verbal stimulus sets the occasion for a corresponding spoken verbal response, the behavior is *textual*. Thus, you might say aloud what's on a menu or read a bedtime story to a child. In textual behavior, the arbitrary correspondence between verbal stimuli and responses is more obvious than in either echoic behavior or transcription, because the stimuli and the responses are in different modes. A written word is a visual stimulus; it has no sound. A spoken word is an auditory stimulus; it has no shape. Yet these correspondences are so familiar that we rarely notice the arbitrary nature of the relations between verbal shapes and sounds.

As with transcription, textual behavior is usually taught explicitly, and some controversies about teaching it are based on assumptions about the units appropriate to various stages of instruction,, such as whether the teaching of reading should begin with individual letters, syllables or whole words (Gleitman & Rozin, 1973).

As with the other formal classes, we must distinguish textual behavior from other kinds of responses to written verbal stimuli. For example, if a sign says STOP, reading the word aloud is textual but stopping is not. In the mature reader, textual responses become less important than other kinds of responses to written verbal stimuli. Vocal responses diminish in magnitude, become subvocal and perhaps disappear completely as a child becomes a proficient reader. Reading is behavior, but textual responses are at best only one part of reading. For example, a father reading a bedtime story aloud to a child might finish a page and suddenly realize that he doesn't know what just happened in the story, even though the child he has been reading to does.

This example, without understanding, is a pure case of textual behavior. The colloquial vocabulary doesn't distinguish between reading that is simply the saying of the words on a page and the kind of behavior we call reading for understanding (e.g., Fowler *et al.*, 1981). Most of us have occasionally found ourselves partway through a page unable to say what was in the last couple of paragraphs. When that happened we were presumably engaging only in some reduced form of textual behavior and absolutely nothing else (of course I assume that never happens when you are reading this book – but would a smiley emoticon have helped here?). Such experiences are evidence for the importance of these distinctions (see also Kolers, 1985). Reading for understanding includes other behavior along with or instead of vocal or subvocal speech, so it is more (probably a lot more) than simply textual behavior.

DICTATION-TAKING

Just as a written stimulus can set the occasion for a spoken response, a spoken stimulus can set the occasion for a written response. This class of verbal behavior is called *dictation-taking* (we're concerned with the listener who takes the dictation, not the speaker who dictates). For example, you might write down an email address just given you or take notes at a lecture. The units of dictation-taking are typically entire words or phrases, but individual letters may also serve (e.g., as when children are taught the written alphabet or an unusual name is spelled out for a stenographer).

As with textual behavior, dictation-taking involves stimuli and responses in different modes. Some of its special properties follow from the relatively permanent record produced in the written text. In addition, occasions for dictation-taking are limited relative to textual behavior, because, unlike our vocal apparatus, writing implements are not parts of our anatomy. Perhaps for this reason, we aren't much tempted to pursue its covert manifestations; we're less likely to speak of submanual writing or typing than of subvocal reading. The motor theory of consciousness had argued that thought was just vocal behavior reduced in magnitude (Max, 1934). Nevertheless, textual behavior and dictation-taking are formally similar, and either can be accompanied by other kinds of verbal behavior occasioned by verbal stimuli.

RELATIONS AMONG THE CLASSES

We've limited this account of formal verbal classes to spoken and written stimuli and responses. We could have extended it to other language modes (e.g., sending and receiving Morse code). Sometimes we can ignore the distinction between verbal stimuli and verbal responses, but not always. For example, even though we speak of letters and words without regard to whether they're written or spoken, the mode matters in teaching. A teacher who has only taught a child to name letters written on a chalkboard shouldn't expect the child to be able to write the letters when they're spoken. The stimuli in the former are the responses in the latter, and vice versa. Until they become equivalent, they're just one part of "knowing the alphabet."

To teach reading and writing is to teach equivalences between spoken and written modes of verbal behavior (cf. Chapter 12). Our everyday vocabulary obscures these distinctions, perhaps because equivalences come easily to us, or perhaps because language instruction is designed to eliminate them, or perhaps both. In any case, the relations are summarized in Figure 20-1.

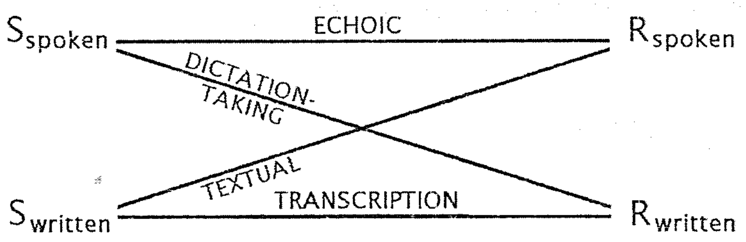

Figure 20-1 Relations among the four formal classes (S, stimulus; R, response). Similar relations can be established for any pair of verbal modes (e.g., printed text and the gestures of American Sign Language, or spoken behavior and Braille). Note that the stimuli here are discriminative; here and in figures to come the superscript D has been faded out.

The Replication of Verbal Behavior

All of the formal classes involve the replication of verbal behavior in either spoken or written form. Darwin discussed parallels between the evolution of species and the evolution of languages (Darwin, 1871); under the pressures of modernization some languages have become extinct and many others are threatened. Organisms replicate some features of themselves and their species in phylogenic selection. Similarly, the replication of features of verbal behavior is critical to the maintenance of human languages.

The replication of verbal behavior is not mere reproduction. Based on a long history in which we've learned relations among spoken and written stimuli and responses, we say that words are the same whether spoken or written. Thus, saying what has been written or writing what has been said are instances of replication even though they carry across different modalities. In this way the effects of verbal stimuli become extended over time and space as verbal behavior produced by some is passed on to others.

Once some individuals repeat what they or others say, verbal behavior can be maintained by cultural contingencies and survive across generations. Some effects of replication are fairly straightforward. For example, if a single utterance has no effect repetitions may produce one. Furthermore, repetitions allow instructions to be followed in the original speaker's absence, later and elsewhere, as when we repeat to ourselves the details of a task someone has asked us to complete. We'll consider some further implications in Chapter 24. Replication must come first, but once in place powerful contingencies maintain it.

Parallels in Music

Thus far we've emphasized classes defined in terms of verbal modes, but other formal relations may be differentiated even within a mode. For example, if a student has learned only to translate from English to German, the student may have difficulty when asked to translate in the other direction, from German to English (cf. Chapter 26 on transfer). The problems are compounded when both spoken and written languages are involved. Language instruction usually recognizes these distinctions; a course in conversational French is expected to emphasize the spoken mode, whereas one in scientific Russian is expected to emphasize the written mode. Bilingualism might best be defined by the extent to which equivalence classes extend across both languages (cf. Caramazza & Brones, 1980; Kolers, 1966).

Each mode, spoken and written, has special characteristics. For example, spoken verbal behavior varies more freely in stress, rhythm and intensity than written verbal behavior, but it is also more transient. Some properties of verbal behavior are independent of mode. For the immediate consequence of being able to dial a phone number, it hardly matters whether you hear the number or see it on a cellphone screen, or whether you repeat it aloud to yourself or write it down. Some consequences may affect how you look for the number next time (e.g., how hard it was to find it or whether you were able to retrieve it later), but they are not essentially verbal. Thus, once equivalences are in place it sometimes may be relatively unimportant to distinguish among the formal classes.

But we may need the distinctions for studying certain topics, such as the language pathologies called *aphasias*, in which deficits are characterized by the verbal classes affected (e.g., Sidman, 1971). Some aphasias affect only one of the four possible relations illustrated in Figure 20-1, so we cannot always assume equivalences. But we don't often run into aphasias. Music provides a more familiar example. In music, the mastery of a subset of the possible formal relations among musical stimuli and musical responses does not guarantee that the others will be carried along.

Figure 20-2 substitutes singing and reading or writing music for the verbal classes of Figure 20-1. The word *behavior* has the same stress patterns as the word *Maria* in the song by that name in Leonard Bernstein's musical, *West Side Story*, so it has been used here as an illustration. This formal correspondence was the inspiration for a musical take-

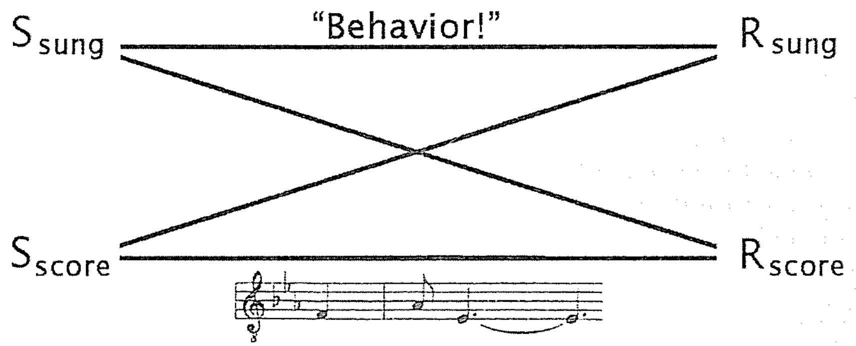

Figure 20-2 Relations among four formal classes of musical behavior. They involve singing given a heard song or a musical score, or writing the music down given a heard song or a musical score (S, stimulus; R, response). The example is the word "Behavior!" sung to the melody of the song "Maria" from the musical *West Side Story*.

off about the conversion of young student named Tony; in the course of the story he is persuaded to move from a mentalistic discipline to the behavior side. The show has occasionally been performed at professional meetings (Catania, 2012).

Some musicians play or sing by ear and others prefer to do so from a score. Some musicians can sing or play a heard melody but many who can do so cannot write the music down. The musical behavior analogous to echoic behavior is common among musicians, but they are variable with respect to the analogs of dictation-taking, transcription and textual behavior, with some proficient in all three and others proficient in only one or two.

Some talented musicians manage without any musical scores at all on either the stimulus or the response side, so with regard to reproducing music they are limited to echoic music making. If they had been correspondingly limited with regard to spoken and written verbal behavior they'd have been called illiterate, but the honed skills of sophisticated musicians together with contemporary recording technologies so outweigh the formal limits on how they can replicate music that distinctions like those we've been considering here simply don't matter to their audiences.

In verbal behavior, a speaker's response is a listener's stimulus and a writer's response is a reader's stimulus, and a speaker or writer at one time becomes a listener or reader at another. Music is a valuable supplement to our other verbal examples because it is more obvious in these musical cases that mastery of one or more of the formal relations does not guarantee the integrity of the others. They are separate classes of behavior that become integrated only given appropriate contingencies.

Chapter 21
Verbal Function:
Intraverbals and Verbal Learning

Intraverbal *breaks down into* intra-, *between, plus* –verbal, *word. Verbal learning procedures used in the study of intraverbals include serial and paired associates learning, free recall and the special case of verbal discrimination called verbal recognition. Most of these terms have been related to words considered elsewhere:* associates *to* consequence*;* recall *to* class*;* recognition *to* cognition*; and* transfer *to* differentiation. Serial, *through the Latin* serere, *to arrange or attach, is related to* series, sort *and, perhaps as an arrangement of words,* sermon.

Thematic Intraverbal Classes
Intraverbal Chains
Intraverbal Chunks
Verbal Learning
Addendum A: Serial Learning
Addendum B: Paired Associates Learning
Addendum C: Free Recall
Addendum D: Verbal Discrimination and Verbal Recognition

The formal classes involve verbal responses occasioned by verbal stimuli. They are each characterized by one-to-one correspondences of verbal units. In transcription, for example, each word of a text has a unique equivalent in the transcribed version. But we learn many verbal relations that do not involve such formal correspondences, and these relations interact in powerful ways, as when combinations of basic arithmetic intraverbals such as *two plus two equals four* or *three times five equals fifteen* combine in our solution of a mathematics problem. Skinner called such relations *intraverbal* (Skinner, 1957). They come in three different varieties: through membership in common thematic or semantic classes, as when hearing *cow* and *barn* makes us just a bit more likely to say *farm*; through effects of verbal stimuli that set the occasion for specific verbal responses, as when we're likely to say *blue* if someone says *red, white and...*; and through the temporal integration of verbal units into larger sequences, as when we say *One picture is worth a thousand words* or *A word to the wise is sufficient* (and more intraverbal behavior is involved in deriving the implication that *A picture to the wise is redundant*).

THEMATIC INTRAVERBAL CLASSES

"The intraverbal relations in any adult repertoire are the result of hundreds of thousands of rein-

KEY TERMS: Intraverbals; Formal and Thematic Intraverbal Relations; Serial Learning, Paired-Associates Learning, Free Recall, Verbal Discrimination and Verbal Recognition.

forcements under a great variety of inconsistent and often conflicting contingencies. Many different responses are brought under the control of a given stimulus word, and many different stimulus words are placed in control of a single response" (Skinner, 1957, p. 74). For Skinner here, reinforcements are not ordinarily edibles or other items but rather the simple consequences of day-to-day interactions, such as a nod or a change in expression or appropriate action; many of these consequences are verbal, such as acknowledgments or other replies. Reading a word in a particular way because it occurred in a certain context might sometimes qualify too. All sorts of small-scale consequences that follow from listening or reading can accumulate over time; these are Skinner's *hundreds of thousands of reinforcements*, and they can lead to many subtle interactions among verbal responses.

Many groups of words tend to be encountered in common situations, and when they do so they come to vary together as members of a verbal class. Such classes are said to be thematically or semantically related. For example, transportation words include *car, bus, plane, train, vehicle*, as well as groups within groups, such as *automobile, taxi, sedan* within the class *car*, and these classes can be even further differentiated by makes and models. The complex interlocking of such classes is evident in literature and poetry, and many verbal games appeal to thematic classes (e.g., Botticelli, Wordsmith, MadLibs). The discussion of free association below provides further examples of thematic relations. Thematic classes are the counterparts, for tacting, of the formal classes we discussed in Chapter 20. A verbal response may be occasioned by combined thematic and formal properties of a verbal stimulus. For example, the formal heard-word rhyme may provide a slight echoic contribution toward my saying "I heard that word" instead of "I heard what you said." Thematic and formal classes are ubiquitous, and the available examples are as vast as our extended human verbal environments.

We might say that thematic classes consist of words with common meanings, but to do so we would have to define the behavioral significance of meanings. Meanings aren't properties of words themselves; they are properties of our responses to them. For example, if you repeat a familiar word like your own name over and over, you may find that it loses its meaning; this implies that some response to it drops out after several repetitions. A red traffic light means *stop* and a green one means *go*, but when I stop on red and go on green, I'm responding to red and green and not to their meanings. Written words in a language that no one understands have no meaning, and when we say a word has many meanings (as in *bat* in a belfry and *bat* at a baseball game), its meaning changes only in the sense that we respond differently to it in its different contexts.

In these examples, words have functioned as both stimuli and responses, so they cannot be identified with just the response terms of operants or just the stimulus terms of discriminated operants. Meaning is sometimes a convenient term to refer to verbal classes that function as either stimuli or responses within verbal contingencies.

INTRAVERBAL CHAINS

We learn to recite the alphabet, to count, to give answers to arithmetic problems, to recite poems, to define terms and to state facts. The sequence of letters in the alphabet is no more orderly than the one on a computer keyboard: Some similar letters of the alphabet are close together and others are widely separated (e.g., *M, N* but *D, T*), some voiced consonants precede and some follow the voiceless (e.g., *B, P* but *F, V*), and so on. Despite its arbitrary character, we learn the alphabet because so much is ordered according to it: dictionaries, telephone books, indexes. We're less able to recite the order of letters on a computer keyboard because we do not have to behave with respect to keyboards the way we do with respect to alphabetized lists.

The same points apply to chronologies, geographies and much else of our everyday knowledge. No one now living ever saw Washington crossing the Delaware. I might argue that I saw the scene in a painting, but even if I'd recognized Washington, could I have known that the setting was the Delaware River without a label that told me so? We don't ordinarily learn historical details by experiencing them. Instead, given names or dates, we learn to say when or in what order events occurred. "Most of the 'facts' of history are acquired and retained as intraverbal responses" (Skinner, 1957, p. 72).

In the chaining of intraverbal behavior, one verbal stimulus sets the occasion for another verbal response. The relation between stimulus and response is arbitrary; there are no systematic correspondences between them. Free association is an example (Galton, 1879). The immediate consequences of free associations are usually minimal, and any given verbal stimulus may occasion a variety of different responses, so like the verbal summator the procedure is assumed to tap verbal responses of relatively high probability in the speaker's verbal repertory. In discussing intraverbal behavior, Skinner (1957) treats free association as follows:

> One verbal response supplies the stimulus for another in a long series. The net effect is revealed in the classical word-association experiment. Here the subject is simply asked to respond verbally to a verbal stimulus, or to report aloud any responses he may "think of."... Such an experiment, repeated on many subjects or on one subject many times, produces a fair sample of the responses under the control of a standard stimulus in a given verbal community.... Many different responses are brought under the control of a given stimulus word, and many different stimulus words are placed in control of a single response. For example, educational reinforcement sets up many different intraverbal operants involving the cardinal numbers. *Four* is part of the occasion for *five* in learning to count, for *six* in learning to count by twos, for *one* in learning the value of π, and so on. On the other hand, many different verbal stimuli come to control the response *four*, e.g., *one, two, three...* or *two times two makes....* Many different connections between verbal responses and verbal stimuli are established when different passages are memorized and different "facts" acquired. The word-association experiment shows the results. (Skinner, 1957, pp. 73–74)

Word associations were the basis for some attempts to measure meaning (Galton 1879). If listeners produced longer and more varied lists in response to one word than another, the first word was said to be more meaningful than the second. It was assumed that words closely related in meaning would occasion overlapping lists of associates. For example, common associations to both *infant* and *baby* might include *crib* and *bottle* but probably none of these would be responses to *guitar*. The different degrees of overlap among the associates are consistent with what we already know: *Infant* and *baby* are closer in meaning than either is to *guitar*.

Word associations can be based on several words as stimuli instead of single words. Such sequences are the basis for constructing different *orders of approximation* to English. Suppose I give you the first three words of a sentence and ask you to complete it by giving me the next word. Now I have four words. I drop the first one and present the remaining three words to someone else, who now adds one more word. Again I drop the first one and present the remaining three to still another person, and so on. An example of text produced this way is *the first list was posted on the bulletin he brought home a turkey will die on my rug is deep with snow and sleet are destructive and playful students always* (Miller & Selfridge, 1950). This is a fourth-order approximation to English, because at each point a participant, seeing only the last three words, must add the new fourth one. A zero-order approximation has words chosen randomly; a first-order approximation has words chosen randomly but in proportion to their frequencies in the language. The more closely a text approximates English the easier it is to remember, but even with high-order approximations grammatical sequences are hardly ever generated.

In its simplest forms, intraverbal behavior has been the focus of much research on human verbal learning, perhaps because of the ease with which verbal materials can be manipulated as stimuli. The classic experiments of Ebbinghaus (1885/1964) were specifically concerned with the learning of arbitrary verbal combinations. Paired-associates learning (learning word pairs) and serial learning (learning ordered lists, as in learning to count) represent fairly basic cases of intraverbal behavior.

Intraverbal relations are an important component of standard educational practice, as when a child who is mastering the multiplication table gives 42 as the response to 6 x 7. But the response is strictly intraverbal only if it doesn't depend on other intervening arithmetic behavior (e.g., adding six sevens, counting by sixes or counting the boxes in a 6-by-7 rectangle). We'll return briefly to cases in which the answer is derived rather than learned as an intraverbal chain when we've considered some environmental determinants of verbal behavior and some of its consequences.

INTRAVERBAL CHUNKS

It is appropriate to speak of intraverbal chains only when successive parts of an utterance serve as discriminative stimuli for later parts. When extended utterances function as independent verbal units, as in the rapidly emitted sequences discussed in Chapter 10, it is inappropriate to say that these sequences are intraverbal chains. In earlier editions of this book I called such integrated sequences temporarily extended units, but this was an unwieldy label. Here, by analogy with similar units in the analysis of remembering, I call them intraverbal *chunks* (cf. chunking in Chapters 26 and 27). Such maxims as *Haste makes waste* and *He who hesitates is lost*, best regarded as verbal units in their own right, are therefore intraverbal chunks.

Some examples of large intraverbal chunks are Hamlet's soliloquy, the Pledge of Allegiance, and the lyrics of Simon and Garfunkel's *Sounds of Silence*, assuming in each case that the words have become very well learned. Some short verbal sequences can be learned very quickly, but when they are as extended as these they may take some time. At the beginning, completion of the sequence will depend a lot on chaining, perhaps in combination with prompting or with reference to a written copy. But even early on, some parts begin to hang together. When I get as far as *To be or not to* I don't need to consider what I've said so far before I complete the line with *be*. Whether the completion of that line occasions *That is the question* will tell me whether this part of the soliloquy is still chained or I've begun to integrate it into a larger chunk. Throughout my learning, the soliloquy will be a mix, partly chained and partly chunked. And even for a seasoned Shakespearean actor, much less me, the soliloquy may remain a segmented work where major boundaries separate chunks that are chained together. Ay, there's the rub.

VERBAL LEARNING PROCEDURES

The area traditionally called verbal learning is concerned with what happens as we learn word sequences, word combinations and word contexts. Its literature is extensive. Here we will briefly examine four major classes of verbal-learning procedures: serial learning, paired-associates learning, free recall and verbal discrimination, including verbal recognition as a special case). We'll illustrate each class of procedures with a sample of some characteristic findings.

In a typical verbal learning procedure, we present verbal stimuli to the learner and then record verbal responses. We take a lot for granted in such arrangements. Describing them simply in terms of verbal discriminative stimuli and differentiated verbal responses omits significant features. For example, the learner's performance is only rarely generated by differential consequences; instead, it's usually generated by giving instructions. The consequences may be hard to specify. The experimenter designates responses as correct and incorrect or right and wrong. When the learner responds, the experimenter may say "Right" or "Wrong" or the learner may simply be shown the item designated as correct. These procedures are sometimes said to provide the learner with *feedback* or *knowledge of results*. It is tempting to assume that telling or showing a learner that a response was correct is a reinforcer, but such feedback might function only as an instruction (telling the learner that a response was correct is like saying, "Respond the same way next time"). Thus, it could be misleading to speak of the reinforcing properties of being correct and the punishing properties of being incorrect.

Other circumstances may affect the learner's behavior. For example, if an experiment ends as soon as a list is learned, the consequences of finishing early will differ for a learner satisfying a course requirement than for one whose pay depends on time spent learning the list. These and other experimental details that determine what the learner does are sometimes called *demand characteristics* (Orne, 1962). Here we won't have much to say about the consequences of verbal learning, but even when unspecified we can usually assume that they are fairly uniform throughout an experiment.

Verbal learning materials have ranged from simple items such as numbers and letters through nonsense syllables and words to complex materials such as sentences and extended texts. Within these categories we must distinguish between *nominal stimuli* and *functional stimuli* (cf. the nominal or descriptive and the functional classes in Chapter 9). Nominal stimuli are the verbal items as defined by the experimenter; functional stimuli are the features that occasion the learner's response. For example, if the items in a learning task are the three-letter sequences OED, EAB and PBK but the learner attends only to the first letter of each, then the three-letter sequences are the nominal stimuli but the first letters are the functional stimuli.

Verbal learning experiments use verbal stimuli in both written and vocal modes, but written stimuli dominated in the early days of this research area. The advantages of written stimuli are that they are uniform, easily described and can be presented at well-defined rates and durations. An early device for presenting verbal stimuli was the memory drum, illustrated in Figure 21-1. Words were typed or printed on a paper loop and at a controlled rate a roller advanced each item to a window where it became visible. In contemporary research, computer displays have displaced the memory drum. In early research, vocal verbal stimuli were less favored because it was harder to control their uniformity and rate of presentation; an experimenter who read a list to different learners might change the inflection or loudness of words over successive readings. Tape recording eliminated this problem, and contemporary research uses either written or vocal verbal stimuli, depending on the suitability of one or the other for a particular experiment. Similarly, the choice of written or vocal responses varies with experimental requirements. Learners can usually respond more quickly vocally than in writing, but the vocal response must be recorded whereas the written response is its own record.

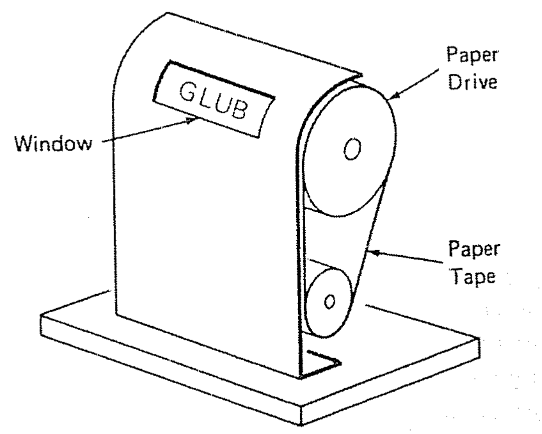

Figure 21-1 A memory drum. Words on a paper tape are shown in the window, each for a fixed time. The learner might be instructed to anticipate each item before it appears or to recite the whole list at the end of its presentation.

Verbal-learning procedures experimentally realized the associationist principles developed by such philosophers as David Hume and James Mill. The associationists had advocated that human thought was based upon the association of ideas. Ideas were said to become connected or associated in various ways (e.g., through similarity, common elements, contrast), but especially through contiguity in space or time. Later, with the beginnings of modern chemistry, analogies were drawn between the formation of associations and the chemical combinations of atoms into molecules. Hermann Ebbinghaus, a German investigator, was in effect the founder of research on remembering, using himself as the experimental subject. He saw the possibility of measuring the formation of arbitrary associations (Ebbinghaus, 1885/1964), and invented the nonsense syllable as an item that had not yet acquired verbal functions and therefore would not be contaminated by existing associations.

Table 21-1 summarizes the major classes of verbal learning procedures. As its name implies, *serial learning* is the learning of a sequence of verbal items, as when a child learns to count or to recite the days of the week. A sequence can be learned through *serial recall* or *serial anticipation*. For example, suppose we try to teach someone the first 10 presidents of the United States through serial recall. We show each name briefly and the learner tries to recite the entire list in proper order only after all 10 have appeared; we repeat the list until the learner gives one or more correct repetitions. Alternatively, in serial anticipation, as each name appears the learner tries to say which will appear next (e.g., on seeing *Jefferson* saying "Madison" and then on seeing *Madison* saying "Monroe" and so on); we also repeat this list until the learner gives some number of correct repetitions. Once the list is learned, we might ask the learner questions about the positions of names in the list, such as "Who was the ninth president?" (Harrison), "Which president was Tyler?" (tenth) or "Which president came before Van Buren?" (Jackson). A well-learned serial list, the alphabet, illustrates the roles of context and position. We can all recite the alphabet and can probably quickly say which letter precedes or follows any other, but without counting how easy is it to give the eleventh letter or the numerical position of *S*?

Paired-associates learning generates correspondences between items on two lists, as in learning the dates of historical events or the equivalents in one language of words in another language. Each verbal stimulus sets the occasion for a verbal response. For example, if the pairs to be learned are titles and authors, the sequence might be as follows:

Table 21-1 Classes of verbal learning procedures

Name	*Description*	*Examples*
Serial Learning	Items in a verbal sequence are learned in order.	Reciting the alphabet or months of the year; learning a poem
Pair-associates learning	Each item in a list occasions a different verbal response; list order may vary.	Given a country, naming its capital; given words in English, providing equivalents in another language
Free recall	Items in a list are named without regard to order.	Naming the teams in a sports league; naming players on a team
Vebal discimination	Responses are occasioned by the classes within which verbal stimuli fall; in other words, the discriminated dimension is verbal.	Identifying nouns or verbs in a text; sorting name cards into two stacks, one of women's names and one of men's
Verbal recognition	A special case of verbal discrimination. The property defining the discriminated class is whether the item appeared in some earlier list or context.	Picking out names of people you know from a list; distinguishing between old and new terms while reading a text

Stimulus Item	*Feedback*
Don Quixote	Don Quixote–Cervantes
Candide	Candide–Voltaire
Moby Dick	Moby Dick–Melville
Faust	Faust–Goethe
The Odyssey	The Odyssey–Homer

Each title appears alone for a few seconds. Whatever the learner's response, the author is then added. The full set is repeated, perhaps with the order changed, until the learner correctly names each author before the name appears. Questions can then be asked about effects of changing items or reversing stimuli and responses. Paired-associates experiments, however, don't typically use familiar items and pairings; they are more likely to use nonsense syllables or arbitrary combinations.

In *free recall*, the learner is asked to name, in any order, items presented earlier, as when you're asked to say what was on a misplaced shopping list or what questions were on a test. In verbal learning studies, the list consists of verbal items, but the procedure is like recall of nonverbal stimuli, as when a witness names those present at the scene of a crime. Whether the stimuli are verbal or nonverbal, the responses in both cases are verbal.

Verbal discrimination is simply discrimination along some verbal dimension of stimuli, as when a child learns to distinguish vowels from consonants or nouns from verbs or grammatical sentences from ungrammatical sentences. Pairs of items are presented, and the learner chooses one item of each pair by naming it, pointing to it or perhaps pressing a left or a right button.

One type of verbal discrimination is that between new and old items, as when a student in a language course distinguishes words already learned from new ones. This is called *verbal recognition*. The learner is given an initial list. Later, the learner is given another list with both new items and items from the first list. The learner must identify (recognize) the items from the first list. Except that it uses verbal materials, verbal recognition is similar to familiar instances of recognition (e.g., as in recognizing friends in a group of people or places you've visited.)

Each of these procedures seems superficially simple, but while mastering any of these tasks, the human learner inevitably does a lot more. Thus, it is inappropriate to regard what has been learned as mere verbal sequences; human verbal behavior has other structural properties besides the temporal ordering of events (e.g., the thematic relations among its components). The transition from serial and paired-associates learning to free recall occurred in part to allow the learner to demonstrate some of the other things learned about lists besides the orders of specific items in them.

ADDENDUM 21A: SERIAL LEARNING

To study serial learning, Ebbinghaus constructed nonsense syllables so that learners would be mastering material with which they were not already familiar. His typical nonsense syllable was a three-letter consonant-vowel-consonant sequence (a *CVC trigram*). He excluded trigrams that were already words (e.g., DOG and CAT are CVC trigrams, but in English they are not nonsense syllables). He then arbitrarily made up lists to be learned later. He learned different lists over a period of years, recording such data as the number of repetitions he took to reproduce a list without error. For a list of up to 7 trigrams, he required only a single reading for a correct reproduction; beyond that point, the repetitions needed for a correct repetition increased with list length, up to about 55 repetitions for a 36-trigram list. Not only total learning time but also the learning time per item increased with list length. Another of Ebbinghaus's findings was that he could learn meaningful material more rapidly than nonsense syllables. In contrast to the 55 repetitions required for a 36-item list of nonsense syllables, Ebbinghaus learned 80-syllable stanzas of poetry in about 8 repetitions.

Together with other verbal-learning procedures, serial learning provided a baseline for studying factors that influence human learning. For example, the finding that spaced practice is typically more effective than massed practice (e.g., Underwood, 1961) is often cited in support of distributing study evenly throughout a semester instead of cramming it in at the end, just before exams. The finding is so familiar that it is surprising that it was once regarded as counterintuitive. The superiority of spaced over massed practice eventually contributed to accounts of verbal learning in terms of interference between items.

One prominent feature of serial learning was the serial-position effect (Robinson & Brown, 1926). In a list, items at the beginning are usually learned most easily, followed by items at the end, and finally by items in the middle. The serial-position effect raised questions about the role of position in the list. Serial learning was regarded as a procedure in which each item is the stimulus that occasions the next. But learners can also answer questions about position (e.g., "What was the next-to-last item?").

Ebbinghaus raised other questions about serial learning. Could associations be formed only between successive items, or were remote associations possible (e.g., between every other or every third item)? Were associations formed both forward and backward through the list? Ebbinghaus examined these issues by relearning lists of nonsense syllables one day after original learning with item order the same as or different from that of the original list. This was called the method of *savings*: The list was ordinarily learned more quickly on the second day than the first, and effects of changing the order of items were measured by the time saved on relearning. Ebbinghaus studied savings with 16-item lists in original, scrambled and reverse orders, and with lists of every second, third, fourth or eighth item (e.g., in a list based on every third item, the order *A B C D E F G H* becomes *A D G B E H C F*). As shown in Figure 21-2, the greatest savings in relearning occurred with items in the original order and the least with items in scrambled order. The reversed list produced more savings than any of the remaining ordered lists. From these data, Ebbinghaus concluded that both remote and backward associations had been formed during the original learning of each list.

If so many different kinds of associations can emerge within a serial list, serial learning might not be the way to study associations. It might be better to examine pairs of associates independently of a serial order. Serial learning could then be treated as a special case of paired-associates learning. For example, consider the five-item serial list, NAJ BEF LUJ PES CED. In the presence of NAJ the learner must learn to say BEF; in the presence of BEF, LUJ; in the presence of LUJ, PES; and in the presence of PES, CED. The same relations could be created as a paired-associates list consisting of NAJ-BEF, BEF-LUJ, LUJ-PES and PES-CED. The learner who mastered the original serial list would probably quickly learn this paired-associates list, and vice versa.

ADDENDUM 21B: PAIRED-ASSOCIATES LEARNING

Despite the relation between paired-associates and serial learning, paired-associates procedures were developed independently of Ebbinghaus's work on serial learning, by Mary Calkins, an American psychologist. Her earliest experiments used colors rather than verbal items as stimuli. Her response items were two-digit or three-digit numbers (Calkins, 1894). Verbal materials later became fairly standard for both stimuli and responses (Calkins, 1896), and paired-associates procedures, like serial-learning procedures, examined effects of a variety of verbal properties.

The learner's task appears simpler with a paired-associates list than with a serial list: Each item occasions a unique response and the order of items is irrelevant. But even this seemingly simpler task can be broken down into more fundamental components. Consider paired-associates learning as an instance of discrimination learning. In contrast to procedures with a single response, as when a pigeon's pecks are reinforced during green but not red, paired-associates learning involves many stimuli and many responses. The several stimuli must be discriminated, the several responses must be differentiated, and each response must be occasioned by its particular stimulus. Discrimination among the items is usually taken for granted; discriminations among various letter and word sequences are well established in human learners. But before learning associations, the learner must also learn which responses are appropriate.

Even after such effects are sorted out the learning of associations remains complicated. Presented by itself, a single stimulus-response pair can be learned very quickly. To what extent, then, is the study of association complicated merely because many pairs are learned together at one time? (The circumstances resemble those in the history of maze learning, when mazes evolved through simpler and simpler forms, culminating in T-mazes and straight-alleys; cf. Chapter 5.) It became appropriate to study the learning of single associations rather than of many associations in a single setting (e.g., Estes, 1964). The issues rekindled a long-standing dispute in the psychology of learning, the controversy over *continuity* versus *discontinuity*. The question was whether learning occurred gradually and continuously or took place discontinuously, in an all-or-none fashion. Were associations learned a little at a time or all at once?

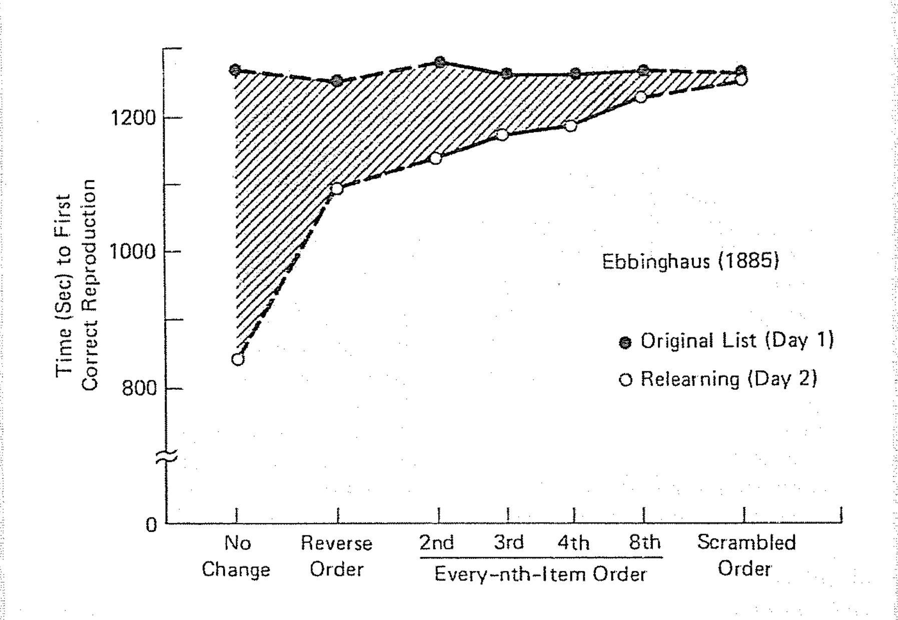

Figure 21-2 Original learning time (day 1) for 16-item lists of nonsense syllables and relearning time (day 2) after various reorderings of the lists: no change; reverse order; orders based on every second, third, fourth, or eighth item; scrambled order. The learning times shown are averages across at least 10 different lists learned on different occasions by Ebbinghaus. The shading shows savings as a function of the different types of reordering. (Adapted from Ebbinghaus, 1885/1964, Chapter IX)

No doubt associations can at least sometimes be learned on a single trial. It may be more important to ask about the properties of the association than to ask how quickly it can be formed. For example, is it symmetrical? Once the stimulus item consistently occasions the response item, will the response item occasion the stimulus item? The naming of written letters is a paired-associates task with written stimuli and vocal responses; if a child who can't write yet learns to say "A" when shown a written *A*, we wouldn't expect the child to write an *A* in response to the spoken letter (cf. Newman, 1972; Polson, Grabavac, & Parsons, 1997).

The issue of symmetry has practical implications. In learning Russian, for example, learning to give the equivalent English word in response to a word in Russian doesn't guarantee you'll be able to give the corresponding word in Russian when the English word is the stimulus item. You'd be well-advised to learn the symmetry explicitly by studying the vocabulary in both directions (e.g., English to Russian and Russian to English). Such symmetries are among the defining characteristics of bilingual skill (e.g., Kolers, 1966).

By definition, an association has been learned when a verbal stimulus occasions its paired verbal response (notice that this places the association in the learning situation and not in the learner). Typically, such responding is produced by instructions rather than by differential reinforcement, but that isn't the only reason why paired-associates learning is seldom treated in the language of stimulus control. Another reason is that any human verbal learning that isn't completely arbitrary will inevitably include much more than the learning of associations.

ADDENDUM 21C: FREE RECALL

One demonstration that a learner who has mastered a paired-associates list has learned more than associations is that the learner can usually name some response items even without any stimulus items. If the order of items is unimportant, this type of performance is called *free recall*. Recall experiments can be conducted with nonverbal stimuli. For example, we could present a collection of objects and then, after its removal, ask someone to name the items. A commonplace example is when someone asks us to name the people we met at a party or some other gathering. In experimental settings, however, free-recall procedures typically involve lists of verbal items. These procedures usually present a list once, followed by an opportunity for the learner to name the items without the list. Thus, free recall resembles the first trial of a serial-learning procedure, except that the learner needn't name the items in their original order.

Given the similarity to serial-learning procedures, it is no surprise that serial-position effects also occur in free recall. Learners are most likely to recall items at the beginning and at the end of a list. These effects are often described in terms of two principles: primacy and recency. The principle of *primacy* states that the first items of a list are more likely to be recalled than later ones; the principle of *recency* states that the most recent items (i.e., those at the end of the list) are more likely to be recalled than earlier ones. It follows that items in the middle of the list are least likely to be recalled. But these principles don't explain serial-position effects in free recall; they just summarize what learners do.

Free recall makes an effective class demonstration. If you'd like to try it, stop reading as soon as you get to the end of this paragraph and then turn to the next page, at the top of which you'll find Table 21-2, a 75-item list of words. Before going there, get yourself something to write on and something to write with and be sure they'll be handy when you need them. Try to do this exercise at a time and in a place where you're not likely to be interrupted. Once you're ready, turn the page and carefully but quickly read just once through the list, starting on the left and going down each column. Don't linger over any of the words. Immediately after finishing the list, close your book or cover that page and write down as many of the words in the list as you can remember. Don't worry about the order in which you write them down. And now, maybe you should read these instructions through once more before you get started, just to check in case you've missed something. But once you're sure you're ready, go for it!

Don't read past this point if you want to do the Free Recall demonstration. Go back to the last paragraph.

What follows assumes that you have handy the list of words that you recalled from the free-recall list. If you haven't read through the list but think you might like to do so later, stop reading now and come back to this section later.

If you are a typical unpracticed learner you probably recalled considerably fewer than half the items in this list, given its length. In fact, it's unlikely that you remembered more than 15 or 20 words out of the 75 and you may even have remembered fewer than 10. You might or might not have recalled the earliest ones (VIOLIN, SEAGULL, LEG), but you were more likely to recall them than the later ones. You probably recalled one or more items from the end (MOUTH, GUITAR, RAVEN) and you probably also recalled the distinctive one in the middle (**DINOSAUR!**). But yours are data from only a single person, so in comparison to findings polled across a large group of list learners, you may find that yours illustrated only some of the phenomena we'll look for. Some verbal learning research uses techniques for quantifying repeated recalls and recall summarized over groups of learners (e.g., Tulving, 1962).

If you especially recalled the earliest and the latest items in the list, you demonstrated a primacy and a recency effect. The position of an item in a list is only one of several factors influencing its likelihood of recall. The greater likelihood of recall of unusual items in a list is called the von Restorff effect (VonRestorff, 1933). It works with a variety of distinctive features, such as color (e.g., a word in red on a page of standard print) or size (e.g., a word in large type). Semantic novelty may also be effective (e.g., the name of a flower in the middle of a list of carpentry tools, or a common verb embedded in a list of animal names).

But did you recall ORCHID, which came just before **DINOSAUR**? Sometimes a distinctive event reduces the likelihood of the recall of items that came before it. If this happened to you, you might regard it as a small-scale version of *retrograde amnesia*, the forgetting of what happened just before a traumatic event such as an automobile accident (Tulving, 1969).

A problem with explanations in terms of distinctiveness is that it is sometimes difficult to predict which items will be distinctive. Suppose that recall probability is much greater for one item on a list than for neighboring items. We might be tempted to say that the item must have been distinctive. But this doesn't explain anything; it simply demonstrates one condition for calling an item distinctive. But I don't have this problem when I do lists like these in class. I shout the distinctive word while jumping up and down, and so far that has always worked.

The important dimensions of free recall are in the learner and not in the list. Learners aren't passive; they behave with respect to what they are learning. For example, in free recall learners often rehearse recent items by repeating them vocally or subvocally. Learners have more opportunities to rehearse early words than later ones, so rehearsal may contribute to the primacy effect. But opportunities to rehearse near the end of lists is limited, so it is unlikely that rehearsal is implicated in recency effects. And if you followed the instruction to go through the list quickly, you probably didn't have time for any rehearsal anyway.

The difference between serial recall and free recall is in the instructions to the learner. In serial recall, the learner is instructed to name the items in their original order; in free recall, the learner isn't told that order is important, or perhaps is specifically instructed to ignore order. But this doesn't mean that freely recalled items are recalled in arbitrary orders. In fact, the order of item recall typically differs systematically from the original order in the list.

Consider our 75-item list. Words in five different semantic categories were distributed throughout it: flowers: e.g., ROSE, DAFFODIL, TULIP; academic disciplines: e.g., CHEMISTRY, BIOLOGY, SOCIOLOGY; birds: e.g., DUCK, SEAGULL, CANARY; musical instruments: e.g., VIOLIN, FLUTE, PIANO; and body parts: e.g., ANKLE, SHOULDER, TOE). In the list, a word in a given category never appeared next to another word in the same category. But if you check your list of words recalled, you'll probably find that you grouped words from a given category together. At recall, the items within a category are likely to be recalled in clusters. This *clustering* depends on properties of the list and the learner's verbal history. For example, if you recalled DOVE that might have occasioned some of the other bird words; if you play a musical instrument you probably recalled more items in this category than would someone without musical interests; and so on.

In the recall of a list like this one, learners are likely to report some items that are semantically related to words in a cluster but didn't actually appear on the list. Such *intrusions* become more likely with longer lists and with delays between list presentation and recall. For example, only some of the following fifteen words were on the original list: VIOLET, OWL, MATHEMATICS, MEDICINE, TROMBONE, HAND, DAISY, GEOLOGY, ARM, ORIOLE, PSYCHOLOGY, CELLO, KNEE, IRIS, HAWK. Can you say which ones they are? Of the others, were any on your recall list? If so, you remembered something that never really happened. As we'll see in Chapter 26, false memories such as these are not uncommon, and sometimes they involve events far more significant than the words presented in a recall list. And now I'll let you know that fewer than half of the words above appeared in Table 21-2. You'll find an easy way to check that out before the end of this section.

Clustering in free recall is most easily demonstrated with lists that include semantic groupings, as in the 75-item list that was our example. But it occurs with all sorts of lists and along various dimensions of verbal items, such as words related in spelling or pronunciation or words that are common free associates (e.g., Bousfield, 1953). If you recalled VIOLET or IRIS, neither of which was on the list, the former might have been strengthened by its echoic relation to VIOLIN, which was on the list, and the latter by its thematic relation either to flowers or to body parts, and in particular to EYES, which was on the list.

It is hard to see consistencies of structure in the recall of just one list by a single learner, so perhaps you saw only a few of the phenomena we've just reviewed. For example, you either recalled or didn't recall CROW and ORCHID, the words that came just before the distinctive von Restorff word, DINOSAUR. But that was also true of a lot of your other words, so we couldn't attribute it to retrograde amnesia if you didn't recall them.

But group data will sometimes show such effects, though a definitive study would have to control for other effects, as by randomizing the locations of words in the list across different learners. When I've used these sorts of lists in class demonstrations I've pooled the class results, so now let's look at some data. Figure 21-3 shows the recall results for a class of 45 students to whom I read the list from Table 21-2 (with suitable animation of Word 37). The group data show both primacy and recency effects, with the former somewhat larger than the latter. Virtually all students recalled the von Restorff word, and the low

Table 21-2 Free recall demonstration. Start with the left column and read through each column carefully but quickly. When you're done, set the book aside and write down as many words from the list as you can remember.

violin
seagull
leg
algebra
banjo
daffodil
nose
canary
neurology
shoulder
rose
harp
duck
philosophy
lilac
ankle
saxophone
geranium
eagle
lips
engineering
tulip
fiddle
hand
chicken
astronomy
clarinet
petunia
toe
dove
geography
woodpecker
trumpet
head
crow
orchid
DINOSAUR!
history
penguin
cheeks
piano
geometry
honeysuckle
turkey
thumb
cello
economics
daisy
vulture
bassoon
arm
calculus
flute
pigeon
gardenia
tuba
eyes
psychiatry
falcon
drum
thigh
chemistry
goose
biology
oboe
pansy
knee
medicine
horn
poppy
owl
sociology
mouth
guitar
raven

Reminder: If you are doing the free recall demonstrations, stop here and write down your recalled words right away; then go back to where you left off on the last page.

recall levels for the words just before are consistent with retrograde amnesia. The substantial variability among words is fairly typical of demonstrations like this one; some words are simply more likely to be recalled than others. For example, Baltimore students who are sports fans may be more likely to recall RAVEN than other bird words.

Table 21-3 presents our free recall words arranged alphabetically within each category. Two additional columns, Category Intrusions and Non-Category Intrusions, list some but not all of the words that were not on the list but were falsely recalled by at least one student; there were seventy or more intrusions in the class as a whole, with some students having none, a few having several, and an occasional response being illegible. The most frequent intrusion was PSYCHOLOGY, with eight instances, and with some students quite certain they had heard it. But in a course on learning and with words like PSYCHIATRY amidst other academic disciplines, I was able to set up the list to make that intrusion highly likely. The students who recalled PSYCHOLOGY were remembering something that hadn't happened. Taken together, these data demonstrate that learners aren't passive when asked to recall verbal lists. Clustering the list items is just one of several kinds of things that they are likely to do during free recall.

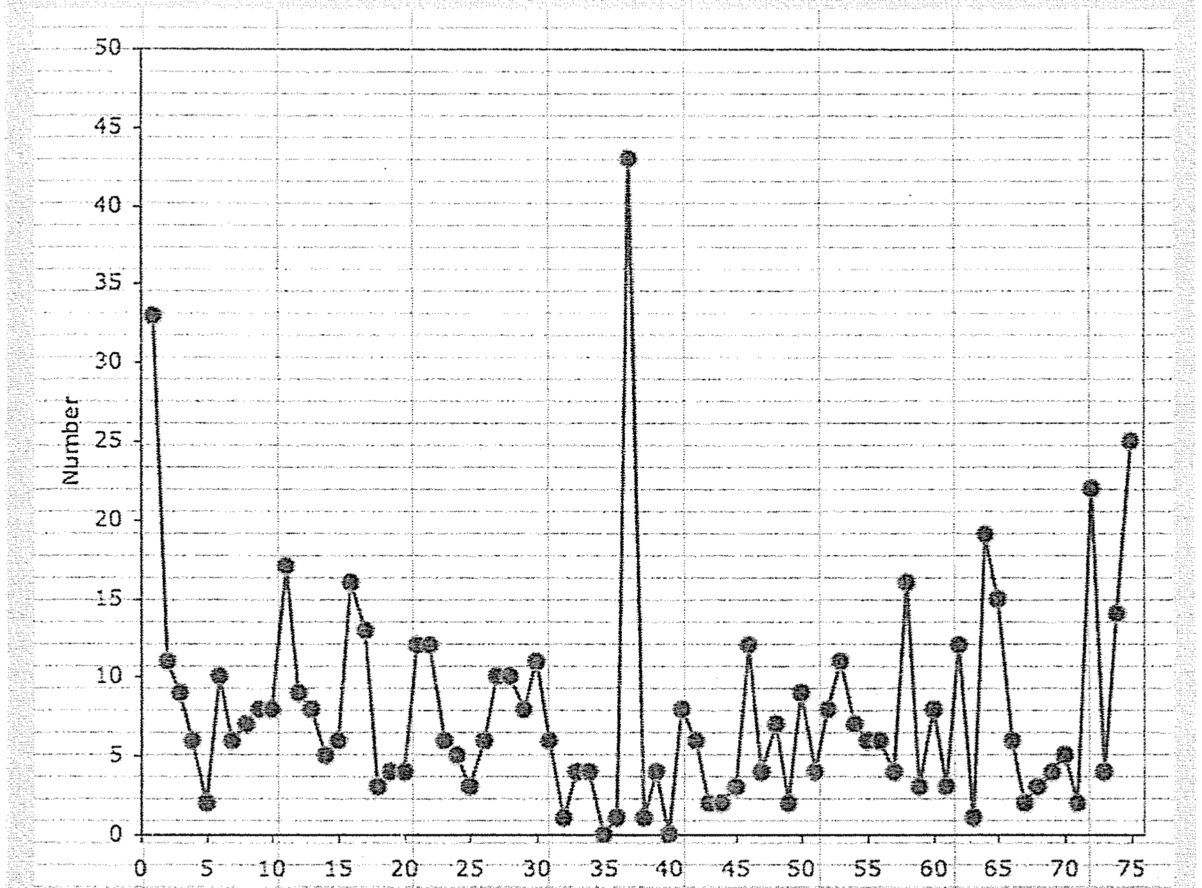

Figure 21-3 Free recall data from a classroom demonstration with a class of 45 students, where number shows the number of students who recalled the words in positions 1 through 75 of the word list in Table 21-2.

Table 21-3
Words from the Free Recall list of Table 21-2 sorted alphabetically within categories. The Intrusion columns list words not on the list that were falsely recalled. Most intrusions were words appropriate to the categories (Category Intrusions). Other types of intrusions, including the category names themselves as intrusions, are shown in a separate column. The number to the left of each list word shows its position in the list; see the data in Figure 21-3.

	Body Parts		Musical Instruments		Academic Disciplines		Birds		Flowers	Category Intrusions	Non-Category Intrusions
16	ankle	5	banjo	4	algebra	8	canary	6	daffodil	psychology	bird
51	arm	50	bassoon	26	astronomy	25	chicken	48	daisy	violet	flower
40	cheeks	46	cello	64	biology	35	crow	55	gardenia	mathematics	giraffe
57	eyes	27	clarinet	52	calculus	30	dove	18	geranium	geology	sneeze
24	hand	60	drum	62	chemistry	13	duck	43	honeysuckle	trombone	bee
34	head	23	fiddle	47	economics	19	eagle	15	lilac	lily	coffee
67	knee	53	flute	21	engineering	59	falcon	36	orchid	sunflower	hobo
3	leg	74	guitar	31	geography	63	goose	66	pansy	hawk	corn
20	lips	12	harp	42	geometry	71	owl	28	petunia	robin	calculation
73	mouth	69	horn	38	history	39	penguin	70	poppy	oriole	poetry
7	nose	65	oboe	68	medicine	54	pigeon	11	rose	dandelion	language
10	shoulder	41	piano	9	neurology	75	raven	22	tulip	physics	purple
61	thigh	17	saxophone	58	psychiatry	2	seagull			sparrow	red
45	thumb	33	trumpet	72	sociology	44	turkey		Distinctive	iris	horse
29	toe	56	tuba	19	philosophy	49	vulture	37	**DINOSAUR**	archeology	toad
		1	violin			32	woodpecker			trigonometry	giant

ADDENDUM 21D: VERBAL DISCRIMINATION AND VERBAL RECOGNITION

The distinction between simultaneous and successive discrimination (Chapter 2) is relevant with verbal as well as nonverbal procedures. Consider verbal discrimination involving a simultaneous procedure: Two verbal items are presented and the learner chooses one (by naming it or pointing or pressing one of two buttons). The learner is then told whether the choice was correct. Unlike discriminations studied in animal laboratories (e.g., a rat's lever pressing reinforced in light but not dark), the procedure usually includes many different stimulus items in both the correct and the incorrect classes. Now consider a successive procedure: Given verbal items presented one at a time, the learner is instructed to respond "yes" or "no" to each one; alternatively, the learner is instructed to sort a deck of cards with verbal items printed on them into two or more piles. Another case of verbal discrimination is *visual search*, as when someone is instructed to find items in a list or a text (e.g., Healy, 1976). Examples are checking a word in a dictionary or proofreading a manuscript for typographical errors.

In simultaneous verbal discrimination, the items and their assignment by the experimenter to correct and incorrect classes are sometimes arbitrary (e.g., pairs of nonsense syllables, with the items to be correct chosen randomly). Such arbitrary discriminations may not be relevant to verbal discriminations that are already part of our everyday language (e.g., discriminations among nouns, verbs and other parts of speech, or along semantic dimensions, or along dimensions of alphabetical order or spelling). Some experiments have explored the acquisition of verbal discriminations in which some feature is always a property of the correct item. For example, a correct item might be defined as one that includes a particular letter or letter combination or one that doesn't include a vowel or one that falls into some semantic class. Such procedures are cases of concept formation that involve verbal stimuli.

Verbal learning procedures usually repeat a list until the learner meets some criterion of correct responding. This is appropriate when we're interested in the relative difficulty of different types of lists. But when we're interested in teaching, we can design lists so that, through a progression of changes, they will efficiently produce a verbal discrimination. In other words, verbal discriminations are well suited to fading procedures (cf. Chapter 11). The following progression of items is an illustration; the correct item of each pair is indicated by bold italics (adapted from Goldiamond, 1966):

A – ***B***
B – D
B – V
O – ***B***
FS – ***BF***
JB – EL
EAB – VOX
ABE – JANE
MARY – ***BERT***
BILL – CONNIE
DAISY MAE – ***LI'L ABNER***
ROBERT – MARIA
ANNA – ***BOB***
BOY – GIRL
DANNY BOY – BETTY ANN
BYRON – BETH
BETSY – ***WILLIAM***
DOROTHY – ***GEORGE***
KENNETH – ANTONIA
JOHN – ELIZABETH

Early in the list, the discrimination is based only on the letter *B*. At the transition from letters to words, the semantic difference, male-female, is introduced, with the male item always paired with the letter *B*. Later the *B* is included in both items, and in subsequent pairs the *B* is removed so that the correct item depends only on gender. The gender discrimination in the final pair also includes a reversal of the letter discrimination that was created at the beginning of the list. The fading from pictures to written items is an effective method for teaching children early reading skills (Gleitman & Rozin, 1973).

The discrimination between items in a list may be based upon which ones appeared in some earlier list. When this is so, we speak of the discrimination as verbal recognition. In experiments on recognition, the response might be pointing at an item or naming it; in a less formal setting, it might be saying, "I recognize this one." The response depends not just on the stimulus but also on the past circumstances in which the item had been presented.

Recognition isn't restricted to verbal items. Discriminations based on whether people or places or objects are familiar or unfamiliar are examples of nonverbal recognition. But verbal recognition doesn't require discriminations between novel and familiar stimuli;

instead, it's based on whether the verbal item occurred in some previous setting (i.e., a particular list).

A personal incident persuaded me of the differences between recognition and other kinds of remembering. When I was a high-school student living with my parents in a small upper-Manhattan apartment, television was new and my parents often watched evening shows. One way in which I could mask the television sounds so I wouldn't be distracted by them was to play music. A friend had introduced me to tape recorders, which at the time were new but affordable, so I got one and taped music from WQXR, the classical radio station of the time, to serve as my masking sound tracks.

I often included the announcer's introductions on my tapes, but I didn't know at the time that some of the pieces I had recorded were unusual. One in particular was a tone poem by Tchaikovsky called *Fatum*, an early work of his that had rarely been recorded. Along with other more familiar classical music I listened to it dozens of times. At least two decades later, after I had recorded many other things over my old tapes and then moved to other kinds of sound systems, I was reminded of the Tchaikovsky piece for some reason. I recalled it as starting with a series of strong chords from the full orchestra and then fairly abruptly switching to a flowing melody, with the opening themes repeated later. But when I tried to locate a recording of it I found that a piece by that name did exist but recordings of it were no longer available.

Some years passed when I discovered that a new recording of the work had just been issued on long-playing records, along with some other Tchaikovsky tone poems I was unfamiliar with. I got myself a copy of the disk at my earliest opportunity, placed it on the turntable, and started listening. It had all of the properties I had remembered verbally, but nothing happened. I played the other works on the disk in case the names had somehow been mixed up, and then played *Fatum* again. Again, not much of a reaction, but it sounded good enough that I played it one more time, and only then did I have a full-blown reaction. For the next week or two I could not shake its melodies, which repeatedly intruded into my activities; what once was music I had used for a masking function now needed to be masked itself. This was clearly the piece I had listened to many times, and though my verbal recall of its musical structure had survived pretty well, when at last I had a chance to hear it again after its very long absence my listening wasn't ready for it; my recognition came only slowly. Listening too is behavior, and in this instance it took some repetitions to bring it back to where it used to be.

Verbal recognition requires responding in the presence of items whereas free recall requires producing items in their absence. We can expect these two classes of behavior to have different properties. Nevertheless, many accounts of verbal learning have assumed that recognition and recall are closely related. One account (Anderson & Bower, 1972) suggests that recall occurs in two stages: First the learner generates items, and then the learner produces only those recognized as being from the list. Colloquially, we might say that the learner must think of the item first and then decide whether it was on the list.

One deduction from such accounts is that recognition accuracy should always equal or exceed recall accuracy. This is often the case, but conditions can be arranged in which recall accuracy is consistently superior to recognition accuracy (Tulving, 1974; Watkins & Tulving, 1975), and such experimental demonstrations are far more persuasive than anecdotal cases like mine. The experimental outcomes not only make untenable the model of recognition as a component of recall; they are also likely to surprise those who regard recognition as somehow simpler than recall. Perhaps the most important point is that recognition and recall are different classes of responses occasioned by different classes of stimuli; they are measured in different ways and therefore it may be best simply to treat them as different kinds of behavior with different properties.

Chapter 22
The Contact of Verbal Behavior With the Environment

Some verbal behavior terms are Skinner coinages. Tact *is derived from the Latin* tactus, *touch, in the sense that here verbal behavior makes contact with its environment.* Mand, *by analogy with* command *and* demand, *is derived from the Indo-European* man-, *hand, as in* manual *and* manuscript.

Autoclitic *combines the Greek roots* auto-, *self, and* klinein, *to lean, as verbal behavior that depends on other verbal behavior. Relatives include* clitic *in linguistics,* incline, clitoris, *and perhaps* climax.

Abstraction
The Extension of Verbal Classes
  Metaphor
The Language of Private Events
Verbal Classes and Naming
Verbal Behavior Conditional Upon Verbal Behavior
  Descriptive Autoclitics: Discriminating Our Own Verbal Behavior

Verbal behavior would never have evolved had it made contact only with other verbal behavior. At some point, it must make contact with environmental events. We speak of this contact as *tacting*; a *tact* is a verbal response occasioned by a discriminative stimulus (cf. Skinner, 1957). For example, if a child learns to say "apple" in the presence of an apple, the child is said to be tacting the apple. The tact doesn't involve any new process; it's just a name for stimulus control as it enters into verbal behavior.

Superficially they may seem alike, but tacting differs from naming in somewhat the same way that textual behavior differs from reading. For one thing, tacting is sometimes a component of naming. Unlike naming, however, it can only occur in the presence of the tacted stimulus. We can name an absent object but we can't tact it. One reason for the distinction is that, as we'll see in Chapter 26 on remembering, our responses to past events aren't determined directly by those events; instead, they're determined indirectly by our previous behavior with respect to them. We'll return to naming later in this chapter, where we'll treat it as a higher-order class of behavior and consider some of its special properties.

An unlimited number of tacts is available to the mature speaker. We tact objects (chairs and tables, cellphones and books), living things (flowers and

KEY TERMS: Tacting versus Naming; Abstraction, The Extended Tact, Derived Tacts, Metaphor; The Language of Private Events; Autoclitics, Descriptive Autoclitics.

trees, birds and insects), weather conditions (rain and snow, sun and clouds), activities (walking and running, working and playing) and innumerable other features of the environment. Some tacts are general (e.g., man, woman) and others are restricted to relatively narrow circumstances (e.g., someone's name). The wealth of available tacts may be taken as a remarkable feature of human language, but we mustn't let it obscure the simple relation that defines an instance of tacting. **The relation between a tact and something tacted is precisely the same as the relation between a stimulus and the response that it occasions in a three-term contingency.**

It is useful to consider how tacts might be taught to a nonhuman organism; such examples force us to be explicit about their properties. Imagine a food-deprived pigeon in a chamber with a window on which we can project different colors. Next to the window are three keys. When the window is red, pecks on the left key produce food; when it is blue, pecks on the middle key do so; when it is green, pecks on the right key do so (cf. the adduction example in Chapter 13). The pigeon will eventually peck left during red, middle during blue and right during green. We can then call the pigeon's performance the tacting of red, blue and green. The stimuli are highly specific, and we wouldn't expect the pigeon to respond to red roses or flags or sunsets as it does to the red window. But the generality or specificity of the relevant stimuli is not a criterion for the tact relation. If we doubt whether the pigeon would generalize its response to red in other contexts, we can just call the pigeon's left peck a tact of a red window in this particular chamber.

What of the consequences of the pigeon's pecks? The pigeon would stop pecking if it weren't food-deprived or if food weren't a consequence. But even human tacting depends on its consequences. We don't go around tacting everything we see. Consequences are not criteria for tacting. An instance of tacting may produce approval, a grade, pay or the help it gives a listener; and, just as other responses aren't always reinforced, sometimes it may produce no consequences at all.

Finally, we might object that the pigeon has no audience. Its pecks are appropriately occasioned by red and blue and green, but the pigeon isn't speaking to anyone. We can easily rectify this difficulty by arranging signs, RED, BLUE and GREEN, that respectively light after the pigeon pecks the left, middle or right key. If our pigeon's performance was accurate and we couldn't look into the chamber, we could watch the signs and let the pigeon tell us which light was on.

But this arrangement isn't relevant to the issue. We don't need to check whether someone is listening to decide whether we've tacted something. Audiences are important in creating and maintaining tacts, but they aren't criteria for the tact relation. Tacting can be modified by audience variables and by consequences, but these variables don't define tacting. The pigeon's pecks in the presence of red and blue and green are relatively simple instances of discriminative responding, but any time the environment occasions what we say, stimulus control plays a role in our verbal behavior. The point warrants one more repetition: **When we speak about tacting we're just speaking about stimulus control as it enters into verbal behavior.** Now I'll have to convince you that this way of talking about words matters.

ABSTRACTION

In many instances in which stimuli occasion responding, we're interested in specific discriminative stimuli. For example, when we study a pigeon's pecks on red and green keys, we're not likely to worry about the pigeon's responses to red and green stimuli in other contexts. But in verbal behavior we'll often be concerned with responding occasioned by properties of the environment rather than by particular stimuli or classes of stimuli. In the tacting of red, for example, the property of color is the critical determinant of the verbal response *red* whether it's occasioned by red fire engines, red traffic lights or red noses. We call verbal discrimination based on a single property of a stimulus *abstraction*.

The tacted property is defined by the practices of the verbal community; it doesn't depend on whether we have some independent physical measure of it. For example, stage lighting sometimes uses a color contrast effect called *Hering shadows*, in which we see a shadowed area as some color even though there's no light of that color illuminating it. No range or distribution of wavelengths exists such that all visual stimuli within the range are called red whereas all those outside aren't, so even discriminations based on relatively simple stimulus properties (e.g., all red stimuli as opposed to specific red objects) may be to some extent arbitrary. **We don't need to specify some physical dimension of the stimulus to decide whether a particular verbal relation is a tact**; in defining our terms, we start from behavior and not physics.

A concept formation experiment illustrates how arbitrary classes can be generated (Hull, 1920) (Hull, 1920; cf. Chapter 8 on concepts). Learners mastered names for the members of sets of Chinese characters (three are shown in Figure 22-1). A name was consistently related to a radical that appeared in one character in each set, but the configuration within which it appeared varied from one set to another. Each set of 12 characters was presented until the learner gave the appropriate name for each and then a new set of 12 was presented. By the fifth set, learners named more than half the characters upon seeing them for the first time, and sometimes they gave the name even though they couldn't sketch the radical or describe the basis for their naming.

Tacting in this experiment differs from tacting in natural languages in that we can't define the basis for many of our tacts so explicitly. For example, we can't say exactly what properties make an object a chair. The word *chair*, whether spoken or written, is an arbitrary class by virtue of the common verbal consequences arranged by verbal communities. A chair may have four legs or stand on a single pedestal, it may have a flat or a contoured seat or back, and it may be made out of a variety of materials. We can't even appeal to its function,

Figure 22-1 Radicals common to the Chinese characters in six lists (columns). Learners had to give the radical name to the character that contained it, but no character was repeated across lists. Learners became able to name characters upon seeing them for the first time in new lists, and sometimes did so even when unable to define the common radical by sketching it. (From Hull, 1920, Figure 1)

because we call some objects chairs even though we can't sit on them (e.g., a toy chair in a set of dollhouse furniture). To make things still more complicated, they can also be embedded in other classes, as when they're pieces of furniture. This is not a traditional way of thinking. For example, this view of classes of objects and events contrasts drastically with Plato's essentialism, the idea that classes of things in the world share some common essence; the contrast also illustrates the antiquity and ubiquity of the problem.

We can tact extremely subtle properties of events. They include relations among stimuli. For example, *above* and *below*, *near* and *far*, and *larger* and *smaller* tact properties of stimuli in relation to each other or to the speaker. Relational tacting occurs when you say two objects are alike or different, or when you note that one item in a set is an odd item (cf. Chapter 12). Such terms rarely stand alone. We'll consider their joint dependence on other properties of events and on other verbal responses when we consider the verbal relations called *autoclitic*. We also tact complex events extended in space or time, as in identifying a musical piece as one by Debussy or a painting as one by Monet. Identifying a musical work on the basis of just part of a melody shows how large a tacted unit can be. At another level of complexity, we might say that Debussy's music and Monet's paintings have something in common, even though it would be difficult to specify the common dimensions. Yet calling them both impressionists might be regarded as tacting common properties.

Sometimes the properties that occasion a tact can be identified more with the speaker's own behavior than with any particular stimulus feature. For example, if a painting or a musical composition or a situation occasions the word *marvelous*, this tact presumably depends on the responses generated in the speaker rather than on physical properties common to all of these stimuli. Saying that we've just done something voluntarily or involuntarily probably depends on many factors, such as what we know about the antecedents and consequences of our behavior, how much variation was possible, and whether it was elicited or coerced (cf. Neuringer & Jensen, 2010; Neuringer, Jensen, & Piff, 2007). And it would only qualify as tacting if we said it at the time rather than an hour or a day later.

An interesting verbal case is the *tip-of-the-tongue* phenomenon. When I say a word is on the tip of my tongue, I'm tacting the near-threshold availability of an appropriate verbal response in my own verbal behavior. Sometimes I can even report properties of the unrecalled word, such as its length or part of its spelling (Brown & McNeill, 1966).

The vocabulary of emotion is similarly based on complex relations involving situations and behavior (cf. Chapter 18). Tacts of love, hate, joy and sorrow, whether in ourselves or others, depend on overt manifestations such as laughter or tears and on the circumstances that generated the behavior we observe. The variability of the language of emotions is itself evidence of the subtlety of the relations that we tact. But these tacts cannot be based on what is inside us, because if this were so a verbal community could not maintain any consistencies in its vocabulary of emotions. No matter what some politicians might say, we cannot feel each other's pains.

Events and situations obviously have many properties that might be tacted. Whether any are tacted and which are tacted will depend on other variables acting upon the speaker. For example, we may tact the color of an apple in one circumstance and its smell in another. The situation presents no difficulties; verbal responses are determined in multiple ways.

Consider again the pigeon example. Let's alter the lighting of the window beside the three keys so that we can present bright, moderate or dim colors. During a tone, pecks on the left key produce food when the light is bright, those on the middle key do so when it is moderate and those on the right key do so when it is dim, in each case regardless of color. Without the tone, everything is as it was before. If the pigeon's pecks become appropriate to intensity when the tone is present but remain appropriate to color when it is absent, we could say that the pigeon tacts intensity during the tone and color during its absence (cf. Chapter 11 on attention and Chapter 12 on conditional discrimination).

The presence and absence of tone in the pigeon example serve the same function as questions in human verbal behavior. A question, like the tone, may be a conditional stimulus that occasions the tacting of some stimulus property (e.g., "What color is the apple?" or "What does it smell like?"). Audiences, previous verbal behavior and other factors can affect tacting. We don't tact indiscriminately. Some things we tact only in some circumstances, and we also learn that sometimes (e.g., in remarking on someone's dandruff or bad breath) it isn't tactful to tact at all.

We tact temporal dimensions of stimuli when we say that something lasted a long or a short time. Often we respond verbally to stimuli that are no longer present, but such responding counts as tacting only if it occurs in the presence of or very soon after the events that occasioned them. Responses that occur long after require special treatment because they usually include other behavior besides tacting (cf. Chapter 26).

Sentences in different tenses can sometimes be regarded as tacts occasioned by temporal properties of the environment. Consider "It's raining," "It's beginning to rain" and "It just stopped raining." Each sentence is a response to rain, but they're distinguished by different temporal properties. A past-tense response such as "It rained" is likely to be determined by other kinds of stimuli, when it isn't appropriate to call it a tact. For example, it might be uttered as an echoic or a textual response. If it is occasioned by present stimuli such as wet streets, it may be derived from other verbal behavior, such as "The streets are wet; therefore it must have rained."

In dealing with the formal relations, we argued that the symmetry of stimulus and response relations favored a vocabulary in terms of words rather than one in terms of specific vocal or written modalities. Similar correspondences exist in relations between tacts and environmental events. These correspondences may be important when we speak of meaning, because the language of meaning is independent of whether words function as stimuli or as responses. In tacting, some correspondences relate words to environmental events, as illustrated in Figure 22-2.

This figure shows both a nonverbal stimulus, the sight or sound of rain falling or how it feels on your skin (*RAIN*), and a verbal one, hearing someone say "It's raining," and both a nonverbal response, taking an umbrella (UMBRELLA), and again a verbal one, saying "It's raining." The only difference between taking an umbrella after seeing that it's raining or being told that it's raining is that the former discriminative stimulus is nonverbal whereas the latter is verbal. Given that you are

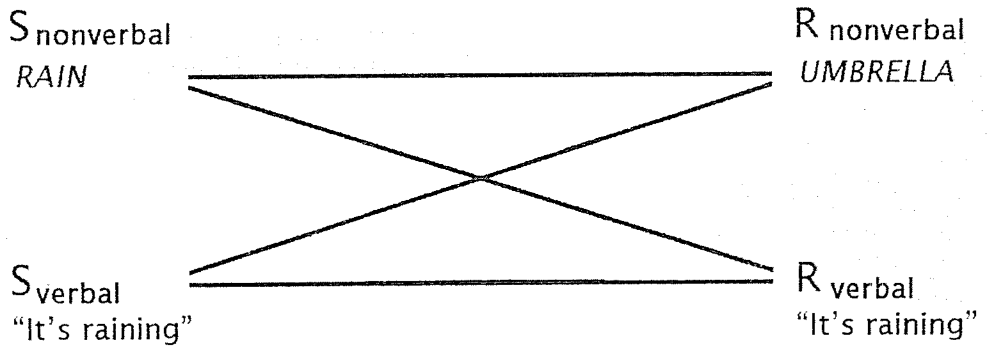

Figure 22-2 Relations among verbal and nonverbal responses occasioned by verbal and nonverbal stimuli (S, stimulus; R, response). Naming, meaning and other properties of verbal behavior depend on such consistencies in the relations among verbal and nonverbal events (cf. Figures 20-1 and 20-2).

about to go outdoors, the consequences of having the umbrella are similar; you will be able to avoid getting wet. The nonverbal case simply illustrates a nonverbal contingency; the verbal one is an example of what we'll treat as verbally governed behavior in Chapter 23.

Now consider the other two cases, in which the verbal response, "It's raining," is occasioned either by the rain or by someone saying "It's raining." The former is an instance of a tact and the latter is an instance of echoic behavior. This network of relations ties verbal behavior to the environment in somewhat the same way as those illustrated in Figure 20-1 tie together spoken and written verbal behavior.

All four cases involve responses occasioned by discriminative stimuli in three-term contingencies; for convenience, accounts of such relations often leave out the third term, which in these cases are the nonverbal consequences of carrying an umbrella and the social consequences of telling someone else about the rain. I may look out a window, see rain and pick up an umbrella on the way out or I may pick up the umbrella when I've just heard a weather report predicting rain. Given either the rain or the weather report, I might also phone someone who works in a windowless office and pass the news on. The power of tacting rests upon how solidly it is grounded in its relation to the environment, and that depends on the practices of a verbal community.

THE EXTENSION OF VERBAL CLASSES

The tact is a flexible relation. In some verbal communities, the stimulus properties that occasion a tact are sharply defined. A student in a science laboratory, for example, is taught to be consistent in tacting apparatuses and materials and procedures. This precision is less common in everyday discourse. We often tact properties of the behavior of our acquaintances, but the conditions under which we say that someone is warm or reserved, energetic or lazy, interesting or dull and so on, vary considerably from one speaker to another. The etymologies or word histories introducing each chapter recognize the ways in which vocabularies have evolved over time. "Etymology is the archaeology of thought" (Skinner, 1989b, p. 13).

In discriminations involving nonverbal behavior, we say a response has generalized if a response maintained during one stimulus occurs when some new stimulus is presented. For example, if our tacting pigeon pecked the left key when a novel amber light was presented, we'd say that its response to red had generalized to amber. A similar generalization of verbal responses to new stimuli occurs in the extended tact.

One type of extension of the tact occurs when new words are formed by combining existing ones (e.g., the words *dish* and *wash* predated the invention of dishwashers). Vocabularies change with environmental changes that are important to the speakers of a language. The ways in which tacts can be extended are so varied that a detailed account isn't feasible. Stewart (1975) offers interesting examples in an account of the origins of place names, which are more likely to be based on unusual than common features of a region. A valley in a forest of fir trees wouldn't be named Fir Valley, but it might be named Oak Valley if a single oak tree stood there. Similarly, no stream where wolves are common is likely to be called Wolf Creek, but a stream where wolves are rare might get that name once a lone wolf is sighted there.

We've already noted the restriction that tacting must occur in the presence of or very soon after the event tacted. Then what about words that seem superficially to be tacts but cannot occur in the presence of what they name? Do we ever see political units like states or nations, subjects like economics or politics, properties like ownership or indebtedness, processes like creation or evolution, and so on? And how about months or days of the week? Such entities must enter into our verbal behavior in other ways; they do not exist in a form that can be tacted. The point is that **tacting is not defined by parts of speech or other linguistic categories; it is a type of behavior.**

Metaphor

Simile and metaphor are other familiar instances of extended tacts. We may say that someone is as busy as a bee or as sly as a fox (simile) or that someone is a hawk or a dove (metaphor). These extended tacts presumably originated through generalization across subtle properties sometimes shared by the behavior of humans and bees and foxes and birds. Language grows and changes through metaphor (cf. Esper, 1973; Jaynes, 1976; Skinner, 1989b).

We tend not to notice as terms evolve from specific origins to technical language. Files and scrolling and memory in computer terminology once had more commonplace senses, and surfing came from a totally different domain. We all know what slides are in a contemporary computer presentation, but how many remember talks in which a glass or plastic transparency was manually inserted into a projector at the back of a room? Metaphor is so pervasive that the surprise of puns and other forms of verbal humor typically comes from overly literal interpretations of standard usages. In those days before remotes, when slides literally slid, colleagues usually worked the projectors for those giving talks. You could get a laugh (or worse) if you heard "May I have the next slide?," and instead of showing it you responded literally by pulling it from the projector and running up to the front of the room to hand it to the speaker.

Once particular metaphors become effective within a verbal community, they're likely to evolve and interact and spread to a variety of situations. Language itself provides an important example: We speak of language in the metaphor of communicating ideas. According to this metaphor, ideas and meanings are objects placed into words and then delivered to someone else (Lakoff & Johnson, 1980; Reddy, 1979). We *put our ideas into words*; we *have* ideas and *get them across to others*; our words *carry meaning* or are *empty*; our sentences *contain* or are *filled with* ideas; ideas can be *grasped* or *dropped* or *kicked around*; and so on. This metaphor is so well established that, as these chapters on verbal behavior testify, it is difficult to speak of language in other ways.

Other common systems of metaphors in our culture are those of time as money (e.g., we *spend* time or *save* it, and ask whether something is *worth* our time); of understanding as seeing (e.g., we *get the picture* or *look at things differently*, and describe arguments as *clear* or *opaque*); and of more as up and less as down (e.g., prices can *rise* or *fall*, and someone can be *under*age or *over*charged).

Systems of metaphors may be coordinated with each other. For example, saying that a theoretical argument can be demolished by attacking its weak points combines the metaphor of discussion as war, in which arguments are *marshalled* or *undermined* or *shot down*, positions are *defended* or *given up*, points are *won* or *lost*, with that of theory as a building, in which theories are *constructed*, rest on *solid* or *shaky foundations*, can be *supported* or *buttressed*, and *stand* or *fall*. On the other hand, systems of metaphors needn't be consistent. For example, the metaphor of discussion as war may only occasionally make contact with that of discussion as exploration, in which issues are *gone over in depth* and *at different levels* and *from different approaches*, the *ground is covered*, the speakers *map out their territories* and make *direct* or *roundabout* arguments (Lakoff & Johnson, 1980).

Metaphor is a pervasive property of language. Children learn it readily and adults can't ignore it (Glucksberg, Gildea, & Bookin, 1982; Winner, 1979). In judgment tasks, reaction times are often shorter for metaphorical than for literal usages (Foss, 1988). It isn't just the stuff of poetry; it is a fundamental aspect of verbal behavior. It is an inevitable product of the extensions of tacts in complex verbal environments. If there isn't a word for it, the verbal community will invent one.

The phenomenon of metaphor tempts us to talk about abstract properties captured by words. Yet its most important feature is that **it allows us to deal with the abstract in terms of the concrete** (Johnson, 1987; Lakoff & Johnson, 1999). For example, the language of abstract dimensions like good-bad or happy-sad becomes that of a more accessible dimension, up-down, from raised spirits and the heights of cloud nine to lowered expectations and the depths of depression. Hardly any dimension can be more abstract than time, but through metaphor it becomes a concrete spatial one: Tomorrow versus yesterday becomes in-front-of versus in-back-of. We're so used to saying our past is *behind us* and our future is *before us* that it's hard to imagine taking an about-face so that the future is in back and the past lies ahead; our time line turns with us.

The ubiquity of metaphor is evident in the etymologies that introduce each chapter. Much of our technical vocabulary evolved metaphorically from everyday sources, and much of the everyday language that applies to our own behavior arose through metaphorical extensions. *Mind* was once a verb, and many cognitive or mentalist terms have mundane origins (cf. Skinner, 1989b). These etymologies demonstrate some of the very concrete sources of our fundamental concepts. The creative aspect of metaphor, in other words, is in making the abstract substantial, specific and solid or, to mix more metaphors, in bringing it down to earth (Lakoff, 1987).

The similarities of some metaphorical extensions across cultures suggest common contingencies in the human verbal histories that shaped them. For example, the word *psychology* comes from the Greek *psukhein*, to breathe, and *psukhe*, breath. When the psyche was said to leave the body of a mortally wounded warrior in the Homeric epics, it was understood to mean the warrior's breath and not his soul. A curious parallel exists between this word and *spirit*, from the Latin *spirare*, another word meaning to breathe, but no evidence links these Greek and Latin words etymologically. A similar relation between air and spirit exists in the Latin *anima*, originally a breath of air but later soul or spirit in such words as *animate* and *animosity*, and in the Greek *atmos*, vapor or air as it enters into the English *atmosphere* but breath or soul in the related Sanskrit *atman*. In all human cultures air is invisible but it can be felt in breath and wind. Air and breath are indispensable to life itself, so it should come as no surprise that the corresponding words have so often been inextricably intertwined.

We haven't explicitly defined *metaphor*. As in other cases, our failure to do so does not imply that the term is meaningless. We don't need to independently specify the stimulus to identify discriminative relations, and so also we don't need to independently specify verbal classes to identify the relations between words and events that we call meanings. Meaning is a shorthand for saying what they have in common as a class. We speak of classes of responses in terms of operants and classes of stimuli in terms of discriminated operants or concepts. Words can function either as responses or as stimuli, so it may sometimes be appropriate to speak of thematically related verbal classes in terms of meanings (see also Lowenkron, 2004).

THE LANGUAGE OF PRIVATE EVENTS

One more crucial extension of the tact is to private events. Tacted stimuli are sometimes accessible only to the speaker, as when we say we have a headache. Such tacts depend on the verbal community for their origins and maintenance. The problem is how the verbal community can create and maintain these responses when it doesn't have access to the stimuli. A parent can teach a child color names because the parent can see the colors that the child sees and therefore can respond differentially to the child's correct and incorrect color naming. (So many varied consequences follow from color naming that it usually doesn't matter whether the parent teaches the color names explicitly or simply allows them to be learned through casual day-to-day interactions; we'll come back to naming in the next section.)

With private events, however, the vocabulary can be taught only through extension from tacts based upon events to which the verbal community has access. For example, the child may learn to report pain because the parents have access to overt manifestations such as the event that caused an injury or the child's crying or facial expression; if the child has learned the names of body parts, the two kinds of verbal responses may be extended to the tact of pain in a particular place (cf. Skinner, 1945).

A toothache is a discriminable event, but the person with the toothache has a different kind of access to it than the dentist called upon to treat it. Both respond to the unsound tooth, but one does so by feeling the tooth and the other by looking at it and probing it with instruments. Their different contact with the tooth might be compared with the different ways a seeing person and a sightless person make contact with a geometric solid if one is teaching its name to the other; the seeing person does so by sight and the sightless person by touch. One kind of contact isn't necessarily more reliable than the other. For example, in the phenomenon of referred pain, a bad tooth in the lower jaw may be reported as a toothache in the upper jaw. In this case, the dentist is a better judge than the patient.

We probably think of private events such as our feelings and thoughts as ones to which we have privileged access and therefore of which we have special knowledge. But we learned the relevant words from others and in teaching them to us they had access only to the public correlates. If we can be mistaken even about the location of a toothache, what assurance do we have that our other reports of our private events are reliable? Skinner (1963) makes the point by describing some students who had watched a pigeon in a classroom demonstration and then described what they saw in terms of the pigeon's expectations:

> They were describing what *they* would have expected, felt, and hoped for under similar circumstances. But they were able to do so only because a verbal community had brought relevant terms under the control of certain stimuli, and this had been done when the community had access only to the kinds of public information available to the students in the demonstration. Whatever the students knew about themselves which permitted them to infer comparable events in the pigeon must have been learned from a verbal community which saw no more of their behavior than they had seen of the pigeon's. (Skinner, 1963, p. 955)

Some verbal responses that superficially seem to tact private events may be occasioned instead by the situation in which our behavior occurs. For example, if on sitting down to a meal you suddenly find yourself eating voraciously, you may say, "I must have been very hungry." You haven't tacted some private hunger pang; you are saying of yourself what you would have said of others if you had observed them eating that way.

Some time ago my colleague, Philip Hineline, collected students' reports of their feelings under various emotional circumstances. I saw an opportunity to extend his results when snow disrupted the exam schedule for one of my courses. I had explicitly told the students that I would not give an exam on the first day after a snow closure, but nevertheless I arrived that day with some papers and exam booklets under my arm. As I distributed the papers face down I instructed the students not to turn them over until everyone had copies. When I allowed them to turn over the paper and start, it read: "This is not an exam. Describe your feelings when you saw me start to hand out the exam booklets." After a few minutes I collected what they had written. I had asked them not to put their names on their papers, so their anonymity was preserved when I shared a few passages with the class.

Despite the instructions, the majority of the reports described not feelings but the situation ("There wasn't supposed to be an exam today" or "I wasn't prepared for a pop quiz") or a question about the situation ("I wondered what's going on"). The reports of feelings were almost exclusively limited to names ("I felt nervous" or "Fear!" or "I felt anxious"). A very few reported physiological events ("My heart raced"), and among those most were metaphorical ("My heart fell to the floor"). Had I asked again after a bit of discussion, I'm sure "Relief" would have been high on the list of reported feelings. But the point is that when we talk about our feelings very often what we say is more a commentary about our environment than about something happening inside of us.

This sort of cross-section of results has come up over several classroom variations, including some in which the instructions specifically encouraged descriptions of feelings rather than descriptions of the situation. But how many of us have been explicitly taught how to describe our feelings or what goes on in our own bodies? There are opportunities here. Given what we know about the history of teaching breast self-examination (Chapter 11), shouldn't we be able to do better at teaching people how to detect the symptoms that may precede a heart attack and to discriminate those from heartburn and sore muscles and other varieties of pain. How many other medical conditions exist for which patients might be able to contribute to their own diagnoses by learning to discriminate properties of their own bodies?

Private events can include behavior as well as physiology. Once we learn to tact properties of the public behavior of others, we may come to tact the same properties of our own behavior, whether public or not. If one person works hard at something with little compensation and another does so only with substantial compensation, we'll usually assume that the task was more important to the first person than to the second. But the same observations of our own behavior may also lead us to say what's more or less important to us (Bem, 1967). Speaking of our beliefs or our understanding of the causes of our actions may follow more directly from our discriminations of our own public behavior than from anything private (e.g. Kiesler, Nisbett, & Zanna, 1969). This isn't to deny private events. It is instead a cautionary note: The language of private events can easily distract us from the public causes of behavior.

The fundamental problem with tacts of private events is that it is hard for the verbal community to shape and maintain them, because it has inconsistent access to the events and their public correlates. For example, when someone says "I have a headache" and leaves a social gathering, it isn't clear whether the verbal response tacted some private event or just allowed the speaker to escape from unwanted company. On such grounds, some have argued that private events have no place in a science of behavior (Baum, 2011; Catania, 2011b). But like the language of public events, that of private events depends on the public practices of the verbal community. The implications have been explored in detail by Skinner and by Wittgenstein. Day (1969) discusses parallels between their treatments of private events.

The implication is that **we know what is public better than we know what is private**, because the verbal community can teach us the language of public events more consistently than it can teach us the language of private events. How did it teach us when to say that we understand something, or that we're in love, or that we're uncertain, or that we're happy or sad? We know from examples like breast self-examination that sometimes we need to be taught even about our own bodies, so shouldn't we be similarly concerned about how we were taught about our own feelings and states of mind?

For our purposes, it is enough to note that verbal behavior doesn't ordinarily require stimuli to be simultaneously available to both speaker and listener. In fact, some important consequences of verbal behavior occur when the speaker tacts an event unavailable to the listener. For example, if I'm telling someone over the phone about something I've found at a website, the other person can't see what I'm seeing but may be able to do something about it (e.g., by switching to the same site that I'm looking at). In other words, the relation between tacting a public event and tacting a private event is pretty much the same as the relation between tacting when both speaker and listener have access to what is tacted and tacting when only the speaker has such access. Although the language of private events has its own special difficulties, we don't need new categories of verbal responding to deal with it.

By themselves, tact relations are only one part of verbal behavior, but through them verbal behavior contacts the environment. Without them there would be nothing about which we could speak. The question of truth is behavioral. Some of what we call truth depends on how verbal communities maintain correspondences between verbal behavior and environment. Those who lie do so because the consequences of lying differ from those of telling the truth, but lying can be effective only within verbal communities in which such correspondences are reasonably reliable ("Unless social interaction is to break down, the lie must always be the exception," (Bolinger, 1973, p. 549; cf. Dawkins, 1976, pp. 82, 112).

VERBAL CLASSES AND NAMING

We've already made the point that we have to keep tacting distinct from naming. Naming may include tacting as one of its components, but the everyday vocabulary extends to a sufficiently wide variety of cases that any formal definition will probably be unsatisfactory for some applications of the term. Let's again consider the pigeon whose pecks tact red and green and blue lights. It is a good bet that most readers would object to any argument that suggested the pigeon was naming those colors or that the pecks on the different keys were the pigeon's names for the colors. And they would be right to object. But the reason is not that we do something different when we name colors; instead, it is that in naming we do something like what the pigeon does but we also do a great deal more.

There are a few cases in which tacts are clearly very different from names. For example, *hello* is a characteristic response that is occasioned by a special social situation: we're introduced to someone or someone we haven't seen recently comes up to us, so we say *hello*. It is difficult to say exactly which situations qualify for this response and which don't, but these sorts of social stimuli occasion the *hello* as surely as a dog might occasion *dog* or a cat *cat*. Thus, though we wouldn't call it a name, *hello* qualifies as a tact. Another example is *thank you*. To the extent that we say *thank you* only given certain social stimuli that involve something that someone has done for you, *thank you* too qualifies as a tact, but here too we wouldn't call it a name. **Tacts, in other words, are examples of stimulus control in which the response is verbal; they are not defined by properties of the stimuli that occasion them.** Naming is something more.

Naming is a higher-order class that involves arbitrary stimulus classes (things or events with particular names) and corresponding arbitrary verbal topographies (the words that serve as their names) in a bidirectional relationship. Tacting is at the heart of naming, but tacting alone is not enough. Consider what a child must do before we're satisfied that the child knows the name of something, such as *shoe* as a name for one of those things that goes on your feet. Not only do we expect the child to say *shoe* when we hold up a shoe and say *What's this?* We also expect the child to look around when we say *Where's the shoe?* and then point to the shoe if it's in sight.

Superficially this looks like an equivalence relation, or at least a symmetry relation, because the stimulus object and response word in the first part of the example seem to switch places in the second. But that is assuming too much, because, unlike a pigeon's pecks on samples and comparisons, a seen object cannot be exchanged with a point at it and a heard word cannot be exchanged with a spoken one: "the relation between a name and that which it names is fundamentally asymmetrical" (Horne & Lowe, 1996, p. 234). Furthermore, a point will not be effective unless the child looks not at the caregiver but at where the caregiver is pointing. Thus, joint attention is a critical prerequisite for naming (Dube *et al.*, 2004; Tomasello & Farrar, 1986). It is a component of the behavior of the listener.

Starting by about 6 months of age, before they have begun tacting, children begin to master listener behavior (Gurteen, Horne, & Erjavec, 2011). Between the ages of 1 and 6 years, they increase their vocabularies at an average rate of five to eight words per day; by the age of 6, a child is likely to have a productive vocabulary of thousands of words (Wagner, 1985). Many function words (e.g., *more*, *bye-bye*, *allgone*) are included along with common nouns (e.g., *mama*, *milk*, *chair*) in the early vocabulary. What common features must the learning of each new word include to allow the child to learn so many new words so quickly?

Another prerequisite for naming is echoic behavior. The child is otherwise unlikely to produce a new word spontaneously. Echoic behavior might seem to be a variety of generalized imitation, but generalized imitation typically depends on a history of naming, so the echoic behavior must already have been in place before generalized imitation developed (Camões-Costa, Erjavec, & Horne, 2011; Horne & Erjavec, 2007). And what about equivalence classes? They seemed to be prerequisites for naming and other higher-order verbal classes, but they may instead have naming as their prerequisites (Horne, Hughes, & Lowe, 2006; Horne, Lowe, & Harris, 2007; Horne, Lowe, & Randle, 2004; Lowe *et al.*, 2002; Lowe, Horne, & Hughes, 2005; Randell & Remington, 1999)

Naming has still another feature. If a child already does some naming but hasn't yet learned the name *glove*, we might hold one up and say *This is a glove*; the child might immediately say *glove* and point at it. We then wouldn't be surprised a little later if the child pointed to the glove when we said *Where's the glove?* or said *glove* when we held it up and said *What's this?* even though the child had never answered these questions with respect to gloves before. The prerequisites for doing so include at least three components: (1) *listener behavior*, in looking for things and pointing based on what someone has said; (2) *echoic behavior*, in repeating names when they're spoken; and (3) *tacting*, in saying the names given the objects (Horne & Lowe, 1996). The child's responses to the questions about the glove are instances of novel behavior, and naming as a higher-order class enables them to emerge. The various social and nonsocial consequences of naming are individually small, but collectively they can sustain a lot of behavior. Soon the child may initiate naming instead of waiting for it to be initiated by others. It is common for a child who has learned some naming to turn the questioning around. On encountering something that doesn't have a name yet, the child points and asks *What's this?*

Naming is generated from the ordinary interactions between children and their caregivers. Once it is available as a higher-order class within a child's repertory, it allows for expansions of vocabulary in which the introduction of new words in particular functional relations (such as tacting) involves those words in a range of other emergent functions. These include but are by no means limited to intraverbal behavior, echoic behavior and orienting toward or pointing at named objects. As we learn more about the several verbal functions that constitute naming, the kinds of contingencies in the natural environments of children that shape and maintain those functions, and the role of naming in enlarging and amplifying other functions of verbal behavior, we will be in a better position to extend our analyses to applications, as in teaching naming to children on the autism spectrum (e.g., degli Espinosa, 2011).

VERBAL BEHAVIOR CONDITIONAL UPON VERBAL BEHAVIOR

Verbal behavior, like any other event, can be tacted. No new kinds of relations are involved, but the complexities created when verbal behavior is built upon other verbal behavior need special comment. Verbal behavior that involves the tacting or discrimination of our own verbal behavior and that modifies the effects of other verbal behavior is called *autoclitic*. A *descriptive autoclitic* is an autoclitic that tells the listener something about the rest of what the speaker is saying, as when we prefix a comment with a qualifier such as *I just saw that...* or *I've just been told that....* These autoclitics modify the effect on the listener of the rest of the sentence by letting the listener know that what follows is first-hand or second-hand information.

Skinner discussed another category of autoclitics that he called *relational autoclitics*. These involved verbal units that could not stand alone because they had to be coordinated with other verbal behavior (an example is the plural *-s*). This category does not depend on discriminations of one's own behavior and will therefore be separately treated in Chapter 25. Note that intraverbals also depend on and may modify the effects of other verbal behavior, but we distinguish them from autoclitics because intraverbals do not require discriminations of one's own behavior. For example, if you say *Red, white and...* I don't have to discriminate anything about my own verbal behavior to finish off the intraverbal sequence by saying *blue*. The verbal stimulus *Red, white and...* sets the occasion for the concluding word whether you say it or I do.

Descriptive Autoclitics: Discriminating Our Own Verbal Behavior

Many verbal responses tact the conditions under which other verbal behavior is emitted and thereby modify the responses of the listener. Consider the phrases *I doubt* and *I'm sure* in "I doubt the coffee is ready" and "I'm sure the coffee is ready." Each one modifies the way in which the listener is likely to act upon the statement that the coffee is ready. For the listener, *I doubt* and *I'm sure* are analogous to the conditional stimuli of a conditional discrimination: In both cases the listener has heard *the coffee is ready*, but the listener is less likely to pour the coffee after *I doubt* than after *I'm sure*.

Now consider the speaker. What has *I doubt* or *I'm sure* tacted? It can't be just the readiness of the coffee. It must be some property of the speaker's own tendency to say "The coffee is ready," and the relation of that statement to the actual state of the coffee. You can't use *I doubt* or *I'm sure* effectively unless you can discriminate your own behavior. In a situation in which you'd like to be able to say *The coffee is ready*, you must be able to tell whether it is appropriate to do so.

We don't tact everything we see, and conversely we sometimes respond as though tacting when the stimulus is absent. The qualifying autoclitic accompanying such verbal behavior is typically some form of the verbal response *no*. For example, we don't continuously say *The coffee isn't ready* under conditions of unready coffee. Instead, this verbal response occurs when circumstances set the occasion for saying *The coffee is ready* (e.g., the smell of coffee, the question *Is the coffee ready?*) even though that response would be inappropriate. If you do say *The coffee is ready* when the coffee in fact isn't ready, adding *Not!* makes a big difference.

Assertion, like negation, is also autoclitic, but the verb *is* serves many functions. Sometimes it specifies that it accompanies a tact (*This is a book*); sometimes it prescribes equivalences between verbal responses (*A human is a featherless biped*); and sometimes it specifies temporal properties (*It's cold now*). The particular function of *is* often depends on other verbal responses or, in other words, on context. Not only does it function as a conditional stimulus with respect to the effect of other verbal behavior, but its function may in turn be conditional on other verbal behavior.

Autoclitics can have quantitative as well as qualitative effects. Examples are *few*, *some* and *many*, and the plural forms of nouns and verbs. The effect of *often* in *This text is often misunderstood* can be paraphrased as a statement that it's often appropriate to say *This text is misunderstood*. We have no independent nonverbal means for characterizing these relations, but paraphrase is useful because it makes explicit the conditional relations among the components of the utterance. This is most obvious when autoclitics specify the listener's verbal behavior. For example, *vice versa* is conditional upon some verbal behavior that preceded and may be interpreted as an instruction to the listener to generate a new verbal response reversing the order of components in the original verbal response.

Descriptive autoclitics depend on discriminations of our own verbal behavior, and we know only a little about how they're learned or can be taught. We can illustrate this by considering what we'd have to do to teach an autoclitic to a pigeon. Suppose we arrange a series of discrimination trials ranging from easy to hard for a pigeon's key pecks and then try teaching the pigeon to report its certainty after each trial. We could add two new keys to the chamber, designating one as the *certain* key and the other as the *uncertain* key. If we were successful, the relation between the pigeon's report of its certainty and its response on the discrimination trial would be like that between the autoclitics *I'm sure* or *I'm not sure* and whatever it is we're sure or unsure about.

The trouble is that we have to know whether the pigeon is certain or uncertain in each discrimination trial before we can reinforce pecks on one or the other of the new keys appropriately. This is again the problem of teaching a tact of a private event. Perhaps we can base our reinforcers on some measure of the difficulty of the discriminative stimuli; another possibility is to base them on the latencies of the pigeon's responses in those trials (but see Blough, 1992). Our verbal community somehow taught us to tact our own certainty, so its teaching too was presumably based on public dimensions such as the properties of problems and whether we responded to them promptly or with hesitation. It would certainly be difficult and maybe even impossible to teach this autoclitic to a pigeon (cf. Smith *et al.*, 1995; Smith & Washburn, 2005), but the analogy is helpful because it forces us to be explicit about the details of these complex properties of our own behavior.

Our verbal behavior would be impossible without autoclitic processes. In saying *I recall that it rained yesterday* or *I read that it rained yesterday* or *I heard that it rained yesterday*, I specify the source of my verbal response, *It rained yesterday.* Other descriptive autoclitics may tact some feature of my reaction to the rest of what I'm saying, as in ***I'm sorry to report that*** *you missed the point* or ***I'm pleased to say that*** *you did very well on the exam.*

More important, descriptive autoclitics demonstrate that the analysis of verbal behavior is an issue of behavior and not logic. Saying *This is so* or *That's probable* or *It can't be* is verbal behavior with respect to other verbal behavior. To reduce such sentences to symbolic logic or to the mathematics of probability may be useful in solving problems of logic or mathematics, but that reduction eliminates a crucial feature of human language. Discriminations of our own behavior are prerequisites for what we call consciousness or self-awareness, and we owe those discriminations to the practices of the verbal community. "It is only through the gradual growth of a verbal community that the individual becomes 'conscious'" (Skinner, 1957, p. 140) and "It is because our behavior is important to others that it eventually becomes important to us." (Skinner, 1957, p. 314).

Chapter 23
Verbal Goverance

Among the labels that have been applied to behavior guided by verbal antecedents are verbally governed *and* rule-governed, *and among those applied to behavior that depends on nonverbal contingencies are* contingency-governed *or* contingency-controlled. Govern *is related to* cybernetics *via the Greek* kubernan, *to steer or pilot, which is probably related in turn to the Sanskrit* kubhan-, *dancing, maybe in a sense of adjusting one's steps.* Control *combines the Latin* contra-, *against, with* rotulus, *roll, which as a diminutive of* rota *implied a roll of paper as well as a wheel. Via the Indo-European root* ret-, *to run or to roll, it is related to* rotate *and* round, *and to* role *perhaps in the sense of taking a turn. In their etymologies,* govern *implies* guidance *whereas* control, *almost literally as working against movement, implies* constraint.

The Consequences of Verbal Operants
  The Mand
  Multiple Causation
Verbally Governed and Contingency Shaped Behavior
  Instructional Control
  Insensitivity to Contingencies
Shaping Verbal Behavior
  Effects of Shaped Verbal Behavior on Other Behavior
  The Development of Correspondences
Verbally Governed and Contingency-Shaped Verbal Behavior

With the exception of the tact, the verbal classes we've considered so far were primarily ones that related verbal responses to each other. We defined the formal classes of echoic behavior, transcription, textual behavior and dictation-taking in terms of one-to-one correspondences between arbitrary spoken and written units (phonemes, letters, words). The relations in intraverbal classes did not involve such correspondences, and those in autoclitic processes added coordinations some of which depended on discriminations of properties of our own verbal behavior. It was only in tacting that we began to connect verbal behavior to environmental events, but that connection was only to the nonverbal antecedents of verbal behavior and not to its consequences. To be functional, verbal behavior must be able to do things. The classes discussed so far are merely the raw materials that are combined in functional verbal behavior.

By itself, verbal behavior doesn't do things. It is effective through the mediation of other peo-

KEY TERMS: Multiple Causation; Verbally Governed Behavior, Contingency-Shaped Behavior, Instructional Control, Say-Do Correspondence; Intrinsic versus Extrinsic Consequences; Verbal Shaping.

ple. The mode isn't critical; you can tell someone what you want when shopping by asking for it or by pointing and you can also do so in writing, as when you place an order online. But mediation by others characterizes all social behavior, so we must add another proviso. The social contingencies that shape verbal behavior don't just create and maintain the conditions for speaking. They also create verbal repertories with a special property: In the ordinary give and take of everyday talk, as speakers become listeners and listeners become speakers, the behavior of the listener reinforces the behavior of the speaker. Some nonhuman behavior may minimally qualify, as when a horse is taught to turn in response to a touch of the reins to its neck (the turns then reinforce the trainer's behavior). But the horse and trainer differ crucially from a child and an adult teacher. For the latter only, verbal contingencies soon become reciprocal: The child learns to ask as well as answer, and to say *thank you* as well as *you're welcome*. Thus, in some respects all verbal cultures are mutual reinforcement societies (cf. Skinner, 1957, pp. 224-226).

Consider a simple verbal exchange: A says, *Hi*; B replies, *Hi, how are you?*; A continues, *I'm fine, thanks. And you?*; B answers *Okay*; and so on. The first reply of B is both a consequence of A's greeting and an occasion for A's continuation, but so is A's response with respect to B's antecedent and consequent verbal behavior. If there is any doubt about whether consequences that maintain verbal behavior operate here, just imagine A's or B's verbal behavior if the verbal behavior of the other simply ceased. Verbal behavior involves both listener behavior shaped by its effects on the speaker's behavior, and speaker behavior shaped by its effects on the listener's behavior. These reciprocities define verbal behavior. Verbal behavior is shaped and maintained by the practices of a verbal community, and this chapter considers some of the varied consequences that follow from those practices.

When we define verbal behavior in this way, by its function, we distinguish it from language (cf. Skinner, 1957, p. 461). A language is defined by structure and not by function. The definitions, spellings and pronunciations in dictionaries and the rules in a grammar book describe the standard structures of various verbal units in a language. In so doing, they summarize some structural properties of the practices of a verbal community. The verbal behavior of a speaker occurs in the context of those practices, but **those maintaining practices, language, must not be confused with what they maintain, which is verbal behavior.**

THE CONSEQUENCES OF VERBAL OPERANTS

Verbal behavior has consequences. As with all operant behavior, these consequences affect subsequent verbal responding. In a speech episode such as a simple two-person conversation, each person provides an audience for the other. Audiences are varied in their properties: We speak into telephones, write messages, or address large groups of people. Often the consequences for the speaker are simply what a listener says later. It doesn't require a laboratory experiment to demonstrate that a listener's response can maintain a speaker's talk. We tend to stop talking to people who don't react to what we say. To this extent, we can say that the listener's responses reinforce the speaker's verbal behavior. Simple consequences such as *yes* or *uh-huh* can reinforce verbal classes such as plural nouns or the substantive content of conversations (Greenspoon, 1955). One of the most general consequences of verbal behavior is that through it a speaker changes the behavior of a listener. Words are ways to get people to do things.

The consequences that serve as reinforcers of human verbal behavior are many and varied. Sometimes they are nonverbal – someone comes when called; sometimes they are verbal – someone answers a question. Sometimes they are fairly reliable – *Thank you* is often followed by *You're welcome*; sometimes they are not – not all requests are granted. Often they are very specific – answers to particular questions, taking appropriate action. But the shaping of verbal behavior, whether by natural or by artificial contingencies, will often involve generalized reinforcers. In addition, the tendency to speak may depend on some consequences while what is said may depend on others. One difficulty with experiments on verbal reinforcers was that they sometimes attempted to modify the substance of verbal behavior with the same consequences that were supposed to keep the speaker talking (cf. Azrin *et al.*, 1961; Greenspoon, 1955; Rosenfeld & Baer, 1970). Verbal behavior is maintained by varied consequences within interacting natural contingencies that can be thought of as complex nested schedules of reinforcement (Chapter 11). It's no wonder that we talk a lot.

The Mand

One obvious kind of consequence of verbal behavior is illustrated when we're given something we ask for. If a child says *milk* and receives a glass of milk, we might say that the milk reinforces the verbal response. The response needn't occur in the presence of the reinforcer. For example, a child may ask for milk even if a glass of milk isn't present. Verbal responses that specify their reinforcers have been called *mands* (Skinner, 1957); demands and commands, for example, specify what the listener is to do.

An analogue from animal research may be helpful. Assume a rat in a chamber with one lever that produces food and another that produces water. If the rat presses the first lever only when food-deprived and the second only when water-deprived, we could argue that the presses are, respectively, food requests and water requests. Although it wouldn't be essential, we could make the analogy more convincing by arranging signs that lit up for the experimenter when either lever was pressed, saying *Please give me food* and *Please give me water*. The rat's vocabulary is limited to two levers, but the relations between the presses and their consequences are similar to those between human verbal requests and their consequences.

Yet this account isn't quite satisfactory. For example, imagine a child who sees a new toy, learns it's called a woozle and then asks for a woozle even though asking for woozles could never have been reinforced before. As a category of verbal behavior, the mand can't consist of many separate response classes corresponding to each of the many consequences that could be manded. Rather, it must be a single class of responses in which a reinforcer is specified by the verbal responses that tact it in other circumstances; in other words, manding is a higher-order verbal operant. This mand is not a tact, but here it does function like a name (cf. Chapter 14). No such class existed in the rat example.

Within the class of mands, some subclasses specify stimuli – *May I have an apple?* – and some specify the behavior of the listener– *Please wait for me*; others called questions specify the listener's verbal behavior – *What's your name?* or *Do you have the time?* or *Did you see where I left my keys?* These classes may be further subdivided according to a variety of features. For example, we speak of a *prompt* when the appropriate verbal response is already known to the speaker – as in giving a hint to a child who is unable to solve a riddle – and a *probe* when it isn't known – as in a police interrogation. In everyday discourse, we also distinguish different kinds of mands on the basis of the contingencies they signal; for example, bribes specify the consequences of compliance whereas threats specify the consequences of noncompliance.

Multiple Causation

One problem of verbal behavior is that particular verbal topographies (e.g., words) can share in many different functions:

> we cannot tell from form alone into which class a response falls. *Fire* may be (1) a mand to a firing squad, (2) a tact to a conflagration, (3) an intraverbal response to the stimulus *Ready, aim...*, or (4) an echoic or (5) textual response to appropriate verbal stimuli. It is possible that formal properties of the vocal response, especially its intonation, may suggest one type of controlling variable, but an analysis cannot be achieved from such internal evidence alone. In order to classify

> behavior effectively, we must know the circumstances under which it is emitted. (Skinner, 1957, p. 186)

Any verbal topography potentially can take on any verbal function. Furthermore, probably no instance of verbal behavior is ever uniquely determined by just one variable. For example, when a child says *milk* in the presence of milk, this verbal response may function simultaneously as a mand and a tact (especially if the child is already naming). When two or more events act together to produce behavior, we speak of their combined effect as *multiple causation*. We considered it in the adduction examples of Chapter 13.

Humor often depends on multiple causation. For example, if the end of this paragraph is funny, that is mainly because of how it combines two different vocabularies with which we talk about things that we value. So now let's change the subject to the Olympics. Is there any truth to the rumor that there was once an Olympic team whose members were so proud of their gold medals that they had them bronzed?

Audiences provide especially good examples of multiple causation. Tacting is occasioned by the tacted stimulus, but we might not tact unless an audience is present, and an effective audience might or might not be a person. For example, we might say something to a pet even though the pet never answers. Pets have enough in common with human audiences. Not only are most reasonably animate but some make eye contact when you talk to them, and most dogs and cats respond differentially to different things we say. Our cat, Chesapeake or Chessie for short, weighs less than seven pounds, but except for his small size and the standard number of toes he has all the features of a Maine Coon cat. If I watch his ears I can see them turn when I say his name (unless he is sleeping, of course). I have never encountered any pet owners who claim not to talk to their pets. It should be no surprise that our talking often generalizes from human audiences to them.

What I say to my wife usually differs from what I say to Chessie. What we say given one audience may differ from what we say given another audience or given both together. If an absent-minded professor comes to class wearing a new suit with price tags still hanging from the jacket and the pants, a student might tact the price tags given another student as an audience, but only provided that the professor is not in a position to hear the tact too. If the student just whispers the tact, the causes of the whisper are the price tags and the other student and the professor's proximity in combination; take away any one and there will be no whispered tact. (By the way, as an undergraduate I had a philosophy professor who once did that. In those days, when some faculty members routinely smoked during their lectures, I also had a math professor who usually held a cigarette in one hand and a piece of chalk in the other; seeing him try to puff on the chalk gave us something to whisper about, but we didn't have to say a word when he tried to write on the chalkboard with his lit cigarette.)

As is appropriate in an analysis, we'll usually treat causes one at a time, but once we start looking for the multiple causation of behavior we find it virtually everywhere. Creating examples of multiple causation in verbal behavior is a useful exercise. I've often used such an exercise in class. If I ask you to invent some instance of verbal behavior in a fast-food restaurant that includes at least five different causes, what could you come up with? Sometimes I've gotten answers that string several pieces together: a student might offer saying *orange juice* on seeing the word on the menu as textual behavior and then manding orange juice from the server, and so on. But that won't do it, because reading the menu and placing an order are two separate utterances and this exercise is asking for just one.

Better would be to suggest standing at the counter and saying to the server *I'll have the chicken burrito* when (1) you've just heard someone say *My chicken burrito was really good!* (echoic *chicken burrito* to *chicken burrito*), and (2) you just saw it on the menu on the wall behind the server (written word occasioning textual behavior), and (3) you just saw another customer with a burrito platter (the burrito occasioning a tact), and (4) your saying it is a mand (it is lunch time, and this burrito will surely be a reinforcer), and (5) your server is there as your audience. All of these come together to make it highly likely that you will say *I'll have the chicken burrito*. Some contribute more than others. You might still place the order even if you hadn't overheard the other customer or seen another's dish, but you wouldn't place the order at all if no server was standing behind the counter.

Not only does this example illustrate the ubiquity of multiple causation; it also shows why we mustn't think of the analysis of verbal behavior as a task of labeling the various bits and pieces of a verbal instance with the names of various functions or verbal classes. It is fairly unambiguous to call the server your audience, but if you're standing in line with a friend at a busy time, your friend and the other patrons are part of your audience too. We cannot disentangle the echoic and textual and tact and mand components of what you said, so why try to say that any part of your sentence was a mand or a tact? When we try to interpret the verbal relations in everyday situations such as these, the best we can do is to make plausible guesses about their various contributions. That shouldn't bother us. We don't expect physicists to tell us where particular leaves will land during a wind storm or biochemists to tell us what created the smells emanating from a trash bin sitting in the sun while it awaits collection, so why should we expect to do any better in our post hoc interpretations of specific real-life situations?

VERBALLY GOVERNED AND CONTINGENCY-SHAPED BEHAVIOR

Verbal behavior can have either verbal or nonverbal consequences, but in one way or another the consequence is usually a change in the listener's behavior. For example, if you tell a friend who is about to go outdoors that it is going to rain, your friend may be more likely to take an umbrella.

The verbal community maintains correspondences between verbal behavior and environmental events. The listener can act on the speaker's verbal behavior only if such correspondences are consistent. If the speaker's verbal behavior is occasioned by environmental events inaccessible to the listener, it may become a potent discriminative stimulus for the listener's behavior. For example, a listener's response to the tact *fire* may have important consequences even if only the speaker has seen the fire. Through the verbal behavior of others, we can respond indirectly to events distant in space or in time (the tale of the boy who cried wolf illustrates how control by the speaker's verbal behavior may weaken if the speaker tacts unreliably).

Sometimes what people do depends on what they are told to do; people often follow instructions. Behavior mainly determined by verbal antecedents is called *verbally governed* behavior (sometimes also *rule-governed* behavior, but that terminology is too often misunderstood); its properties differ from those of *contingency-governed* or *contingency-shaped* behavior, behavior that's been shaped by its consequences (Skinner, 1969). Some instructions affect nonverbal behavior (*Come here*, *Sit down*, *Go away*); others affect verbal behavior itself (*Tell me a story*, *Ask nicely*, *Be quiet*).

Instructional Control

The most general function of language is instruction; we tell each other what to do and what to say. Language isn't an instrument of reason or a vehicle of truth; those properties are only corollaries of its primary function, changing the listener's behavior. Orders are given, advice is offered, laws are enacted and so on; each case involves instructional control. This is easiest to overlook when the instructed behavior is itself verbal. A script is a set of instructions to an actor and a text is a set of instructions to a reader. In both cases the instruction specifies verbal behavior or, in other words, what is to be said. The instructor who defines a term, for example, specifies the conditions under which the term and its definition will be appropriate in the student's future verbal behavior (especially in those contexts called exams).

An important characteristic of instruction is that it substitutes verbal antecedents for natural contingencies, as when a parent tells a child *Don't touch the stove or you'll get burned.* This property of instruction has far-reaching implications. Instructions can change the listener's behavior in situations in which natural consequences by themselves are ineffective or are effective only slowly. If we invite friends for a visit, for example, we give them directions rather than letting them search for our place on their own (these days, of course, GPS navigation may substitute for our directions, but that too provides instructions).

Notice that these verbal antecedents aren't necessarily discriminative stimuli. We might act appropriately when we see a sign saying *Hot surface! Don't touch!* In that case the sign would qualify as a discriminative stimulus. But the instruction isn't present to function as a discriminative stimulus when the child avoids the hot stove today after a warning about it yesterday. Many instructions alter the functions of other stimuli instead of functioning as discriminative stimuli (Schlinger & Blakely, 1987). For example, if the warning about the stove is effective, it changes the stove to a discriminative stimulus correlated with an aversive event, getting burned, but the warning itself, no longer present, is not a discriminative stimulus.

Contingencies operate for following instructions. To the extent that instruction-following is characterized by correspondences between the instruction and the listener's behavior and thus is more than the following of particular instructions, it's another higher-order class of behavior held together by common contingencies. The higher-order contingencies that operate on the following of orders in general, which are usually social and verbal, are different from those that operate on the following of a particular order, which may be completely nonverbal. For example, following orders in the military is a product of extensive and powerful social contingencies; they are often artificial applications of aversive control. But there is nothing artificial about the immediate consequences of obeying a command to attack a heavily armed position across an open field of fire in actual combat.

Varied social contingencies maintain instruction-following. Among the consequences that might reinforce the granting of a request might be such immediate consequences as changes in the requester's expression or posture or what the requester then says, or more remote consequences such as that the granter thereby avoids hurting the requester's feelings, or puts the requester in debt so that the favor will probably later be returned, or gets the requester to leave, or annoys someone else who didn't want the request granted, and so on, and so on. The social possibilities, whether subtle or patently obvious, are endless.

Sometimes the contingencies that maintain instruction-following depend on the relation between verbal formulations and nonverbal contingencies, as when someone successfully makes a repair by following a service manual, or as when someone avoids injury by acting on a warning. The term *pliance* has been suggested for instruction-following that depends on social contingencies – *Please do this for me as a favor,* – and the term *tracking* for instructional contingencies that involve correspondences between verbal behavior and environmental events – *You have to put the battery in with the plus side facing this way for the flashlight to work* (Zettle & Hayes, 1982). Some instructions function as establishing operations, *augmenting* the effectiveness of reinforcers, as when an advertisement that shows pictures of food makes eating more likely, or as when a sales pitch that describes special features of a car makes buying the car more likely (cf. Hayes, Zettle, & Rosenfarb, 1989, p. 206).

Consider an example of verbally governed behavior. Without instructions someone types with the index fingers, one letter at a time, in the method called hunt-and-peck. This is faster for the novice typist than touch typing, in which each finger has its own resting position at the keyboard. The immediate consequences of the two typing methods favor the former: With the familiar hunt-and-peck method, the job will get done more quickly. But in learning touch typing, the long-term consequences of following instructions to place the fingers appropriately and to type each letter with a particular finger will eventually outweigh the short-term advantages of the hunt-and-peck system: once it is mastered, touch typing is much faster than hunt-and-peck typing. What is going on involves much more than just learning a particular method of typing. To follow instructions successfully, the learner must ignore the natural consequences, in this case that the text is at first produced more slowly by touch typing than by hunt-and-peck typing.

The student who is a hunt-and-peck typist may find it hard to switch over to touch typing, especially when a paper is due. Under this time pressure, the slowness of novice touch typing cannot compete with hunt-and-peck typing (we'd probably say that the student who sticks with touch typing under these circumstances exhibits self-control: cf. the contingencies in Chapter 16). I was once a hunt-and-peck typist, and in my student days I got around the problem by making the transition a gradual one, adding just one or two letters at a time to the proper placement for touch typing while continuing hunt-and-peck for the rest . I introduced new letters added only when the touch typing of the earlier ones was well differentiated. I found that it was important always to return to the base finger placement, with the left index finger on the F and the right on the J. Despite the tactile marking on the F and J keys of computer keyboards, I still occasionally make errors by placing my fingers incorrectly on the keyboard, usually after I've moved my hand away to use the computer mouse. Probably I should switch to an ergonomic keyboard, which makes misplaced finger positioning unlikely (this example illustrates that giving an instruction to oneself is sometimes less effective than an instruction that has come from someone else; later we'll address some of the variables that alter the effectiveness of instructions).

Because of the practical advantages of instruction, the verbal community shapes the behavior of following instructions across a substantial range of activities throughout a substantial portion of our lives. This can happen only if the contingencies that maintain instruction-following are more potent than the natural contingencies against which they're pitted (we seldom bother to ask people to do things they'd do on their own anyway). Thus, instructions may begin to override natural contingencies: People then do things when told to do them that they would never do if just the natural contingencies operated.

A major achievement of human verbal behavior is that it allows behavior to be controlled by descriptions of contingencies, in the verbal behavior of others, as well as by direct contact with the contingencies themselves. As we have seen, an instruction as a verbal input can produce a range of outputs, either verbal or nonverbal, in the various things that the listener may then do. This corresponds to what evolutionary biologists have called weak linkage, as when "different kinds of nerve cells with different receptors and neurotransmitters can be produced within an organism using the same basic cellular design" (Kirschner & Gerhart, 2005, p. 139). The nervous system mediates between the sensory end, which can involve stimuli along many different dimensions, and the response end, which can involve varied configurations of skeletal and autonomic behavior; the nervous system is a mediator, in that it serves in place of a direct link between the two systems. In a sense the stimulus tells the muscles or glands what to do, but it does so in a different way than do the processes of selection during development. Biologists distinguish between instructive signaling, in which the stimulus determines the response, and permissive signaling, in which outcomes are occasioned and selected; that distinction at the physiological level has much in common with the distinction between verbally governed and contingency-shaped behavior.

But the advantages of this unique property of verbal behavior, the control of behavior by descriptions of contingencies, are accompanied by special problems. For example, a history of following instructions may make individuals susceptible to verbal control by authority figures (e.g., Milgram, 1963). In addition to the abuses that can arise when people just follow orders, instructions can create problems in subtler ways (cf. Chase, 1938; Korzybski, 1941).

Insensitivity to Contingencies

Consider the simple task of pressing a telegraph key, with presses earning money according to various schedules of reinforcement. When key pressing is established in humans by instructions instead of by shaping, instructed performances are typically insensitive to the schedule contingencies, whereas shaped performances often are not (Matthews *et al.*, 1977). Schedule effects observed with uninstructed responding in humans (e.g., higher rates maintained by ratio schedules than by yoked interval schedules; cf. Chapter 15) don't occur reliably when responding is instructed. Just telling a human in an experimental setting to press a button produces persistent responding that is insensitive to its consequences (cf. Lowe, 1980; Shimoff, Catania, & Matthews, 1981). Such insensitivities have been observed across a range of schedule contingencies (e.g., Baron & Leinenweber, 1995; but cf. Svartdal, 1992).

This property of instructed performances is relevant to those aspects of skill that we sometimes say can't be taught. Skilled performances are those in which behavior is sensitive to its consequences from moment to moment, as when a woodcarver adjusts to the changing patterns of grain in a woodblock or a quarterback anticipates the movements of an opposing player or a dancer accommodates to slight deviations in the steps of a partner. If verbally governed behavior is likely to be insensitive to its consequences, then skilled behavior must be contingency-shaped instead. We must learn by doing in such cases; instructions cannot substitute for the subtleties of direct contact with contingencies.

In Chapter 9 we discussed the skills involved in shaping. In learning to shape, it makes a difference whether we learn through instruction or through nonverbal contingencies. The properties of verbally governed shaping differ from those of shaping that has been learned through the contingencies of actually interacting with a behaving organism. If you have gotten good at shaping that way, your skills cannot easily be expressed as rules, and if you can't say what the rules are you can't offer those rules as instructions. In such cases, differences between verbally governed and contingency-shaped skills are inevitable. Another complication is that we sometimes generate rules for ourselves. To the extent that we behave consistently with those rules, they may get in the way and make us insensitive to some of the contingencies that would otherwise shape our behavior.

We usually don't want others to do what we say simply because we say it. A parent or teacher who gives instructions to a child might prefer but can't be confident that the natural contingencies will eventually take over the relevant behavior and make instructions unnecessary. For example, a reason for telling a child to put overshoes on before going out to play in the snow is that the instruction may keep the child from coming home with cold wet feet. If the child always obeys the instruction, the natural contingencies will never act on the child's behavior; if the child disobeys the instructions, the aversive consequences that follow from snow and unprotected feet may enhance the control by instructions on future occasions. Thus, if we try to teach by telling others what to do, we may reduce the likelihood that they'll learn from the consequences of their own behavior. There is no easy solution to this dilemma. We must always choose between the immediacy and convenience of verbal instructions and their longer-term effects on the learner's sensitivity to the consequences of behavior.

Cases of insensitivity to contingencies produced by verbal behavior can be found across a wide range of settings. For example, people learn the rules of an artificial language more effectively by working only with sample sentences than by working with those sentences plus statements of the grammatical rules (Reber, 1976); right-left confusions are less likely in spatial tasks that don't involve words than in those that do (e.g., Maki, 1979); individuals can't ignore false labeling even if they did the labeling themselves (Rozin, Millman, & Nemeroff, 1986); remembering faces and some other classes of stimuli that are difficult to put into words is impaired by naming the stimuli (Schooler & Engstler-Schooler, 1990). Some differences in the effects of intrinsic versus extrinsic reinforcers, discussed in Chapter 14 in connection with the so-called hidden costs of reward, depend on whether or not rewards are instructed, in the sense that students are told it is important to earn them. Perhaps it is appropriate to conclude that some things are better left unsaid.

One contribution to the effectiveness of verbal governance may be that verbal communities arrange contingencies to maintain correspondences between what we say and what we do. The correspondence can operate in either direction: If we do something we can say we've done it, and if we promise to do something we can do it (e.g., Paniagua & Baer, 1982; Risley & Hart, 1968; Rogers-Warren & Baer, 1976). To the extent that the verbal community arranges contingencies for these correspondences, **we can change behavior not only by instructing the behavior but also by shaping what is said about it.** If both saying and correspondences between saying and doing are reinforced, doing may follow. Through such contingencies, one's own verbal behavior may become effective as an instructional stimulus (e.g., Jaynes, 1976; Lovaas, 1964), though presumably within limits (e.g., Baer, Detrich, & Weninger, 1988; Hübner, Austin, & Miguel, 2008; Matthews, Shimoff, & Catania, 1987a).

Human schedule performances were sometimes substantially different from those maintained with nonhumans, and as reports of these performances became more common (e.g., Weiner, 1983), some of my colleagues and I took an interest in exploring those performances. We arranged various VR and VI button-pressing tasks with UMBC undergraduates as our participants, and we too found that humans sometimes pressed at high rates without regard to whether the schedule was VR or VI, that students who followed instructions carefully were more likely to respond indiscriminately to these schedule contingencies whereas those who didn't follow instructions well and whose button-pressing had to be shaped were sensitive to them, and that performances were sometimes different depending on instructions or reinforcers, which were usually points exchangeable for money (e.g., Matthews, Catania, & Shimoff, 1985; Matthews *et al.*, 1977; Shimoff & Catania, 1998; Shimoff *et al.*, 1981). As a further complication, with procedures like these it cannot be taken for granted that the consequences identified as reinforcers are actually functioning that way.

We knew that nonhumans were exquisitively sensitive to these sorts of schedule contingencies, and yet again and again we found that humans were instead insensitive to the consequences of their own behavior. Shouldn't sensitivity to one's own behavior be a fundamental skill across all behaving organisms? If so, was something masking our sensitivity? Our task was cut out for us. With human behavior so inferior in this respect to the behavior of pigeons and rats, could we restore human dignity by showing what was getting in the way?

SHAPING VERBAL BEHAVIOR

Let's call him Marty, though that was not his name. He was in his late teens or early twenties and most of his talk consisted of what psychiatrists call word salad: semigrammatical though unpredictable and not particularly meaningful, and sometimes reminiscent of the examples given by linguists (e.g., Chomsky's "Sleeping green ideas dream furiously"). He was tall and broad-shouldered, and I was glad that I didn't have to deal with him alone when he got upset with his relatives on a visiting day. I was a college student spending the summer of 1956 substituting for vacationing ward attendants on a locked psychiatric ward. It was part of the Bronx Municipal Hospital Center, a new teaching hospital that was to be affiliated with the Albert Einstein College of Medicine once the latter was completed and began admitting students. By then I had declared a major in psychology. Given our interests, I and another summer temp, an undergraduate pre-med, were allowed to review charts and to sit in on discussions of patients.

The psychiatric staff was mainly psychoanalytic in orientation and the unit was essentially a triage ward. After admission, patients were supposed to be treated and diagnosed and within a few weeks discharged or referred to the outpatient unit for continuing treatment or committed to a state mental hospital. Some contingencies were fairly obvious. The somewhat unattractive pre-teen girl who mostly cried in a corner was quickly sent off to the state hospital but the perky and attractive young woman with some social skills and a southern accent was still around as a patient when I returned there for another summer of work in 1957 (had she been an actress she would have been a good candidate for the role of Blanche Dubois in *Streetcar Named Desire*). The elderly Jewish man who had survived the Holocaust in Europe during World War II and who had been in the United States for only a decade or so was suspicious of everyone, hiding things and talking little; he was declared paranoid. One young man, recently discharged from the military and diagnosed with personality disorder, was a captivating raconteur, but within a couple of days he exhausted his entire repertoire of stories and began repeating them again, in the same order and without regard to whether you'd already heard them; his chart showed that his IQ was in the 60s. The psychiatrists argued at length over whether one male patient was schizophrenic or manic-depressive (we would now say bipolar); they were so evenly divided that the diagnosis was not resolved while I was there (cf. Rosenhan, 1973).

Ward attendants could not give meds, so our main functions were to interact with the patients, keeping them out of trouble and escorting them when necessary to activities in other parts of the hospital; those ranged from x-rays to rehab to electroshock therapy. This meant that it was usually easy to elect to spend time with one patient or another. Verbal shaping, as in the reinforcement of plural nouns by Greenspoon (1955), had been a topic of discussion in my psychology courses during the preceding academic year. One part of that story was the trouble that Greenspoon had in getting his work published.

Given Marty's word salad I decided to see whether I could make verbal shaping work. Usually I'd find him standing somewhere verbalizing. I sat nearby and showed interest whenever something he said included *I* or *me*; otherwise I looked away and ignored him. Over the course of two or three weeks not only had those words become far more common in his talk; when I was around he began speaking to me in full and coherent sentences, usually in first person and about himself. The psychiatrists hadn't gotten much from him directly in their one-on-one therapy sessions, but what I learned from him once he began talking coherently turned out to be consistent with what was known about him from his parents;. I wish now that I had taken copious notes, so that I could have provided a detailed account including sample utterances and the time course of the shaping. But at least I discovered that the bit of attention I gave when he uttered those first-person words was reinforcing enough to change his talk. I had heard that the literature on verbal shaping was controversial, but any doubts I might have had vanished once I had Marty talking to me about himself. And verbal shaping would come in handy later too.

Effects of Shaped Verbal Behavior on Other Behavior

In one experiment (Catania, Matthews, & Shimoff, 1982), students' presses on two buttons occasionally produced points exchangeable for money; a representative apparatus is shown in Figure 23-1. When a light above the left button was lit, a random-ratio (RR) schedule operated for presses on that button; when one above the right button was lit, a random-interval (RI) schedule operated for its presses. Between alternations of the two schedules, the students filled out guess sheets that included sentences to be completed for the left and right button of the form: *The way to produce points with this button is to....* Their guesses were shaped with differential points worth money. In shaping guesses, the experimenter assigned points to each guess, writing point values next to each one and passing the sheet through a partition to the student.

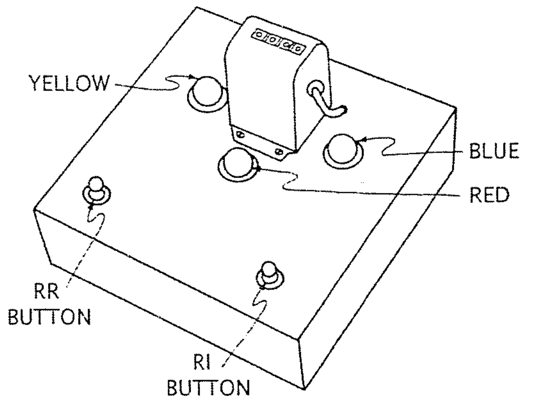

Figure 23-1 A button-pressing apparatus for studies of human behavior maintained by reinforcement schedules. In this example, presses on the left button are reinforced according to a VR schedule when the yellow lamp is lit and those on the right button are reinforced according to a VI schedule when the blue lamp is lit; the red lamp is used to signal other events (e.g., providing a verbal report). Points on the counter exchangeable for money serve as the reinforcer.

When the guess shaped for one button was *press quickly* and that for the other was *press slowly*, response rates on the two buttons changed in corresponding directions, without regard to schedule contingencies. Thus, shaping *press slowly* for the left button and *press quickly* for the right button produced relatively low RR rates and relatively high RI rates, opposite to those usually produced by these schedule contingencies (cf. Chapter 15). What the students said about their responding was a more powerful determinant of what they did than the consequences of their responding (e.g., students who slowed down on the RR button lost points they would have earned if they'd responded more rapidly).

In a procedure in which students were told what to guess, however, correspondences between the guesses and the response rates were inconsistent; sometimes guessing *press fast* was accompanied by fast pressing and guessing *press slow* by slow pressing, but sometimes these guesses were accompanied by equal response rates on the two buttons or by rates that differed in the opposite direction.

Verbal shaping involves treating successive verbal responses as varying along semantic dimensions, but judgments of which ones are closer or farther from the behavior to be shaped can be tricky. It is easy for us to pick out words related to *fast* or *slow* and *time* or *number*, but a student who comes up with a guess like *four fast presses then three slow ones* might just vary those two numbers on all trials thereafter. Because the possibilities for varying the numbers are unlimited, the student may be caught in a sort of verbal trap, in which our further attempts to shape simpler guesses will be unsuccessful. Furthermore, it makes a difference whether what we shape is a description of behavior or a description of the contingencies operating for that behavior.

Figure 23-2 shows data from a variation on these procedures (Catania *et al.*, 1989c). Again, an RR schedule operated for presses on the left button (L) and an RI schedule operated for presses on the right button (R). This time the guesses that were shaped were descriptions of contingencies rather than descriptions of performances. The sentences to be completed for the left and right buttons were of the form: *The computer will let your press produce a point depending on....* The verbal shaping successfully produced guesses that were variations of *number of presses* for the left button contingency and variations of *time intervals* for the right button contingency. The progress of shaping is shown by the shaded areas; shaping was complete when the student's guesses were earning the maximum of 18 points per guess period. The shaped contingency descriptions were accurate, but they were not accompanied by substantial differences in the RR and RI rates of button pressing.

Whether descriptions of contingencies produce corresponding changes in performance depends on other verbal behavior. For example, one student, correctly identifying two schedules as RR and RI, might go on to say that point deliveries increase with higher RR rates but not with higher RI rates. Another, also correctly identifying the two schedules, might instead go on to say that because points in both are unpredictable, point deliveries are unaffected by pressing rates. In most cases the

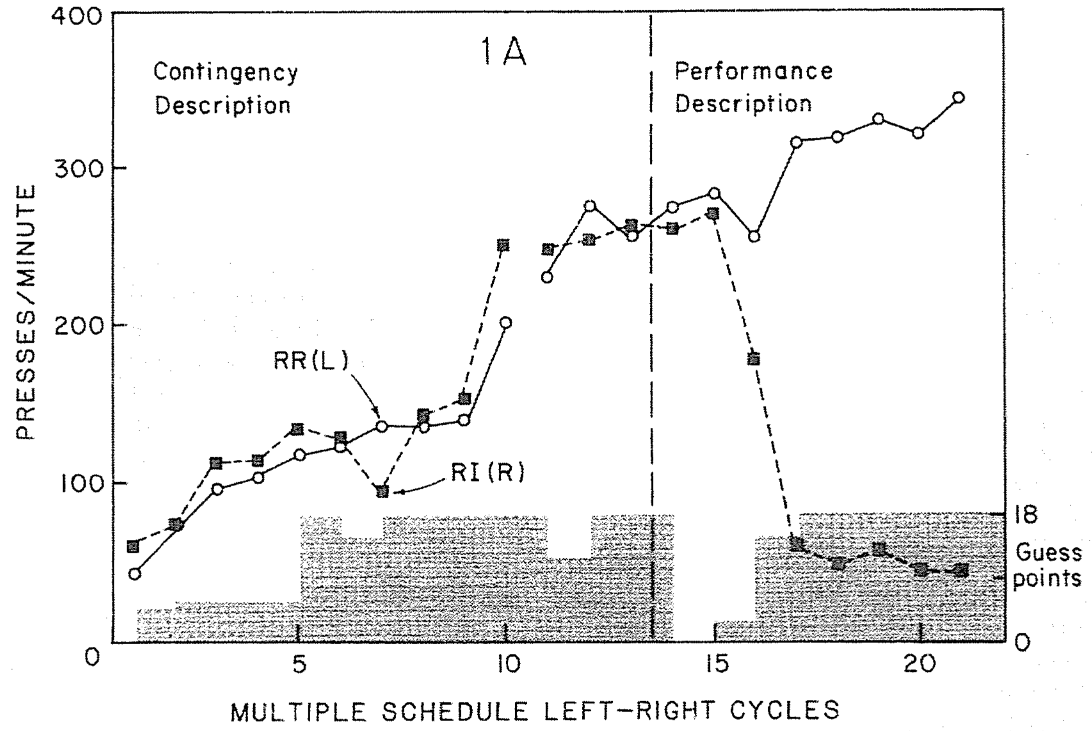

Figure 23-2 Left (L) and right (R) rates over 3-min cycles of multiple random-interval (RR) random-ratio (RI) schedules of point delivery for a student's button presses. Shaded areas show point deliveries for verbal behavior (guesses) during shaping of contingency descriptions and, right of the dashed vertical line, of performance descriptions. An interruption between sessions is shown by the unconnected points. The shaping of contingency descriptions did not produce differential response rates, but the shaping of performance descriptions did. (Adapted from Catania, Shimoff, & Matthews, 1989c, Figure 1)

first student but not the second will show rate differences appropriate to the schedules. In any case, descriptions of what you do in an environment differ from descriptions of how that environment works.

Verbal shaping was quickly effective when it switched from targeting contingency descriptions to targeting performance descriptions (right of the dashed vertical line in Figure 15-1); it produced variations on *press fast* for left RR button guesses and variations on *press slow* for right RI button guesses. The rates of button pressing diverged as these performance guesses became more consistent over cycles. Unlike procedures that merely sample verbal reports during operant performance, this kind of procedure identifies the direction of effect in the relation between verbal and nonverbal behavior, because we know which came first. The verbal behavior changed during shaping; then came the change in response rates. In different circumstances, of course, the direction can go the other way, as when students whose pressing is fast or slow then accurately describe their own behavior.

The Development of Correspondences

We've already suggested that correspondences between shaped verbal behavior and relevant nonverbal behavior such as those illustrated in Figure 23-2 come about because verbal communities differentially reinforce consistencies relating what we say to what we do. These effects begin early and can be demonstrated with children. For example, one experiment with four- to six-year-olds arranged reinforcing consequences consisting of the successive lighting of the lamps in an 8-lamp column; when the complete column was lit, the child earned a present, a small box later exchangeable for pictures to be pasted in the child's scrapbook or a toy that the child could select from a treasure chest. The child occasionally talked with a Garfield hand puppet, who appeared from time to time through a small curtained aperture (Catania, Lowe, & Horne, 1990a).

Figure 23-3 illustrates the experimental setup. When the top window showed a star, presses there worked according to a random-ratio (RR) schedule; when the bottom window showed a tree, presses there worked according to a random-interval (RI) schedule. As is typical only for verbal humans, the two schedules did not produce consistently different response rates even at this age.

After some sessions, Garfield appeared between schedule components and began talking to the child: e.g., *This game looks like fun. Can you teach me how it works?* The reinforcers used to shape what the child said about the performance included lights in the column and enthusiastic reactions from Garfield: e.g., *That's what I'll try when I play the game.* The child was never told what to say or what to do, but the verbal shaping procedure had to be adapted to each child's vocabulary. Fast and slow typically mean different things to children than to adult experimenters. For example, children will usually say fast or slow for how quickly their hands move rather than for how often they press. It was therefore usually more effective to shape alternatives such as *press a lot without stopping* and *press and wait.*

Although performance was more variable than with adults, successful shaping of verbal behavior was typically accompanied by corresponding rates of pressing. As with adults, shaped verbal behavior produced corresponding rate changes even when those changes were opposed to the usual effects of the schedule contingencies, as when the child who came to say *With the star you press and wait* responded slowly during star even if that schedule was RR. As another demonstration of the primacy of the shaped verbal behavior, a reversal of one child's verbal behavior was accompanied by a reversal of response rates, as illustrated in Figure 23-4.

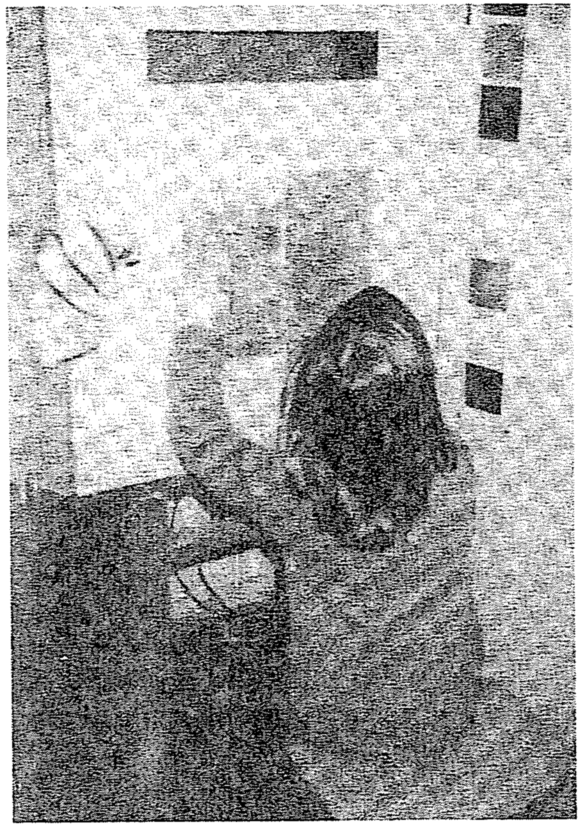

Figure 23-3 Experimental arrangement for a study of the effects of verbal shaping on the schedule performances of children. When a star is projected on the top screen, presses there are reinforced according to an RR schedule; when a tree is projected on the bottom screen, presses there are reinforced according to an RI schedule. This child is pressing the star on the top screen. Garfield, a hand puppet, looks on at the left. Details in text. (Adapted from Catania, Lowe, & Horne, 1990a)

Other developmental changes in verbal behavior are presumably relevant. The characteristic nonhuman fixed-interval (FI) performance is an increasing rate as time passes within each interval. The performance of verbally competent human adults, however, typically consists of one of two patterns:

either a high steady rate or a long pause about equal to the FI. At what point do these human performances emerge? The schedule performances of infants and older children have been studied using simple responses such as button-presses producing simple consequences such as snacks or opportunities to listen to music (Bentall, Lowe & Beasty, 1985; Lowe, Beasty & Bentall, 1983)

For children under 2 years of age, performance is completely contingency-governed, in the sense that it is not distinguishable from performances of nonverbal organisms such as pigeons or rats. For children between about 2 and about 4 or 5 years of age, performances are transitional, with variable properties that seem related both to contingency-governed and to verbally governed behavior. Children older than 5 years typically produce performances with the relatively stereotyped properties that characterize adult rule-governed performances. Verbal behavior is involved in the development of the verbally governed performances: intermediate-age children who have not shown verbally governed behavior on their own with these schedules do so when they are given some simple verbal instructions about the performance (Bentall & Lowe, 1987).

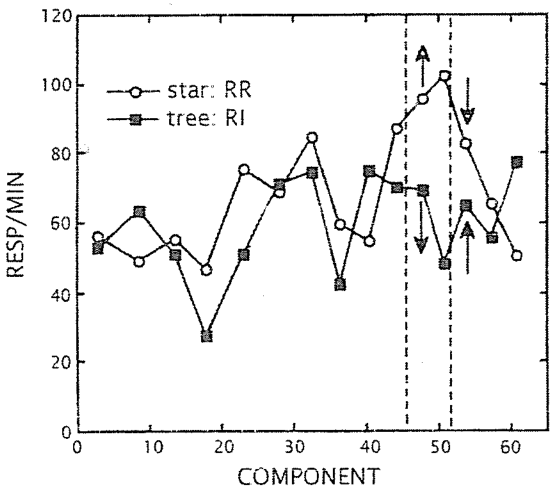

Figure 23-4 Response rates during multiple RR RI components for a five-year two-month old boy, illustrating the effects of a reversal of shaped verbal behavior. At the first dashed line, verbal behavior consistent with the contingencies was shaped (star: *press a lot without stopping*; tree: *press and wait*) and response rates followed. At the second dashed line, the verbal shaping was reversed and response rates also reversed. (Adapted from Catania, Lowe, & Horne, 1990a)

VERBALLY GOVERNED AND CONTINGENCY-SHAPED VERBAL BEHAVIOR

The shaping of verbal behavior is a potent technique for changing human behavior, especially given that the distinction between verbally governed and contingency-governed behavior is relevant to verbal as well as nonverbal behavior (Catania, Matthews, & Shimoff, 1990b). Verbal behavior that is shaped or contingency-governed is, like nonverbal shaped behavior, sensitive to its consequences, but it is also accompanied by corresponding nonverbal behavior: If what we say is shaped, we do what we say. On the other hand, verbal behavior that is instructed or verbally governed is, like nonverbal instructed behavior, relatively insensitive to its consequences, but it is less reliably accompanied by corresponding nonverbal behavior: If we are told what to say, what we do doesn't necessarily follow from what we say even if we reliably say what we were told to say.

The practical implication is that it is often easier to change human behavior by shaping what someone says than by shaping what someone does. If you'd like someone to come up to your place, you probably won't try to differentially reinforce successively closer approaches to your place. Differentially reinforcing successively closer approximations to *Can I come up to your place?* is much more likely to be effective.

Human nonverbal behavior is often verbally governed, but human verbal behavior is usually contingency-shaped. Thus, the therapist may sometimes be effective simply by shaping a client's talk (Truax, 1966). Therapies that invoke cognitive behavior modification or cognitive efficacy are said to modify the client's behavior by changing the client's cognitions, but everything we know about these kinds of procedures suggests that the critical change is in the client's verbal behavior, either through instructions or through verbal shaping. In cases like these the verbal shaping is more likely to be incidental than deliberate and may not even be recognized by those doing the shaping. Therapies like these are sometimes effective, but probably for reasons other than those claimed (cf. Bandura, Adams, & Beyer, 1977; Catania, 1995a; Chadwick *et al.*, 1994; Lowe & Chadwick, 1990). Might it be that some behavior pathologies, (e.g., compulsions) are mainly problems of verbal governance? Might hypnosis be viewed as a powerful form of verbal governance that restricts the operation of other classes of behavior (Hilgard, 1986)?

We don't often talk about the variables that determine our own verbal behavior. Our verbal behavior itself may be more likely to be free of verbal governance than other nonverbal varieties of human behavior because the everyday vocabulary of verbal behavior does not lend itself well to our talking about those variables. The shaping of verbal behavior is now well established in the experimental literature, but we've only scratched the surface of its possible applications.

One problem with shaping the content of verbal behavior is that different contingencies might be needed for maintaining the behavior than for shaping particular instances. In shaping the content of conversation, for example, the experimenter must keep the conversation going while also engaged in shaping. But if continued verbal production can be maintained, as with an automatic writing task (Skinner, 1934a), feedback can be scheduled intermittently (cf. Catania & Shimoff, 1998). If that feedback was contingent on, say, plural nouns according to an interval schedule and on first-person pronouns according to a ratio schedule, would these verbal units be sensitive to these schedule contingencies in the same way as the nonverbal responses of nonverbal organisms?

We'll explore some implications further in Chapter 24, but it may be worth noting here that verbal shaping is also relevant to cultural practices, ranging from positive ones such as honesty and courtesy to negative ones such as prejudice and hate speech. Chapter 11 distinguished discrimination from prejudice on the grounds that the former is a product of direct exposure to contingencies whereas the latter has been taught. Verbal shaping shows us how such practices can be initiated and maintained within verbal communities. If those who speak tolerantly are ignored or even disrespected while those who speak with prejudice get lots of attention, it should come as no surprise if confrontational speech sometimes overcomes civil discourse. The talk of community leaders and educators can have far-reaching effects. For better or worse, the politician's bully pulpit can make a difference.

Our culture puts a large premium on freedom of speech, but that freedom should carry with it a heavy responsibility for recognizing the power of words. Much political behavior is built around getting voters to say *I'll vote for Candidate X* on the assumption that if voters say that they'll behave accordingly once they get into the voting booth. One way to get people to say things like *I'll vote for you* is to engage them in discussion and debate; another is to bombard them through the media with verbal stimuli that occasion echoic and intraverbal and other relatively superficial sorts of responses. The reciprocal relations allowed by the internet may favor the former; the mass media operate mostly in one direction and therefore vastly favor the latter. Neither case allows much room for deliberation and critical thinking. If those with large financial resources are more capable of making use of the mass media that than those whose resources are limited, should we be surprised if these superficial ways of modifying verbal behavior can swing elections? If only money talks in politics, can any speech be said to be free?

One function of education should be to teach judgments about the sources of verbal behavior and its effects. "Human history becomes more and more a race between education and catastrophe" (Wells, 1920, Chapter 41). In education, we sometimes teach by shaping what our students say through questions and discussion. More often we teach not by shaping but by instruction; in lectures, students are in effect told what to say on exams. If the courses include no direct contact with a subject matter, the former type of teaching will probably be more likely than the latter to affect the student's interaction with the subject matter outside the classroom. Our conclusion is worth repeating: **It may be easier to change human behavior by shaping what someone says than by shaping what someone does.** If reinforcement sometimes appears not to work very well on human nonverbal behavior, we might be tempted to discount it as applying to human behavior at all. But if, because of the power of instruction-following as a higher-order class, the effects of reinforcement are stronger when applied to human verbal behavior than when applied to human nonverbal behavior, it would be foolish of us to ignore it. The phenomenon can be put to good use or to bad, and the best defense against its misuse is to learn as much as we can about how it works.

Chapter 24
Verbal Function: Coordinations among Classes

Our language words have diverse sources. Verbal, *through Latin, and* word, *through Old English, are derived from a common Indo-European root,* wer-, *to speak. The Germanic* spek *or* sprek-, *from which comes the German* die Sprache, *speech or language, leads to the English* speak *and* speech. *The Greek* legein, *to speak, and* logos, *word, lead to* lexical, legible *and such relatives as* logic *and* intelligent. *Latin provides* language *and* linguistics, *through* lingua, *tongue, and* vocal *and* vocabulary, *through* vox, *voice.* Glossology *is an obsolete term for the study of language; like* glossary, *a collection of explanatory notes, it is derived from the Greek* glossa, *tongue or language.*

- Three-Term Contingencies in Verbal Behavior
- The Listener's or Reader's Behavior
  - *From Action to Acting to Literature: Creating Worlds*
- Interlocking Verbal Contingencies
  - *Verbal Governance*
  - *Replication*
  - *Attention to Verbal Stimuli*
  - *Verbal Shaping*
  - *The Coherence of Interlocking Verbal Classes*
- When Verbal Behavior Becomes a Closed System
- Addendum A: Verbal Behavior and Nonhuman Language

We've considered some of the properties of verbal behavior: instructional control, correspondences or equivalence classes, and discriminations of our own behavior in autoclitic processes, to mention just a few. In its full complexity, verbal behavior involves interactions among a variety of different processes. Just as a taxonomy of basic processes is required for the analysis of nonverbal behavior (Chapters 1 through 19), we need a taxonomy of verbal behavior. That taxonomy includes classes different from those in the everyday vocabulary. Textual behavior isn't equivalent to reading, although it may be its precursor. Transcription isn't equivalent to pictorial copying, but depends on the establishment of units of written verbal behavior. Tacting isn't equivalent to naming or referring, and yet as stimulus control of verbal behavior it is the point at which verbal behavior is anchored to the environment. The effectiveness of verbal behavior depends on the coordinations of these elementary components. By themselves they aren't even particularly verbal, but in fundamental ways our verbal behavior is built upon them.

KEY TERMS: Interlocking Verbal Contingencies; Behavior of the Speaker, of the Listener, of the Reader; Meanings and Equivalences.

THREE-TERM CONTINGENCIES IN VERBAL BEHAVIOR

Verbal contingencies can be organized in terms of antecedents, behavior and consequences. Table 1 classifies three-term contingencies by sorting them on the basis of the nonverbal or verbal status of antecedents and responses and the nonsocial or social status of consequences. The top half of the table (I) shows cases of contingency-governed behavior. When both antecedents and responses are nonverbal, such behavior can be maintained by either nonsocial consequences (A) or social consequences (B). When the behavior is verbal (C and D), it takes social contingencies to create classes of behavior occasioned by nonverbal antecedents (D), but once established, as higher-order classes, they may be maintained by nonsocial contingencies (C), as when some situation occasions a description of contingencies that alters someone's later behavior.

We speak of verbal governance whenever the antecedents are verbal (II), but the differential contingencies are most obvious when they maintain nonverbal behavior (E and F). Again, the higher-order classes are established by social contingencies, but once established such verbal governance can be maintained either by nonsocial consequences (E), as when someone makes a repair by following a service manual, or by social ones (F), as when someone complies with a request or follows an order. These relations have been distinguished by the terms tracking, for instruction-following based on correspondences between verbal behavior and environmental events, and pliance, for instruction-following based on social contingencies; a case can be made for also extending tracking to case G and pliance to case H.

Verbal governance allows a range of contingencies when verbal antecedents set the occasion for verbal behavior (G and H), probably because verbal units can serve as any term in a three-term contingency. Paralleling cases C and D, social consequences (H) create the higher-order verbal classes that may later enter into contingencies involving nonsocial consequences (G). Social consequences, often but not necessarily also verbal, are the glue that holds together everyday conversation. The contingencies that involve nonsocial consequences of verbal governance (G), which probably are prerequisites for cases that involve verbal behavior occasioned by nonverbal antecedents, are of special interest because they are essential features of science and technology.

Such contingencies have extended the ways in which we act upon our environments. Verbal behavior is typically shaped by social contingencies, but nonsocial contingencies may also be effective, as when an engineer's calculations lead to successive changes in the specifications of a project. Such instances clearly depend on an extensive social history of behaving verbally, but the interactions between writing and reading are functionally similar whether writer and reader are the same or different individuals.

Table 24-1 A Classification of Verbal and Nonverbal Contingencies

	Antecedents	*Behavior*	*Consequences*
I. Contingency-Governed Behavior			
A. Nonsocial Behavior	Nonverbal	Nonverbal	Nonsocial
B. Social Behavior	Nonverbal	Nonverbal	Social
C. Verbal Mediation, Self-Management	Nonverbal	Verbal	Nonsocial
D. Naming, Labelling, Description	Nonverbal	Verbal	Social
II. Verbally Governed Behavior			
E. Tracking	Verbal	Nonverbal	Nonsocial
F. Pliance	Verbal	Nonverbal	Social
G. Logic, Calculation, Invention	Verbal	Verbal	Nonsocial
H. Speaker-Listener Behavior	Verbal	Verbal	Social

THE LISTENER'S OR READER'S BEHAVIOR

Given that the speaker's verbal behavior provides discriminative stimuli for the listener, the listener's behavior is what is occasioned by these verbal stimuli. The listener's responses to verbal stimuli can be as varied as the responses to any other kinds of events. Many of the possible verbal responses have already been considered, in echoic, intraverbal and other classes. Some nonverbal responses occasioned by verbal stimuli are also obvious enough that they don't require special consideration. Whether the critical stimulus is a red light, a traffic officer's outstretched hand, the word *stop* or a tree fallen across the road, the driver's stepping on the brakes illustrates stimulus control. There is typically no need in such cases to distinguish among them on the grounds that some of the stimuli are products of someone's verbal behavior and others are not.

Listeners aren't passive, and they often behave verbally along with the speaker, saying things to themselves, planning replies and so on. Sometimes we act as our own listeners or readers, as when we carefully attend to what we're saying or think aloud or read something we've just written. We'll see in Chapter 29 that products of our own behavior can serve as discriminative stimuli and occasion our own future behavior (Skinner, 1989b).

As with nonverbal stimuli, not all responses to verbal stimuli are operant. For example, if a spoken word is paired with a stimulus that elicits autonomic responses (e.g., shock), the word may come itself to elicit these responses. This phenomenon, sometimes called *semantic conditioning* (e.g., Riess, 1946), is a verbal equivalent of the respondent conditioning of nonverbal responses. Responding generated by these procedures generalizes across semantic as well as phonological dimensions of verbal stimuli. For example, if electric shock is paired with a vehicle word, such as *truck*, the conditioned galvanic skin response is more likely to generalize to other vehicle words, such as *car* or *bus*, than to words that simply have some letters in common with the original word, such as *duck*. Semantic conditioning may contribute to instructional effects such as a child's adherence to warnings about a hot stove.

Yet if we say that a listener has understood something, it seems unlikely that we can provide an adequate account of the listener's response simply by appealing to relations among words (as in the giving of definitions) or of words with events (as in the teaching of tacts). The problem of meaning must reside at least in part in properties of the listener's responses to verbal stimuli. One critical property may be the correspondences between the responses occasioned by a word or utterance and the responses occasioned by the nonverbal events that the word or utterance ordinarily tacts. Many studies of verbal behavior are concerned mainly with how verbal responses occasioned by events vary together in the verbal behavior of a speaker or have common effects on a listener.

Whatever else is involved in the listener's behavior, the response to a tact must share some properties with the response to what is tacted. This is demonstrated when the properties of what is tacted interact with the properties of the relevant verbal behavior. For example, if words are printed in different colors, it is difficult to tact these colors rapidly if the words themselves are incompatible color-names, as when the word *red* is printed in green (Stroop, 1935)); we read words and don't ordinarily attend to physical properties such as the color in which they are printed. This relation between stimuli and the verbal responses they occasion clarifies some logical paradoxes of the language of reference. For example, consider the statement, *This statement is false.* If the statement is true, then it must be false; if it is false, then it must be true. Clearly it can't be true and false at the same time. This is a paradox of logic but not of verbal behavior. One verbal response can tact another verbal response, but it cannot tact itself. In other words, *This statement is false* is not a response to itself, but *It can't be true and false at the same time* might be a response to it.

In dealing with the formal relations, we argued that the symmetry of stimulus and response relations favors a vocabulary in terms of words rather than one in terms of specific vocal or written modalities. Similar correspondences exist in relations between tacts and environmental events, as was illustrated by Figure 22-2. These correspondences may be important when we speak of meaning, because the language of meaning is independent of whether words function as stimuli or as responses. This may be the most important way in which equivalence classes enter into verbal behavior.

When a listener repeats what a speaker has said and we say that the listener has understood the speaker, we aren't usually satisfied in calling the relation echoic. The several relations between *rain* as a verbal stimulus and *rain* as a verbal response in Figure 22-2 are among the criteria for talking about naming, but more is involved here than naming. Presumably we judge understanding or meaning not by any single relation between stimuli and responses, but rather by the integrity of the sorts of relations illustrated in Figure 22-2. We say that someone understands something that has been said when the individual repeats what has been said not because the other person said it but for the same reasons that the other person said it (cf. Skinner, 1968, p. 139). Such behavior implies consistent relations among verbal and nonverbal responses.

From Action to Acting to Literature: Creating Worlds

As we move from watching an actual incident to watching the incident acted in a play or a film and then to reading the script for the actual incident and then to reading a description of the incident in a story, the common feature that holds these cases together must lie in consistencies of stimulus control over verbal and nonverbal behavior. We can understand at least some of what happens when we watch a play or film or a television drama, because in those cases we are responding to stimuli with a reasonable resemblance to what we would see and hear if we were observing actual events. Part of the history of the movies is its progress toward ever more realistic presentations, in the transitions from silent films to sound, from black and white to color, from flat to 3D, and on ever larger screens. But those are produced for us. What happens when instead we must participate in creating a world by reading or listening to a novel?

One key is to recall that we are not passive in reading. Reading is much more than textual behavior (cf. Chapter 20). Reading is something we do, and part of it no doubt has much in common with the things we do that we call imagining and visualizing and thinking. In a work of fiction, or for that matter in most things we read, the author helps us by telling us what to do. Our reading is verbally governed, and the author's text provides the instructions.

The generalization from an actual event to a theatrical performance on stage to a movie or a television show seems obvious enough, but where do we go from there? Consider first that the stage or the screenplay did not occur spontaneously. A script came first, and that script provided instructions, probably varying in detail, to the cast and the set designers and stage hands and the lighting technicians and so on. Let's start with the assumption that what you do as a reader has something in common with what they do when they create the performance.

Some of what you read will be dialog, and when you get to what one character says in reply to the remark of another, your reading will depend not just on your textual behavior in reading that line but also on the effects of the preceding line that set it up. To the extent that you become involved in the story, you will be speaking along with the characters; skillful authors increase the likelihood that you will do so by including appropriate formal and thematic components along the way. "A novel achieves one of its main effects by preparing the reader to join in with the remarks of its characters" (Skinner, 1957, p. 275). Skillful writers will also keep your attention, but occasionally changing things on you in surprising ways; that will be most effective, however, if those changes are consistent with what you've read elsewhere in the story.

But it can't all be verbal. An instruction about the visual context demands seeing instead of saying, which we sometimes call visualizing or imaging (cf. Chapters 28 and 29). Some of us are better at this than others, but most of us can do well enough that long after our reading of a gripping book we can describe some of its scenes in some detail to others. In the course of reading the book we also have the opportunity to behave, though on a small scale, in ways that are otherwise unavailable to us; though it is not verbal, this behavior is verbally governed by what the author has written. With the author's help we can solve crimes or commit them, have love affairs or avenge betrayals, climb mountains or compete in races, travel through time or through space, fight wars or watch gardens grow, or just about anything else you can think of. What we read presumably also tells us something about what our reinforcers are. "One can read without guilt much which one could not say for other reasons. The book itself and the act of reading constitute a tolerant situation in which verbal behavior is freely emitted" (Skinner, 1957, p. 398). Clearly the behavior of reading a novel and the behavior of reading a textbook differ in some important ways, but they do have some features in common. I'll leave it to you to work some of them out.

INTERLOCKING VERBAL CONTINGENCIES

Verbal governance, replication, differential attention, and verbal shaping share common stimulus and response terms and so produce interlocking contingencies in which extensive classes of behavior come to be dominated by verbal antecedents. When words enter into such interlocking contingencies, they may become increasingly removed from nonverbal contingencies and their relation to those contingencies may become distorted by social contingencies (Catania, 2003a, 2004). Now that our view of our species as occupying a central place in the universe has been pummeled by Galileo and Hubble and Darwin and Freud, verbal behavior may be the last bastion of our anthropocentrism. Maybe we should not assume that it gives us bragging rights.

We'll focus here on four topics: (1) *verbal governance*, in the contingencies, mainly social, that lead to instruction following and to correspondences between what we do and what we say about what we do; (2) *replication*, in the echoic and other processes that are prerequisites for initiating and maintaining verbal behavior in individuals and spreading verbal behavior throughout social communities; (3) *attention to verbal stimuli*, in which the reinforcing and aversive properties of these stimuli affect whether they are sought out or avoided and whether they become incorporated into one's own verbal behavior; and (4) *verbal shaping*, in the natural and artificial contingencies that arrange consequences for verbal behavior and thereby raise or lower the probabilities of different verbal classes.

These four categories could easily have been treated in a different order. As we'll see, they so interact that it is difficult to assign priorities. Verbal governance does seem fundamental, as a primary function of verbal behavior is to change someone's behavior, but that behavior will survive only if there is replication, it will work only if there is attention to verbal stimuli, and it can evolve only if it can be differentially reinforced.

These phenomena are not theoretical. They occur in everyday instances of human verbal behavior. We observe verbal governance and replication and differential attention and verbal shaping whenever we see individuals responding to simple requests or repeating something said or attending to one speaker rather than another or changing a topic of conversation upon changes in audience reaction. If the account is theoretical in any sense, it is in treating how these phenomena can come together to produce complex outcomes.

Verbal Governance

Verbal discriminative stimuli set the occasion for verbal and nonverbal responses; these responses may have consequences. Verbally governed behavior is behavior, either verbal or nonverbal, determined by verbal antecedents and maintained mostly by social contingencies. These contingencies can operate globally on higher-order classes, or locally on specific subclasses, or both. Verbally governed behavior is often determined more strongly by higher-order social contingencies than by more local (often nonsocial) contingencies. We do not ordinarily tell people to do what they would do even without being told.

Verbal governance, like any other operant class, may be conditional on other events. We learn to follow instructions from some individuals but not others and we often talk to ourselves, usually silently but occasionally, as when following complex instructions, out loud. Thus, verbal antecedents may lead to other verbal behavior, as when we derive implications and courses of action from something said. And because we often distinguish between what we have been told and what we have arrived at without being told, the most effective verbal antecedents may be those we generate ourselves. In such cases, we may not recognize the remote origins of what we generated, in the verbal behavior of others that initiated our own self-talk.

Replication

Verbal discriminative stimuli set the occasion for verbal responses; these responses can function in turn as verbal discriminative stimuli. We tend to repeat what we and others say. But replication is not mere reproduction. Based on a long and complex history in which we learned the relations among spoken and written stimuli and responses, we say that words are the same whether spoken or written. We repeat what others say or copy what they have written, but we also speak what has been written or copy what has been said. These are instances of replication even when they carry across different modalities.

The relevance of replication so viewed is that the effects of verbal stimuli become extended over time and space as verbal behavior produced by some is passed on to others. Such replication remains important throughout life, as when we repeat something we have heard or jot down a note about it.

Once some individuals repeat what they or others say, verbal behavior may be maintained by cultural contingencies and therefore survive across generations. Some effects of replication are fairly straightforward. For example, if a single utterance has no effect repetitions may produce one. If the replicated verbal response also participates in verbal governance, the speaker's influence is extended by the listener's replication, as when instructions are shared by group members. Furthermore, repetitions allow instructions to be followed in the original speaker's absence, later and elsewhere, thus transferring governance from the speaker's verbal behavior to the listener's replication, as when we repeat to ourselves the details of a task someone has asked us to complete. In other words, repetition extends the effects of verbal behavior over time and space. Replication must come first, but once in place powerful contingencies can maintain it.

Attention to Verbal Stimuli

Verbal or nonverbal responses produce verbal discriminative stimuli; these in turn produce other verbal or nonverbal responses. Some discriminative stimuli are correlated with reinforcers and others with extinction. Only the former are likely to acquire reinforcing functions of their own, when they are called conditional reinforcers. The effectiveness of discriminative stimuli depends on whether the organism attends to them, which in turn depends on their status as conditional reinforcers . Looking at or attending to a stimulus is reinforced when the stimulus is correlated with reinforcers but not otherwise. An informative stimulus is also discriminative, but its informativeness alone is not sufficient to maintain attention to it. A discriminative stimulus correlated solely with extinction or with aversive events may not maintain attention, and discriminative stimuli correlated with aversive events about which something can be done may be more likely to maintain attention than those correlated with uncontrollable aversive events.

Experiments on attention typically make attending explicit by introducing an observing response, a response that produces discriminative stimuli that are otherwise unavailable. Changes in the rate of this response may then be studied as a function of the correlations of discriminative stimuli with various reinforcing or aversive contingencies .

Given that organisms attend to stimuli mainly when the stimuli are conditional reinforcers, the effectiveness of messages depends more on whether their content is reinforcing or aversive than on whether they are correct or complete. In fact, what needs explanation is that humans attend to bad news or sometimes reach conclusions when the answer is not what they wanted to hear. The key probably lies with whether the bad news allows effective action, and even with uncontrollable aversive events many stimuli are sufficiently correlated with reinforcers that they can maintain attention despite their correlation with occasional aversive events.

Attention to a verbal stimulus is a prerequisite for its replication or for governance by that stimulus. We repeat to ourselves what we have heard or read. Once we have done so our own verbal behavior is likely to summate with other verbal antecedents that participate in verbal governance. If these verbal stimuli were reinforcing in the first place, then behavior producing contact with more or similar instances will be strengthened, in turn providing even more initiating instances.

Verbal Shaping

Verbal responses have consequences; the new verbal responses shaped by these consequences may then function as discriminative stimuli. In shaping, reinforcement produces a spectrum of responses that differ from the reinforced response along dimensions such as topography and magnitude. Verbal shaping is of special interest because it also works along semantic and syntactic dimensions. Audiences provide discriminative stimuli that set the occasions for verbal behavior and provide reinforcers that shape it. Anyone who has observed drifts in the content of conversation as attention to particular topics or speakers picks up or flags has seen shaping in natural settings.

The shaping of verbal behavior is a potent technique for changing human behavior. A practical implication is that shaping what people say about their own behavior may be more effective in changing it than either shaping it directly or telling them what to do, perhaps because verbal communities arrange contingencies for correspondences between saying and doing, as in those for lying versus telling the truth or for keeping versus breaking promises.

The fact of verbal shaping implies that verbal behavior itself is sensitive to local contingencies even when governing other behavior. The greater effectiveness of verbal shaping than instructions may depend in part on the speaker's failure to discriminate the sources of the verbal behavior. Even speakers who accurately discriminate among various sources of verbal behavior when following instructions usually call their own shaped verbal behavior self-generated. It matters whether we have been told what to say or have come to say it in other ways .

The Coherence of Interlocking Verbal Classes

Some of the most interesting reinforcers of verbal behavior are themselves verbal: an answer to a question, an acknowledgment, and so on. If some verbal consequences are more effective reinforcers than others, some will command more attention than others. Verbal stimuli are discriminative stimuli, and we attend to them, as we attend to nonverbal ones, not based on the information they carry but rather based on their correlation with reinforcers and aversive events. Having attended to them we may replicate them, but over successive replications they may be subject to further shaping.

And the newly shaped verbal behavior may also begin to participate in verbal governance. In this brief compass we have already touched on each of the four verbal properties considered here. These paradigms summarize them (their nonverbal components have been omitted):

Verbal Governance
VERBAL S^{D}'s → VERBAL R's
VERBAL R's → CONSEQUENCES

Replication of Verbal Behavior
VERBAL S^{D}'s → VERBAL R's
VERBAL R's → VERBAL S^{D}'s

Differential Attention to Verbal Stimuli
VERBAL R's → VERBAL S^{D}'s
VERBAL S^{D}'s → VERBAL R's

Verbal Shaping
VERBAL R's → CONSEQUENCES
NEW VERBAL R'S → VERBAL S^{D}'s

The same sorts of terms enter into each category. A verbal discriminative stimulus participating in verbal governance can be replicated and can command attention, and its replications can be shaped by their consequences. Verbal behavior that has been replicated can participate in verbal governance, can command attention, and can be shaped by its consequences. A verbal stimulus commanding attention can be replicated and can participate in verbal governance, and its replications can be shaped by their consequences. A shaped verbal response is replicated in the course of its shaping and can function as a verbal stimulus commanding differential attention and participating in verbal governance. But if the consequences maintaining these aspects of verbal behavior are primarily social, then verbal communities can create and maintain verbal classes that have ever diminishing contact with nonverbal contingencies. In other words, they can create verbal worlds that become increasingly autonomous.

Differential attention to verbal stimuli may lead to self-generated verbal behavior that is independent of current environments, thereby pervading someone's behavior across a range of situations. As the circle of speakers and listeners becomes more limited and contact with current antecedents and contingencies becomes more tenuous (it is difficult, after all, to disprove claims of heaven and hell, parallel universes and alien abduction), such verbal contingencies may generate idiosyncratic verbal interactions maintained in isolated individuals or small groups, as the social consequences provided by the broader verbal community become less relevant. When such concentrated effects show up in individuals, we sometimes call them interests or obsessions; when they extend across groups, we sometimes call them fads or cults.

Each of the four processes interacts with every other one as well as with itself, and such verbal functions can affect not only the verbal behavior passed on among members of a verbal community but also the verbal behavior of individual speakers or listeners. These processes do not merely fit well together; they are components of cycles within which they strengthen and build upon each other.

These contingencies are illustrated in Figure 24-1. Omit any one and the cyclicity is significantly attenuated. The social reinforcers participating in these functions are ordinarily small, but they operate on behavior day in and day out over weeks and months and years. Should we be surprised at the effects of interlocking verbal contingencies operating over human lifetimes? Verbal governance and verbal replication and differential attention to verbal stimuli and verbal shaping make their own separate contributions, but they produce complex behavior that is much more than just the sum of its parts.

With regard to nonverbal contingencies, verbal behavior need not always get in the way. Social contingencies create higher-order classes of verbal behavior occasioned by nonverbal antecedents, but such classes, once established, may be maintained by nonsocial contingencies, as when a situation occasions a description of contingencies that alters someone's later nonverbal behavior. Similarly, governance by verbal antecedents is established by social contingencies, but once in place it can be maintained either by nonsocial consequences, as when someone makes a repair by following a service manual, or by social ones, as when someone complies with a request.

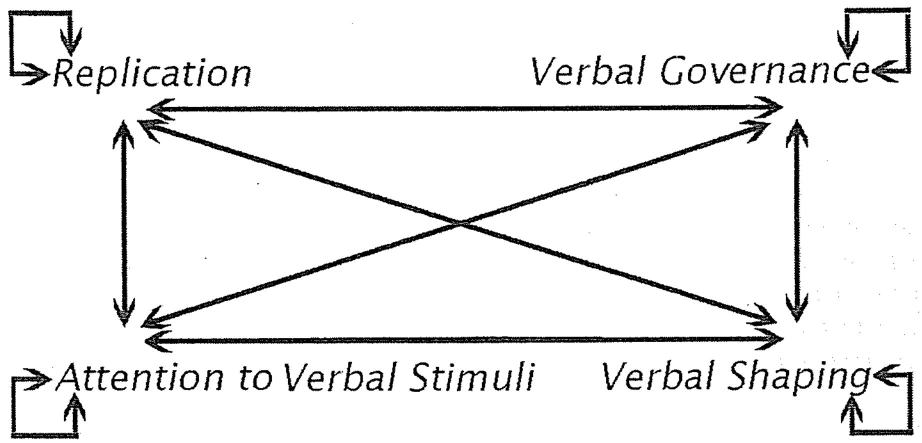

Figure 24-1 Interactions among replication, verbal governance, verbal shaping, and differential attention.

WHEN VERBAL BEHAVIOR BECOMES A CLOSED SYSTEM

In Chapter 23 we discussed the relevance of verbal governance and verbal shaping to cultural practices and how such practices can be initiated and maintained within verbal communities. Consider Holocaust denial as an example. An uncle of mine, while in his eighties, had been at a social event. At one point he suddenly lunged at a younger man he'd been talking to and had to be pulled off him. It later became clear that my uncle had done so because that person was a Holocaust denier sounding off.

My uncle had served in an engineering unit during World War II but until recently had hardly talked to his family about his service. His war experience included the combat that took the bridge at Remagen and led to the crossing of the Rhine and, later and more important, being among the early units that liberated some of the Nazi concentration camps in Germany. Like many others who had seen the corpses and the emaciated survivors (Hirsh, 2010), he didn't talk about it when he returned from the war. We can understand why. The war was over and there had already been plenty of bad news, and the family settings to which those servicemen returned did not provide the sorts of audiences that had been available when they were in combat together. It was hard enough for them to see how terribly some people could treat others; what would it take to persuade people that things so hard to believe had actually happened?

I also had an older cousin who had jumped into France with an airborne unit on D-Day and continued in combat through Hurtgen Forest and the Battle of the Bulge up to the German surrender. Though the families knew each other, my uncle and my cousin only learned about each other's European service just a few years ago. We must stand in awe of what they did, but that's how little they had each talked about it.

We can understand why my uncle tackled that Holocaust denier, and what we know about audiences helps us to understand why he and so many others were silent for so long. But it is perhaps even more important to try to understand how Holocaust denial arose, and the origin of antisemitism in prejudice and discrimination is probably only part of the story. Sadly the silences of those eyewitnesses made it easier for the Holocaust deniers to construct their own stories and to find audiences where, uncontradicted, telling those stories was reinforced instead of being punished. As we have seen, bad news is typically not well received. Inquisitions are based in part on suppressing things that some do not want to hear. Perhaps the denial of the Holocaust has something in common with the denial of a heliocentric solar system or the denial of evolution or the denial of climate change caused by humans (e.g., Garwood, 2007).

Verbal contingencies can create restricted communities in which verbal shaping can overwhelm social systems. What begin as subtle verbal distinctions can sometimes be so amplified and widely disseminated within a culture that stereotyping and persecution and even genocide can become pervasive and systematic (e.g., Evans, 2004, 2005, 2009). Pernicious verbal practices can evolve within churches that shun members who raise questions or fail to adhere to its teachings, or social groups that differentially reinforce prejudicial practices with regard to race or ethnicity or religion or political affiliation, or political groups that spin out conspiracy theories and shout down those who disagree. Many leaders have led their countries to war and destruction and some of what they did depended on their control over instruments of coercion; nevertheless, they could not have succeeded without words (Bix, 2000; Chang & Halliday, 2005; Gellately, 2007; Kershaw, 1998, 2000). But if we cite Mein Kampf or the Communist Manifesto or the Code of Bushido or the Little Red Book of Chairman Mao, what then about the Magna Carta and the Declaration of Independence? Let's at least question the wisdom of the old saying that sticks and stones can break our bones but words can never harm us. We must be wary of the power of words, especially now that they can spread so much more widely and quickly than at any other time in human history.

Social consequences hold everyday conversation together, but contingencies that involve nonsocial consequences of verbal governance are essential features of science and technology. Its anchoring to nonverbal environments through data is the special advantage of scientific verbal behavior. Scientists attend to and replicate what others have found and those replications may govern subsequent research, which in turn may be shaped by interactions in the laboratory and across the scientific community. Very different consequences follow from verbal behavior depending on whether it is anchored to environmental events, as in scientific verbal practices, or becomes independent of it, as in poetry and fiction and, too often, politics.

Those who become engrossed in artificial verbal worlds such as those of Don Quixote or Frodo Baggins or Harry Potter may get themselves into trouble by neglecting other contingencies but might not otherwise make much trouble for others. That is not always so, however. In fundamentalism, the word becomes all powerful. Fundamentalism is defined by and cannot exist without verbal control. It is maintained through the shaping of verbal behavior consistent with the replication of texts that command differential attention and that participate in verbal governance. Verbal behavior can be tightly determined by nonverbal environmental contingencies, as in scientific practices, or loosely determined, as in social practices such as literature and religion.

We considered scenarios for the evolution of human verbal behavior in Chapter 19. Once speakers could instruct the verbal behavior of listeners who could in turn instruct nonverbal behavior, the prerequisites for human political and religious institutions were firmly in place. Human behavior throughout the world has been and still is heavily influenced by records of long-past verbal behavior. Like government, religion has its roots in social control. Religious behavior provides compelling examples of the phenomena reviewed here, in verbal governance as demonstrated by the following of religious precepts, in the replication of verbal behavior through recitations of sacred formulas, in differential attention to prescribed and proscribed texts, and in verbal shaping in the course of religious rites such as devotions and meditations and pilgrimages and confessions and inquisitions. Heaven and hell, like angels and devils, are human verbal creations.

In his concept of superstition Skinner appealed to accidental contiguities of events, but the effects of accidental contiguities on behavior are mostly quite temporary (cf. Chapter 5). Verbal governance has far more powerful effects, and much human behavior that we call superstitious involves verbal practices rather than the accidental strengthening of nonverbal behavior. In their reliance on verbal governance, governmental and religious institutions typically have much in common (cf. Skinner, 1953). Exhortations to behavior in the name of patriotism are not so very different from exhortations in the name of a deity.

Verbal competence is said to separate man from beast. It is said to make human rationality possible. It is therefore a curious cultural practice that allows more control to be exerted on behalf of religious belief than on behalf of scientific findings, because the criterion for religious belief is precisely that it is unverifiable. Freedom of religion implies that you can act on almost any belief whatsoever, unless what you believe happens to be supported by evidence. Science, and especially behavioral science, is at risk if contact solely with words is regarded as more fundamental than other kinds of environmental contact.

We've seen that verbal behavior includes many components: higher-order classes in naming; correspondences in formal classes; stimulus control in intraverbals and tacts; contingencies in instructional control; symbolic behavior in equivalence classes; discriminations of our own behavior in autoclitic processes. The list is incomplete. To the extent that these processes are related to those of nonverbal behavior, they hint at the foundations of human language. The assumption that its primary function was to direct the behavior of others through instructional control suggests how such control might have emerged and might have been sharpened by the contingencies that operated within human social groups (Catania, 2001b; Jaynes, 1976; Skinner, 1986).

We'd do well to recall the constraints on selection. For example, organized complexity can evolve only if it remains advantageous at every stage of selection (cf. Chapter 3). We can then wonder about the selective advantages of infant babbling, or the consistent features of those verbalizations of a mother toward an infant that we call *motherese*, or the changes that make it more difficult for adults than children to learn the different phonetic structure of a new language, and so on. Once the social selection of behavior began to operate on verbal behavior, the way was clear for the development of other functions of verbal behavior as derivatives of its primary function. That is where we should seek the foundations of human concepts such as narrative, communication, meaning and truth.

ADDENDUM 24A: VERBAL BEHAVIOR AND NONHUMAN LANGUAGE

It is probably because verbal behavior involves so many different processes acting together that the issue of nonhuman languages has typically engendered controversy. There certainly exist many cases in which the sounds or gestures of one nonhuman organism influence the behavior of other members of that species; some were considered in Chapter 13. Birdsong, for example, has important effects in mating and the establishment of territories. The songs of birds often depend on the social context in which they sing (West & King, 1980); those of some birds have dialects, and birds who don't hear them at an early age don't sing them as adults (Marler & Peters, 1982). The dependence of birdsong on both genetic and environmental history has some parallels in the development of human vocal behavior (as in the shaping of echoic behavior), but the functions and structures of the two kinds of behavior differ in many important ways. For example, absolute frequency is more important in discriminating among melodies for birds than for humans (Hulse, Cynx, & Humpal, 1984).

Human language involves much more than vocal releasers or stimulus control based upon vocal stimuli. For that reason, discriminations based even on very subtle events, as by the horse Clever Hans, have never counted as instances of verbal behavior (cf. Pfungst, 1911; Sebeok & Rosenthal, 1981).

The domain of animal language has often been taken to broadly encompass any case in which the behavior of one organism serves as either an eliciting or a discriminative stimulus for the behavior of another (e.g., Bright, 1985). Accounts have dealt with the full range of nonhuman communication, ranging from the flashes of fireflies and the chirps of crickets to the rivalry calls of deer and the coordinated songs of whales. Many of these cases involve stimuli or responses of special interest, but their relevance to human verbal behavior is usually limited.

As verbal classes, tacting involves stimulus control as it occurs in verbal behavior and manding involves consequences as they act in verbal behavior. Those and other processes are important and have been studied with several species. For example, some studies have examined whether bird vocalizations can function as tacts or even as members of equivalence classes (cf. Manabe, Kawashima, & Staddon, 1995, on budgerigars; Pepperberg, 1988, on the parrot).

Other research has concentrated on structural rather than functional aspects of human languages as they might enter into nonhuman behavior, such as discriminations of human phonetic categories by Japanese quail (Kluender, Diehl, & Killeen, 1987). Attention, however, has more often focused on aspects of human verbal behavior that aren't obvious components of nonhuman behavior. For example, studies of the behavior of sea mammals such as dolphins and sea lions have shown them capable of sophisticated relational discriminations in both auditory and visual modes, but the sides taken on their verbal competence have depended on judgments about whether it is appropriate to regard the complex properties of their behavior as nonhuman examples of grammatical structure or of equivalence relations (cf. Herman & Forestell, 1985; Kastak & Schusterman, 2002; Kastak, Schusterman, & Kastak, 2001; Schusterman & Kastak, 1993).

In the search for nonhuman language, investigators have most often turned to chimpanzees and other primates (cf. Savage-Rumbaugh, 1986). Some of our information comes from behavior observed in natural habitats (cf. Gouzoules, Gouzoules, & Marler, 1984; Seyfarth, Cheney, & Marler, 1980b). For example, the sound patternings of vocalizations may determine how well one individual can recognize kin or other individuals (cf. Rendall, Rodman, & Emond, 1996).

Early attempts to demonstrate language in chimpanzees were unsuccessful because they concentrated on language in the speech mode (Hayes & Hayes, 1951). But the chimpanzee's vocal apparatus restricts its capacity to produce differentiated vocal behavior, and the chimpanzee Viki learned to imitate only a few human utterances: *mama*, *papa*, *cup*, *up*.

Researchers then switched to languages based on other modalities. The chimpanzee Washoe learned the gestures of American Sign Language (Gardner & Gardner, 1969); the chimpanzee Lana learned a language based on visual displays and key presses at a computer console (Rumbaugh & Gill, 1976); and the chimpanzee Sarah learned a language based on arrangements of plastic chips of various shapes and colors on a magnetic board (Premack, 1970). In these projects, chimpanzees acquired large vocabularies and began to produce word combinations, but as each new aspect of the chimpanzees' capacities was demonstrated, questions were raised about others. For example, after the chimpanzee Nim Chimpsky was taught some of the gestural vocabulary of American Sign Language, the structure of its word combinations was compared with that of the language of a human child (Terrace *et al.*, 1979; see also Thompson & Church, 1980). The sequential structure of Nim's word combinations was less orderly than a child's, and it was concluded that Nim's behavior couldn't be called language because it lacked adequate structure or syntax. Yet there were inconsistencies in the signing across Nim's teachers and structure was deliberately avoided so as not to impose it on Nim's signing, whereas a child's early verbal environment includes the structured speech of adults.

Other studies were concerned with the verbal implications of the chimpanzee's capacity to discriminate complex relational properties of the environment (e.g., Gillan, 1981; Savage-Rumbaugh *et al.*, 1980) and with the emergence of instructional control from the language of tacting, especially in the interactions of the chimpanzees Austin and Sherman (Savage-Rumbaugh, Rumbaugh, & Boysen, 1978). The syntheses of complex interactions in such studies have special value because they force all of the assumptions about what counts as verbal to be made explicit (Epstein, Lanza, & Skinner, 1980). You can't tell a pigeon or a chimpanzee what to do in an experiment; instead, you have to shape every component that will be integrated into the final performance.

The performances of nonhuman organisms will become ever more sophisticated as human investigators provide them with ever more sophisticated environments. Differences will remain, some more obvious than others (e.g., operant control of the vocal apparatus: cf. Chapter 14). For example, Kanzi, an 8-year-old bonobo monkey (*Pan paniscus*), appears capable of sentence comprehension that has so far been beyond the reach of its relative, the chimpanzee (*Pan troglodytes*) and that may be comparable to that of a 2-year-old human child (Savage-Rumbaugh *et al.*, 1993). We've already seen how crucial verbal communities are in the shaping and maintenance of human verbal behavior, so we shouldn't be surprised that the details of verbal contingencies matter. For example, judgments of the nature of the discrimination between requests such as *Get juice from kitchen* or *Take juice to kitchen*, or their ape language equivalents, depend on contextual features of the requests such as whether juice is always present when the requests are made; if it is, then the discrimination is probably based on the difference between *get from* and *take to*, but if juice is present only given the request to take it to the kitchen, then the discrimination may depend only on whether juice is present when a request has *juice* in it.

Once some features of human language had been demonstrated in the behavior of a chimpanzee or a pigeon or any other nonhuman organism, that feature could no longer be regarded as uniquely human; attention then turned to the definition of language rather than to the experimental analysis of its properties. The issue of grammatical structure, which we'll consider in Chapter 25, was the focus of much controversy. Given these debates, we can't say whether chimpanzees are capable of language; the answer depends too much on how we define language. We certainly can say, however, that their behavior includes some critical components of language.

Chapter 25
Language Structure

Three primary terms in the analysis of language have been syntax, *the study of grammatical structure;* semantics, *the study of meaning; and* pragmatics, *the study of the functions of language.* Syntax *can be traced to the Greek* taxis, *arrangement;* semantics *to the Greek* sema, *a sign or thing seen; and* pragmatics *to the Greek* prassein, *to make happen or do.* Pragmatics *is a relative of* practice. Grammar, *through the Greek* graphein, *to scratch or write, and* gramma, *a picture or a writing, is closely related to* graph, program *and* topography.

Syntax: The Grammatical Structure of Language
  - *Constituents and Phrase Structure*
  - *Transformations*
  - *Relational Autoclitics: The Conjunction of Verbal Units*

Semantics: The Meaning of Verbal Units
Language Development
Addendum A: Some Properties of Language
Addendum B: Deixis

This chapter turns from the functions of verbal behavior to its structure. We can arrange words in sentences and note how different words are similar or dissimilar in meaning. These are the topics of *syntax* and *semantics*. *Syntax* deals with how we organize words in sentences; its concern is grammatical structure. In treating syntax, we'll approach language from the point of view of psycholinguistics, noting how its vocabulary is related to the functional account of Chapters 20 to 24 when appropriate.

Semantics deals with the problem of meaning. Our treatments of verbal behavior have questioned traditional concepts of meaning and reference. In this chapter we'll see what can be said about them. In effect, we'll ask what properties of verbal behavior lead us to say that particular words are related in meaning. Psycholinguistics deals with a speaker's vocabulary in terms of the speaker's *lexicon*, the dictionary of words available in the speaker's verbal behavior. Studies of meaning are concerned with the structure of the lexicon. *Phonology*, concerned with the structure of speech, and *pragmatics*, concerned with the uses of language, are two other significant areas of linguistics; some material in Chapter 20 is relevant to the former and the latter is most closely related to the functions of verbal behavior we considered in Chapters 22 to 24.

A critical precursor of contemporary psycholinguistics was the effort to write programs for computer translation from one language to another. That history included the evolution of computers in Allied efforts to break military codes during World War II (Hodges, 1985) and later applications of mathematics to problems of language structure during post-Sputnik efforts to translate Russian technical materials into English. Backed by the prestige of mathematics and the growth of computer technology, linguists made the case that their field was the one that could produce effective computer translations along with other achievements. Presumably that played a role in one linguistic response to Skinner's *Verbal Behavior* (1957).

KEY TERMS: Syntax, Semantics, Pragmatics; Phrase Structure or Constituent Grammar, Transformational Grammar; Surface Structure, Deep Structure; Psychological Reality, Language Relativity.

I was in the audience as a graduate student in the Harvard Department of Psychology in the Fall of 1959 when an MIT linguist named Noam Chomsky gave a colloquium on his new transformational grammar. Skinner was in attendance, but Chomsky didn't mention his forthcoming review of *Verbal Behavior*, which must already have been written because it appeared soon after. It overshadowed many favorable reviews of Skinner's book (Knapp, 1992). Chomsky was mainly concerned about language structure, and for many years it was assumed that Chomsky's critique had destroyed Skinner's position. One reason was that years passed before responses to Chomsky's arguments appeared, so some readers concluded Skinner could not produce an effective reply. But most behavior analysts who read Chomsky's review, including Skinner, had concluded that Chomsky had missed the point of the book, so no reply seemed necessary (cf. Harris, 1993).

One reply was rejected by the journal *Language*, the journal that had published Chomsky's review; it appeared elsewhere some years later (MacCorquodale, 1970). Among its several arguments was that Chomsky's criticisms were relevant to Clark L. Hull's learning theories rather than to Skinner's work. A comparison of Skinner's approach to Wittgenstein's treatment of language games (Day, 1969), some passing comments on Chomsky by Skinner (1972), and occasional reassessments of Skinner's book and Chomsky's review (e.g., Andresen, 1990; Catania, 1997a; Czubaroff, 1988) added to the case. Linguists often still fail to appreciate the functional content of Skinner's account and typically ignore the research that has expanded Skinner's early taxonomy and extended it to the origins of novelty in verbal behavior. They concentrate more on whether what is said is grammatical than on how it functions (e.g., Catania, 2003b; Jackendoff, 2002; Pinker, 1994).

The goal of language translation provided one focus for linguistic work on syntactic structure as the United States and the Soviet Union faced off against each other during the Cold War. Early attempts simply to substitute words in one language for their equivalents in another were unsuccessful for a variety of reasons. For example, many words have multiple equivalents (should *bar* be translated as a noun, either a lever or a place to drink, or as a verb, to stop?), and grammatical distinctions in one language may be absent in another (Russian doesn't use articles, so how does one decide whether an English translation of a Russian noun should be preceded by *a* or *the* or no article at all?).

One test for language translation programs is to translate a text from one language to another and then to translate the new text back to the original language: If the translation program works, you should get back the same sentence as the one you put in. One illustration of the translation problem, probably apocryphal but cited often, with variations, is the translation to Russian and then back to English of *The spirit is willing but the flesh is weak*; it comes back as *Strong vodka, rotten meat.* Another is *Out of sight, out of mind*; it comes back as *Blind maniacs.* And if the program generates *The lions leave by the end of summer*, its source sentence, though fitting, probably isn't obvious: *Pride goeth before a fall.*

The treatment that follows is necessarily of very limited scope. It examines some parallels between evolving linguistic concepts and some aspects of the evolution of computer programs in language translation and concentrates especially on features that most readily make contact with behavior analytic concerns. Those early programs had to define explicit procedures for translation, including rules for substitution and transformation and ways to interpret ambiguous terms on the basis of context provided by earlier text. Although they revealed serious limitations on computer translation (e.g., Dreyfus, 1992; Winograd, 1980), they also led to more thorough descriptions of the complexities of syntax and semantics. When effects of various features of syntax and semantics were demonstrated in verbal behavior, those features were said to have *psychological reality*. When we can show that some structural properties of language make a difference in the behavior of speakers or listeners, it is appropriate to speak of their psychological reality.

SYNTAX: THE GRAMMATICAL STRUCTURE OF LANGUAGE

How are words organized in sentences? We can discriminate between grammatical and ungrammatical sentences. We can classify grammatical sentences into categories such as active voice, passive voice, past tense and so on. But how do we define the dimensions along which we make these distinctions? Any listing of grammatical sentences would be infinite if we didn't restrict sentence length and indefinitely large even if we did. And if we could produce such a list, we still wouldn't know what made a sentence grammatical. Instead of lists, we need an exhaustive description of kinds of sentences. In the language of psycholinguistics, we'd speak of writing a grammar with a finite number of rules.

We'll review two accounts of the structural regularities of grammatical sentences (Catania, 1972; Chomsky & Miller, 1963). The first describes sentence structures in terms of their *constituents* or components. The second examines the *transformations* that show how one sentence structure is related to others. For example, consider the saying *He who hesitates is lost.* When we study how parts of the saying come to be named subjects and predicates or pronouns and verbs, we're dealing with constituents. When we describe its relation to such paraphrases as *He is lost who hesitates*, we're dealing with grammatical transformations.

We might have tried to handle grammatical structure in terms of word sequences, as in the analysis of intraverbal behavior. The trouble is that such analyses cannot handle relations between words separated by varying numbers of other words. For example, consider *The chimpanzee used sign language* and *The chimpanzee taught by the psychologist used sign language.* The words *chimpanzee* and *used* are adjacent in the first sentence but separated by four other words in the second, and yet they are grammatically related in the same way in both sentences.

Another problem with sequences is that common word sequences can occur together in ungrammatical ways (*Haste makes waste not want not*) whereas rare sequences may be grammatical (*Sleeping green ideas dream furiously*). The probabilities with which different words follow each other don't tell us anything about the grammaticality of sentences. Even if we resolved that problem, we'd still have the problem of ambiguous sentences. *Running experiments should be encouraged* might be read as recommending more support for research (*Encourage the running of experiments*) or more research on exercise (*Encourage experiments on running*). We can't distinguish between the two interpretations based on word sequences, because the same words appear in the same order in both readings. It would be more helpful to know whether the sentence was uttered in a laboratory or at a track meet; but that's a matter of function, not structure.

Constituents and Phrase Structure

We determine the constituents of sentences by noting how their parts are related. Consider *A word to the wise is sufficient.* We can name its constituents: *Word* is a noun, *is* a verb, *sufficient* an adjective, and so on. The relations aren't defined by how close the words are to each other. For example, *is* is more closely tied to *word* than to *wise*, even though *is* is positioned more closely to the latter. The relations among words in sentences have been represented in various ways (e.g., Wundt, 1900). Three representations, sometimes called phrase-structure diagrams, are illustrated in Figure 25-1. The examples in the figure provide names for the constituents. But how do we decide whether a word should be called one type of constituent or another? We can't go by the words alone. We have to look at their relations to other words in the sentence. In other words, we can't classify constituents without also identifying the structure of the sentence.

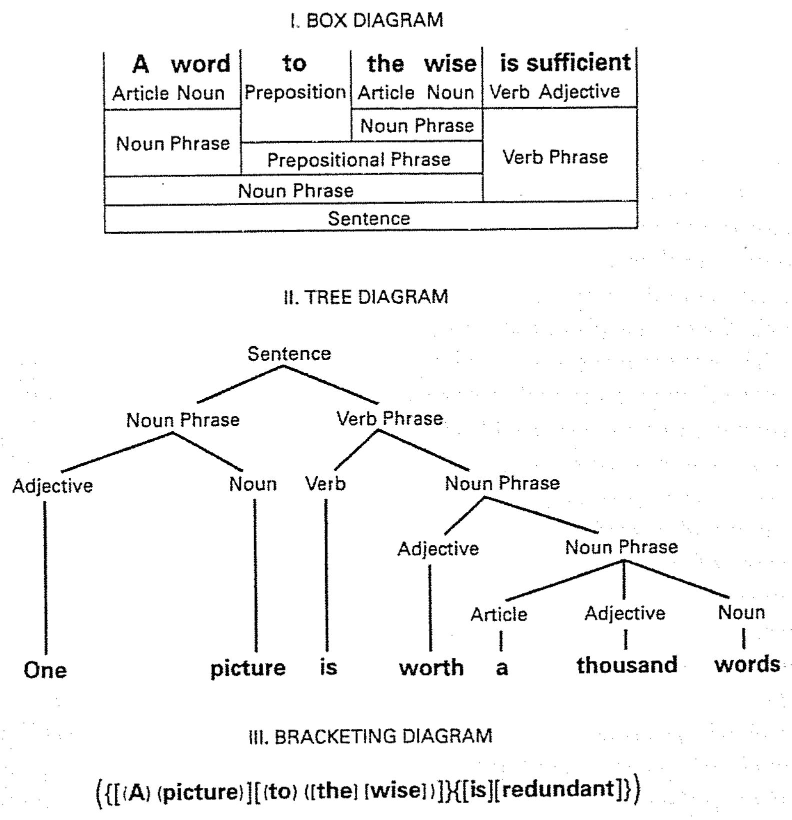

Figure 25-1 Three methods for representing constituent structures. Each shows how a sentence can be analyzed into structural units ranging from individual words to phrases. The sentences in the box diagram (I) and the bracketing diagram (III) have equivalent structures. In psycholinguistics, a tree diagram (II) has been the most common representation. Branches come together at places called nodes; e.g., *verb phrase* in the example is a node for *verb* and *noun phrase*.

Consider *He whose laughs last laughs last. Laughs* appears twice, first as a plural noun and then as a verb; *last* also appears twice, first as a verb and then as an adverb. So what we call *laughs* or *last* depends on the relation of each to other words in the sentence. This poses a problem. We saw earlier that we can't identify a sentence structure without classifying its constituents, but this example shows that we can't classify its constituents without identifying its structure. How then can we ever describe sentence structure? The answer is that the consistencies of sentence structure aren't in particular words or word sequences; instead, they're in various kinds of coordinations among words. We can classify words of a sentence as particular kinds of constituents because we've already learned typical structures (e.g., patterns of agreement between singular and plural nouns and verbs).

We might be tempted to look to the environment as a basis for deciding on sentence structure. But that won't work, because we can name the constituents even of some sentences partly made up of nonsense words. For example, compare *He who gulks merts* and *She merts his gulks.* Even though *gulks* and *merts* aren't standard words, we'd call both verbs in the first sentence but we'd call *merts* a verb and *gulks* a plural noun in the second.

Grammatical classifications of words don't depend on the environment events of which we speak. They depend instead on the sentence structures within which words appear. Verbs, for example, aren't defined as the class of words occasioned by activities; they are defined in terms of conjugation and other grammatical properties. Compare the activity word, *running*, in the sentences *The child is running* and *The child's running is fun to watch*. Only the first is a verb. The case is even more obvious when the activity word changes form with changes in grammatical structure, as with *move* in *The chess player moved the knight* versus *The move surprised the other player*.

This doesn't mean the environment never affects our judgments of structure. For example, we've already mentioned ambiguous sentences. Consider *They are observing responses*. Was this occasioned by a pigeon's key pecks in an experiment on attention (*Those responses are observing responses*: cf. Chapter 16) or by what some researchers were doing while looking into a pigeon chamber (*Those researchers are observing responses*)? Here's another: *Time flies*. Typically, as when this is said at a reunion, we call *time* a noun and *flies* a verb. But if it is said in a biology laboratory as an instruction to record how long some insects take to get from one place to another, *time* is a verb and *flies* is a plural noun. We can also resolve such ambiguities by expanding the structure: Consider *Time flies like an arrow* and *Fruit flies like a banana*.

We're left with a paradox. On the one hand are sentences with structures that allow us to name their constituents without knowing the circumstances in which they were uttered (*He who gulks merts*); on the other are sentences with structures that don't allow us to name their constituents unless we know those circumstances (*Time flies*). In other words, any account of grammar that is either exclusively structural or exclusively functional is necessarily incomplete.

In any case, and perhaps more important, our response to a sentence is typically not a matter of naming its constituents or drawing a diagram of its phrase structure. If someone asks you a question, you don't have to say which words are nouns and which are verbs before you answer. Children learn to speak and to understand sentences long before they are formally taught grammar and parts of speech. We must base our analysis of grammatical structure on something more than the capacity to name constituents or diagram structures. We must demonstrate the relation between sentence properties and the speaker's or the listener's behavior. Experiments that seek to demonstrate such relations are said to be concerned with the *psychological reality* of these dimensions of language (Fodor & Bever, 1965).

In one demonstration of the psychological reality of phrase structure (Fodor & Bever, 1965), listeners wearing earphones heard a sentence in one ear and a click in the other and were asked to locate where in the sentence the click occurred. The same tape recording produced all but the opening words in pairs of sentences such as *(In her hope of marrying)*^{X} *(Anna*^{Y} *was surely impractical)* and *(Your hope of marrying*^{X} *Anna)*^{Y} *(was surely impractical)*. The major sentence boundaries were in acoustically identical parts of the recording, so this ruled out effects of inflections or pauses in the reading of the sentence (Garrett, Bever, & Fodor, 1966). Instead of reliably reporting the click in the middle of *Anna*, listeners typically reported it displaced in the direction of major sentence boundaries: toward X rather than toward Y in the first sentence and toward Y rather than toward X in the second. Whether we interpret the results as displacements in heard locations of clicks or systematic errors in reports of the locations, they demonstrate effects of sentence structure on behavior. That's the sense in which saying that the structures have psychological reality is justified.

Transformations

Some sentences say different things while others say the same thing in different ways. In making such judgments, we discriminate among relations between sentences. For example, consider *He who hesitates laughs last*, *He laughs last who hesitates* and *He who laughs last hesitates* (you need to know what a shoemaker's last is before you can appreciate *He whose laugh lasts is a hallucinating shoemaker*). The first two sentences have something in common that neither has in common with the third. We say that the first two mean the same thing. They differ in word order, but the structural relations among their constituents are the same (e.g., *who* is similarly related to *hesitates* in both sentences). We therefore call the second sentence but not the third a transformation of the first.

Transformations describe changes in sentence structure that preserve certain relations among the constituents. When we transform a sentence from active to passive voice (e.g., "*The rat pressed the lever*" to "*The lever was pressed by the rat*"), we preserve the subject-object relation among the nouns and the verb. When we speak of the transformations that relate one sentence to another, we're discriminating some structural features the sentences share.

We can distinguish among structures by how they can be transformed. For example, consider *He is hard to understand* and *He is last to understand*. They differ only in the adjectives *hard* and *last*. They seem structurally similar, but we can transform the first to *To understand him is hard* but not to *He understands hard*; conversely, we can change the second to *He understands last* but not to *To understand him is last*. In the language of psycholinguistics, the sentences are called similar in *surface structure* (the particular order of constituents) but different in *deep structure* (the underlying structural features that distinguish them).

What advantage does speaking in terms of transformations have over just labeling sentences according to tense and so on? With transformations, as with constituent structure, we again face the problem of behavioral significance or psychological reality. One interpretation is that the transformations correspond to something listeners or readers actually do when responding to spoken or written sentences. If so, transforming a sentence from one form to another is a kind of behavior. Even though it cannot be observed directly, it may be possible to record its duration or its effects on other behavior.

In one demonstration of the psychological reality of transformations, readers were given two sentence lists (Miller, 1962). Each sentence in list 1 was related by some transformation to one in list 2, and the reader matched list-1 sentences to corresponding list-2 sentences. For example, pairs such as *John warned the small boy* and *The small boy was warned by John* made up two lists differing by kernel to passive transformations; *John warned the small boy* and *John did not warn the small boy* differed by a negative transformation. Time per transformation was estimated from the average time per correct match across different pairs of lists. This time was consistently shorter when lists differed by one transformation, in kernel to negative, kernel to passive and passive to passive negative, than when they differed by two, in kernel to passive negative and negative to passive (the transformation from kernel to passive negative breaks down into kernel to passive or to negative and then to passive negative, but that from negative to passive is indirect, by way of negative to kernel and then kernel to passive). These findings argue for recognizing transformations as properties of verbal behavior, but they don't explain verbal behavior. Describing the structural properties of sentences is the point of grammar, but describing a sentence doesn't say how it was produced or understood nor does its production or understanding require a judgment about its grammaticality.

Transformations change not just the form of sentences but also how the sentences or their parts can combine with each other. For example, *The canary sang* and *The cat ate the canary* can combine in *The canary that the cat ate sang* or *The cat ate the canary that sang*. These structures are called *recursive*, because the addition of segments can recur again and again. With the addition of *The cat grinned*, the sentence could become *The canary that the cat that grinned ate sang* or *The cat that ate the canary that sang grinned*. Strictly speaking, these sentences are grammatical, even if some of them don't seem quite right. We can expand sentences by adding phrases at their beginning or end, when their structures are called left-recursive or right-recursive, or by adding parts within or around them, when their structures are called self-embedding or self-enveloping. For example, *The canary that the cat ate sang* is embedded in *The music that the canary that the cat ate sang was off key*, where it is enveloped by *The music was off key*. Whether we regard this structure as self-embedding or self-enveloping depends in large part on whether the conversation started with *I'm glad that the music stopped* or *I'm glad that the cat ate*.

Relational Autoclitics: The Conjunction of Verbal Units

The syntactic properties we've been considering are primarily structural. They do not deal with the origins or the functions of syntax. Skinner's relational autoclitics provide an alternative treatment. Essentially, various dimensions of the environment contribute to an utterance, as when we might say *The dogs are running* or *The cat is sleeping*. The animal determines the noun, the activity determines the verb, and their number determines the singular or plural. But why not *Dog two run* or *Cat one sleep?* Skinner discussed these cases in terms of grammatical frames, essentially generalized structural classes in which various parts can be substituted, just as different particular imitations might occur in generalized imitation (e.g., Skinner, 1957, pp. 336, 346). Thus, any number of nouns can substitute for the dogs or cats of our example, and where in a frame words can fit contributes to whether we call them nouns or verbs or other parts of speech. These then are examples of multiple causation in verbal behavior. We considered some relevant examples in Chapter 13, so we will elaborate on them further only briefly here.

Some verbal responses specify events only through their relations to other verbal responses. For example, words such as *above*, *before* and *of* don't tact particular events. They almost always occur in combination with other verbal responses and they depend on those other verbal responses for their effects. Tense and other grammatical variations are also conditional upon properties of events tacted. A sentence in passive voice and past tense combines tacts of various relational and temporal properties of a set of events. For example, *The duck quacks* and *The bird flew* involve discriminations not only of animal and action but also of tense. In these examples we are more concerned with how the events around us combine to occasion what we say than whether it is grammatical when we say it.

These multiple dimensions are not very different from the ones along which stimuli occasioned response properties in the adduction example of Chapter 13. When environmental properties occasion classes of verbal responses that are invariant even when they occur in combination with other response classes, we can treat such classes as verbal units, as when the present-tense active-voice sentence structure remains invariant across a variety of different tacted events. Such units have complex structures, but we can generate novel verbal behavior under novel conditions only because we can combine verbal units in such novel ways; we can tact new events only on the basis of verbal behavior available with respect to things already known. Even if you've never seen a purple cow, separate tacts of color and cowness will allow you to say *Look, a purple cow!* when you encounter one.

These new combinations are important because our own verbal behavior often occasions later behavior; for example, I might act today on a reminder I wrote for myself yesterday. Sometimes my later behavior is verbal: I reword sentences, draw conclusions, derive solutions. These manipulations are of special interest when, as in problem solving, they lead to behavior that was not available to me earlier. Part of the power of verbal behavior resides in how it can occasion novel responding with important consequences (cf. Chapter 29 on problem solving).

Consider an example from mathematics, a convenient illustration because mathematical notation exactly prescribes the verbal responses that particular verbal stimuli should occasion. We judge the understanding of addition or multiplication by the number of ways in which someone can respond appropriately to relevant verbal stimuli. The person should be able to define the operations, to discriminate between cases where they apply and cases where they don't, to give answers to specific problems, and to derive each from simpler counting operations. Such behavior is verbal, and is necessary and sufficient for saying that someone understands addition and multiplication.

The learning of arithmetic involves intraverbal and autoclitic processes. Its particular advantages come when it combines with tacting of the numerosity of objects or events. A child might calculate the number of objects in a rectangular array by multiplying the number of rows by the number of columns or simply by counting all the objects. Either operation is verbal and the outcome is a verbal response that has a consistent relation to the quantity of objects in the array; it's a derived tact. The structure of arithmetic corresponds to the structure of the environment in such a way that new verbal responses generated arithmetically can function effectively as tacts; *twelve* tacts the number of eggs in a full box of a dozen eggs; we don't have to count the eggs every time.

Much important verbal behavior is derived from other verbal behavior. We mentioned the problem of entities that can't be tacted (e.g., how do we tact philosophy or biology or psychology?). We can progress from individuals and what they do and where we find them to groups and more general activities and broader areas, until we speak of academic institutions, governments, business organizations, religions, political parties, industries, branches of the military and so on. But we can't point to these entities, so we still can't tact them, even though they are related to events we do contact directly. Derivations of this sort, however, are much less explicitly defined than those of logic or mathematics, so correspondences between the world and what we say about it are likely to become less reliable as our verbal behavior gets further removed from its points of direct contact with the environment.

Derived verbal behavior also permits us to respond to properties of the world that we cannot respond to in other ways. We cannot tact noon or Saturdays or February 3rd or the 21st century. These exist only by virtue of clocks and calendars; they cannot stand independently of verbal behavior (cf. Austin, 1962, on speech acts). Examples include pronouncing a couple married or bestowing a title. Thus, these verbal responses acquire their functions only by virtue of the syntactic structures into which they can enter. The analogy between these cases and mathematical derivations explains verbal behavior only in the limited sense of showing how it works and how its properties differ from those of nonverbal behavior.

SEMANTICS: THE MEANING OF VERBAL UNITS

Our treatment of semantics will stress comprehension: What happens when someone is said to understand a word or a sentence? Let's begin with an experiment in which listeners heard a passage of text (e.g., an account of Galileo and the invention of the telescope), and were then asked whether a new sentence was one that had been in the passage (Sachs, 1967). The new sentence was either the same as the one in the passage (base sentence) or it differed in one of three ways: a change in word order that didn't affect grammatical structure (formal change); a change in grammatical voice (active to passive change); or a change in meaning (semantic change). The new sentence was presented either right after the original one in the passage or after 80 or 160 syllables of additional text. The passages and base sentences varied across different presentations and listeners. Examples of sentences used with the passage about Galileo are the following:

Base sentence:	He sent a letter about it to Galileo, the great Italian scientist.
Formal change:	He sent Galileo, the great Italian scientist, a letter about it.
Active to passive:	A letter about it was sent to Galileo, the great Italian scientist.
Semantic change:	Galileo, the great Italian scientist, sent him a letter about it.

When the new sentence followed the original immediately, listeners identified it as either identical or changed with better than 80% accuracy. When it was presented after 80 or 160 syllables of intervening text, accuracy decreased for all sentence types, but it stayed greater than 75% for semantically changed sentences whereas it dropped toward chance levels for the other types. In other words, listeners were likely to recognize a sentence as different only if its meaning had changed; as long as the meaning remained the same, they didn't notice formal or active to passive changes. Thus, the listeners weren't remembering words or word orders. They were remembering something more fundamental: the *gist* of the sentence, or whatever it is that sentences have in common when we say they mean the same thing.

The finding that listeners are more likely to remember semantic structure than specific words or sentences is robust (Bartlett, 1932; Fillenbaum, 1966). One study gave listeners related sentences such as *The ants were in the kitchen* and *The ants ate the sweet jelly* (Bransford & Franks, 1971). Later the listeners heard a mix of the original and new sentences and for each were asked to rate their confidence that they'd heard it before. Some new ones combined original sentences, as in *The ants in the kitchen ate the sweet jelly.* The listeners were usually more confident that they'd already heard the new combined sentences than that they'd heard the simpler ones that had actually been presented (cf. the intrusions in free recall in Addendum 21C). They'd learned something more abstract than particular words or sentences.

If meaning is a feature of verbal behavior, then, as with syntactic structures, we should be able to measure some of its properties. The *semantic differential* was one attempt (Osgood, Suci, & Tannenbaum, 1957). Words were rated along dimensions like *happy-sad, hard-soft* and *slow-fast.* Similarities among ratings were then determined by a statistical procedure, factor analysis, which created a space within which words could be placed. Words close together in this space were said to be more alike in meaning than ones far apart. Words like *good, beautiful, clean* and *pleasant* clustered together, distant from other clusters such as *bad, foul, dirty* and *ugly.* The semantic differential was intended to deal with any word in the lexicon. Other rating methods have concentrated on words in specific categories, such as emotions or probabilities (e.g., Reyna, 1981). Each method describes semantic relations among words, so meaning could be defined as just the relations measured by these methods, whatever their basis.

Events in a novel context can occasion novel grammatical utterances (as when a child gives the standard plural of a nonsense word), but utterances can also be semantically new (as when *drink-fruit* is coined as a name for a watermelon). We speak of such cases in terms of metaphor or analogy (Esper, 1973; Jaynes, 1976). Metaphor is effective only to the extent that relations among events in the world correspond to the relations discussed here as examples of semantic structure. If grammar is a description of relations among syntactic structures, then metaphor is the grammar of semantics (cf. Chapter 22).

LANGUAGE DEVELOPMENT

In their language development, children show an orderly progression from particular words to simple word groupings and then to syntactic and semantic structures of increasing complexity. The usage of irregular verbs, such as *go, come* and *break* (e.g., Kuczaj, 1977), provides an example. Among young children, regular but nonstandard past-tense forms are common: *goed, comed, breaked.* Children don't always start with the regular forms, however. Often, the child first learns the standard but irregular forms as individual words: *went, came, broke.* After the child learns some standard regular past-tense forms, the irregular forms are displaced by the regular but nonstandard ones, even though the irregular forms had been part of the child's vocabulary for some time. Months or years later, the standard irregular forms reappear and become permanent components of the child's verbal behavior. The progression from standard irregular to nonstandard regular and back to standard irregular past-tense forms is consistent with a progression from the mastery of individual words through stereotyped syntactic classes to fluent language.

Children who grow up in Italy speak Italian and those who grow up in Brazil speak Portuguese. Clearly they each learn the specific grammars and vocabularies of their native languages, and most children become fluent in them long before they begin formal education. Although there are many descriptions of language development in children, there is still much to learn about the factors critical to that development (e.g., Brown, 1973; J. G. De Villiers & De Villiers, 1978; Moerk, 1992). Between the ages of 1 and 6 years, the child's syntax progresses from single-word utterances to those of two words or more, sometimes described as roughly telegraphic (e.g., *stove hot, daddy go car*), and then to constructions that more and more closely approximate the syntax of adult speech. Problems of data collection range from those of recording and sampling children's speech at different ages to those of interpreting the speech phonetically, semantically, and syntactically. For example, a young child's *more* probably has the sense of *give me* and not that of expressing a quantity (Moore & Frye, 1986).

Among the controversial issues is the extent to which consequences play a role in the child's acquisition of language, and in particular in the child's acquistion of grammatical structure. In appeals to the *poverty of the stimulus,* some have argued that the child's verbal environment isn't rich enough to support language acquisition, and therefore that some structural features of language are *prewired,* in the sense that they'll emerge even in the absence of relevant contingencies (Culicover, 1992). Others have examined exchanges between parents and children and have argued instead that contingencies play an indispensable role (Moerk, 1980, 1983; Whitehurst & Valdez-Menchaca, 1988).

The case for the poverty of the stimulus argues that verbal environments do not include the negative or ungrammatical instances that should be there in support of claims that a child's grammatical behavior is shaped through natural contingencies, in the sense that such instances don't occur in the speech that children hear, or in the sense that such instances are not corrected when the child makes them, or in both senses. For example, if the muppet who is laughing is a frog, the child may hear questions in the form *Is the muppet who is laughing a frog?* but will not hear them in the form *Is the muppet who laughing is a frog?* (Crain, 1991). When the child then asks questions, they are in the former and grammatical form and rarely if ever in the latter and ungrammatical one. What keeps the child from making the latter kind of error? The simplest reply is that negative instances are not necessary for all kinds of learning. For example, our account of the echoic shaping of phonetic structure did not demand an environment that included nonnative as well as native speech sounds. With regard to syntax, any verbal environment including contrasts between several grammatical structures provides the differences on which learning can be based: for example, both active and passive voice are grammatical but provide a contrast that may enable later discrimination of grammatical versus ungrammatical (see also Palmer, 1996, 1998) .

The evidence on the poverty of the stimulus may go both ways (e.g., Moerk, 1992), but suppose for the sake of argument that the case had been made for grammatical universals by demonstrations that children can't learn certain types of sentence structures, or at least that they learn some types much more easily than others (cf. Pinker, 1984). Those universals would still involve structural rather than functional limitations, and it might even be appropriate to regard them as accidental by-products of systems that had evolved based on other functions (Gould & Lewontin, 1979). Human language has many of the properties of other evolved systems (Pinker & Bloom, 1990), but it is not obvious that it has been around long enough for the selection of those kinds of grammatical constraints.

Just as anatomical features of birds and bats determine the different ways they fly, special characteristics of our species may determine the structure of human language and how it develops. If we show that human language is limited in its structural properties or in how the structure develops, functional questions still remain about the circumstances in which verbal behavior occurs. Either way, environmental contingencies still matter. The anatomical analogy remains valid: An account of the different properties of flight in birds and bats doesn't bear on where or when they take off or land; so too an account of language structure doesn't bear on when we speak or what we talk about.

Another part of the issue is the appropriate level of analysis for verbal classes. The child must learn not only individual phonemes and words and sentences but also larger units such as phonetic and semantic and syntactic structures. Furthermore, the consequences of verbal behavior are often subtle and probably aren't ones that need to be explicitly arranged. Contrived reinforcers such as praise or candies may be less likely to be effective than such natural consequences of verbal behavior as hearing yourself say something similar to what you've heard others say, or getting something you've asked for, or hearing someone else say something relevant to something you've just said, and so on (but the contrived reinforcers are better than none at all). With regard to the question of whether language is innate or learned, the reasonable conclusion is that both phylogeny and ontogeny contribute.

If the significant consequences of verbal behavior range from such direct outcomes as getting something you've asked for to such indirect ones as hearing a remark relevant to something you've just said, these all contribute to the shaping of verbal behavior. If among these consequences we include the correspondences between sounds one has heard and sounds one has produced oneself, as in the ontogenic shaping of echoic behavior (Chapter 20), it is not too great a leap to extend such correspondences from phonemic to semantic and syntactic properties of verbal behavior. It's presumably important to discover that the relations among words and between words and things in our own behavior correspond to those relations in the behavior of others. In this view, by engaging in verbal behavior the verbal community provides the models the correspondence with which shapes and maintains the consistency of the child's verbal behavior (cf. Palmer, 1996, on parity).

Even if it were proved that children don't have to learn all the details of grammar because some are built in, that would not mean that there aren't lots of other things about verbal behavior that they'll still have to learn. In fact, there's plenty of evidence that rich verbal environments in which parents spend lots of time interacting verbally with their children can make vast and lasting differences in their verbal competence (Hart & Risley, 1980, 1995; Moerk, 1992). Contingencies matter. The effects, described as meaningful differences in the everyday experience of young American children, appear in vocabulary growth rate, school performance and IQ scores. What else that has emerged from the analysis of behavior could have greater relevance?

ADDENDUM 25A: SOME PROPERTIES OF LANGUAGE

Properties assumed to be characteristic of all human languages have been called *language universals*. Various grammatical relations have been proposed as such universals (e.g., Greenberg, 1966). For example, it has been argued that utterances in all human languages have subject-predicate structure, distinguish singular from plural and are limited in the transformations that can operate on embedded structures. These properties may depend as much on the conditions under which humans speak as on biological constraints on the kinds of sentences they can produce. The subject-predicate relation is one criterion for calling an utterance a sentence, so it excludes by definition utterances without subject or predicate (e.g., *Hello, Ouch!, Oh?* and *Aha!*). Differences between single objects and collections are presumably important in all human environments and to that extent determine the distinction between singular and plural; that distinction appears in all languages, but in some it's expressed by vocabulary rather than grammar (as in *one book, two book, many book*; cf. Japanese). And limits on levels of embedded phrases may have arisen as derivatives of aspects of remembering such as the limited span of immediate memory (see Chapter 26).

The inverse of the issue of language universals is that of *language relativity*. Much has been made of the many Eskimo names for snow, but the data don't compare with the claims (Pullum, 1991); anyway, skiers have many names for snow, too. It should not surprise us that vocabularies are consistent with functionally important classes of events in different language communities. Language relativity gets more interesting when it seems to involve allowable grammatical relations rather than vocabulary categories. For example, languages in which nouns and verbs are exchangeable (thought versus thinking, memory versus remembering) may lead to different treatments of events and actions than those in which they are not.

Chinese grammar doesn't include convenient forms of the counterfactual conditional (*if A hadn't occurred then B would have*, or *if not A then B*) whereas English grammar does; on the other hand, English grammar doesn't include convenient forms of the exclusive and the nonexclusive *or* (*A or B but not both* versus *A or B or both*) whereas Chinese grammar does. Mastery of these two logical forms of *or* undoubtedly proceeds differently in these two language communities (cf. Braine & Rumain, 1981). Language differences are obvious when they involve formal properties, as in the artificial languages of symbolic logic, calculus and computer programming. To the extent that language is behavior, particular languages will inevitably have different functional properties. Language relativity reminds us that we must deal with each language in the context of the environment within which it was shaped.

ADDENDUM 25B: DEIXIS

One significant feature of the development of a child's language is the evolution of a *deictic* vocabulary; deictic terms are occasioned not by intrinsic properties of events or objects but rather by their relation to the speaker and listener (P. A. De Villiers & De Villiers, 1974; Wales, 1986). Examples are *here* versus *there*, *this* versus *that* and *in front of* versus *behind*. In each case, the appropriate term depends on where one is located; for example, when you are cooking, the kitchen is *this* room and the dining room is *that* one; but when you are eating, the terms are reversed.

The acquisition of deixis follows closely upon that of other relational vocabularies (e.g., *big* and *little*, which involve relative rather than absolute size; the child is big relative to a pet frog but little relative to an adult). In combination with pronouns, the functions of deictic terms in language are analogous to those of variables in algebra; we can speak of things even if we can't name them: *what's this?*, *who was there?*, *is that it?* (cf. Clark & Sengul, 1979).

A special case of the deictic vocabulary is the personal pronouns (e.g., Charney, 1980; Huxley, 1970). Children generally learn *it* before *me* and *you*, and the distinction between first and second person (*me*, *you*) emerges before distinctions within those classes (*I*, *me*, *my*, *mine* and *you*, *your*, *yours*). But *I* and *me* aren't learned like one's name; one is called *you* rather than *I* or *me* when spoken to. How then do children come to say *I* and *me* appropriately when they become speakers instead of listeners?

It is hard to observe the conditions under which this feature of language develops. Personal pronouns are mastered by almost all children (Chiat, 1982). We can describe how the mastery evolves, but we don't know enough of its details to say what aspects of the child's verbal environment are crucial to that evolution. Given the great variability in the ways parents interact with their children, the properties that lead to deixis are probably common to human environments in general.

Personal pronouns and the deictic vocabulary involve discriminations among events as they are related to oneself, and so they're presumably related closely to the discriminations of one's own behavior discussed in the context of autoclitic processes and the concept of self-awareness (cf. Damon & Hart, 1982). Such discriminations are often expressed in mental language, as in the various ways of describing the difference between remembering something and having forgotten it and never having known it (cf. Skinner, 1945). These discriminations have an intrinsic hierarchical structure that presumably determines the order in which they will be learned. For example, a child will probably learn *remember* before learning *forget*, and *knew* before *never knew*. We can describe how and when this language develops (e.g., Wellman, 1990), but one way to find out more about it would be to see what it takes to teach it explicitly.

Chapter 26
The Functions of Remembering

The central metaphor of remembering involves storage, retention and retrieval, and what is stored, retained and retrieved is sometimes called a representation. Representation *combines the Latin roots* re-, prae- *and* esse, *which together have the sense of existing in front of again. Via the verb* esse, *to be, it is related to* is, yes, entity *and* interest.

By way of the Latin instaurare, *to set upright,* storage *comes from the Indo-European root* sta-, *to stand; it has a broad range of relatives, including* circumstance, exist, history *and* system. Retention *combines the Latin* re-, *back or again, and* tenere, *to hold; it is related to* maintain, continue, tenure *and* attention. Retrieval, *from the Old French* retrover, *which combines* re- *with* trover, *to find, is related to* troubador, controversy *and* tropism, *which in turn have connections with the Latin* tropus, *song, and the Greek* tropos, *a turn or figure of speech. Those early forms had a sense of composition or invention, which is fitting in the context of contemporary accounts of remembering as reconstruction rather than reproduction.*

- Mnemonics
- The Metaphor of Storage, Retention and Retrieval
  - *Storage: Encoding and Levels of Processing*
  - *Retention: The Question of Memory Reorganization*
  - *Retrieval: Cue dependency and Accessibility*
  - *Remembering, Cue-Dependency, and Stimulus Control*
  - *Hearsal and Rehearsal*
- Metamemory
  - *Running or Working Memory*
  - *Discriminated Remembering*

Memory is the popular term for the topics we'll treat here, but I've chosen *remembering* for the chapter title. Woodworth has outlined the rationale for preferring one usage over the other:

> Instead of "memory," we should say "remembering"; instead of "thought" we should say "thinking," instead of "sensation" we should say "seeing, hearing," etc. But, like other learned branches, psychology is prone to transform its verbs into nouns. Then what happens? We forget that our nouns are merely substitutes for verbs, and go hunting for the *things* denoted by the nouns; but there are no such things, there are only the activities that we started with, seeing, remembering, and so on.... It is a safe rule, then, on encountering any menacing psychological noun, to strip off its linguistic mask and see what manner of activity lies behind. (Woodworth, 1929, pp. 5-6)

KEY TERMS: Mnemonics; Remembering, Metamemory; Running or Working Memory; Storage, Retention, Retrieval; Cue or State Dependency; Encoding, Hearsal and Rehearsal

In this chapter I'll emphasize the functions of remembering and in the next I'll emphasize its structure, though these topics often overlap. The study of remembering is concerned with how an organism's present behavior can be occasioned by past events, as when a delay is imposed between a stimulus and an opportunity to respond. Accounts of remembering often speak of what the organism does when the stimulus is presented as *memory storage*, of the intervening time as the period of *retention*, and of what the organism does when the response later occurs as *retrieval from memory*. Systematic relations between stimuli to be remembered and the responding occasioned by these stimuli are often discussed as *encoding*. For example, the learner who recites words aloud as they're presented visually is said to be encoding the visually presented words in a vocal mode.

A witness to an accident later describes what happened. How do we deal with verbal responses that depend on stimuli no longer present? What are the functions of this behavior and what structure does it have? Does describing events as they occur differ in important ways from describing them later? When in Chapter 22 we defined tacts as restricted only to those verbal responses in the presence of or very shortly after the tacted stimulus, we assumed the difference was important.

Once again let's start with a pigeon example. Our rudimentary example of tacting used a chamber in which a pigeon's pecks on a left or middle or right key were respectively reinforced during red or blue or green light in a window beside the keys. The window is lit with a color and when the keys are lit white the pigeon has an opportunity to peck the appropriate key. Once the pigeon is performing accurately, we continue to present colors briefly in the window while the keys are dark, and we light the keys only after the color is gone. If red appears for just a moment and then the pigeon immediately pecks the left key, we'll still probably say its peck was occasioned by the brief red stimulus even though that stimulus was gone by the time the peck occurred.

Now let's impose a 2-s delay between the brief presentation of red and the pigeon's opportunity to peck. If the pigeon still pecks the left key, we might still say that the peck is occasioned by the red stimulus. So far we have no reason to assume that the behavior is different in kind from that when pecks occur during the color. Imposing a delay between a stimulus and an opportunity for the response it occasions doesn't necessarily alter the control of the response by that stimulus.

In fact, control of the pigeon's pecks by color diminishes even with delays of 1 or 2 s; responding will probably be near chance levels with delays of 5 s (Blough, 1959). But maybe we can teach the pigeon to remember. Let's shape different performances after red and blue and green and then chain each one to pecking on an appropriate key. For example, we might get the pigeon to peck the left chamber wall after red and then get it to peck the left key if it's still pecking the left chamber wall when the keys light up; (cf. *mediating behavior*, Chapter 10). If each color occasions pecking in a different location, the duration over which the pigeon remembers a color will depend only on how long it can maintain its pecking in that location during the imposed delay. Under such conditions, control by the colors might extend over delays of many seconds and perhaps even minutes. We could then study how the pigeon's remembering depended on the mediating behavior by interrupting the behavior during the delay (e.g., Jans & Catania, 1980).

But maybe we shouldn't call the pigeon's performance remembering if the temporal gap between stimulus and response is bridged by uninterrupted mediating behavior. Somehow it seems more appropriate to say the pigeon's pecks are occasioned by its previous behavior rather than by the now-absent colors as stimuli. Yet we modify our remembering by keeping calendars and appointment books. If I looked up a phone number and then recited it over and over to myself until I had a chance to copy or dial it, I'd probably say that I remembered it even though my use of the number depended more on my vocal repetitions than on the number as I'd seen it in some directory. My vocal repetition illustrates behavior sometimes called *rehearsal*; it justifies the view that **what we remember isn't so much stimuli as it is our own behavior with respect to those stimuli**: "what is reproduced on all occasions after the first is not the original but one's own reproduction of it" (Zangwill, 1972, p. 130).

Questions about what memory is can often be translated into questions about what it is that we remember. No doubt ongoing research in the neurosciences is revealing important details of the physiology of remembering, as neuroimaging techniques such as positron emission tomography (PET scans) identify brain regions active during different sorts of remembering (e.g., Squire, 1992). These techniques may be useful for analyzing memory tasks, as when different areas are active during different tasks or components of tasks. But a researcher who seeks some memory copy of a remembered event will be disappointed in that quest if remembering depends not on the event itself but rather on behavior with respect to the event.

MNEMONICS

Remembering is a complex subject with an extensive history. We began with a hypothetical pigeon experiment. Now let's shift to a contrasting example, as we consider the complexity of human remembering in the context of memory strategies called *mnemonic* systems. Mnemonics are techniques for increasing the likelihood of remembering. Remembering may involve behavior that persists over time, as in rehearsal, but clearly that is just part of the story. If I remember an event that took place yesterday or last week or last year, I can't argue that somehow I've been rehearsing it without interruption ever since then.

Mnemonic systems make the learner less dependent on rehearsal, which is a major component of memorizing material by rote. A simple example of a mnemonic technique is the conversion of a sequence of symbols into a sentence, as when a beginning student of music remembers the notes on the lines of the treble clef, *EGBDF*, as the sentence "*E*very *g*ood *b*oy *d*oes *f*ine."

One well-established mnemonic technique is the *method of loci* (places). It is attributed to the Greek poet Simonides (Yates, 1966). He is said to have left a banquet hall just before the roof collapsed and killed all the occupants. Although the bodies were unrecognizable, he was able to identify them for their relatives by where they had been sitting at the table. This demonstration that an orderly spatial arrangement contributed to accurate remembering is supposed to have led Simonides to invent the method of loci. It became the basis for remembering sequences of topics in speeches and was described by the Greek orator Quintilian:

> In order to form a series of places in memory, he says, a building is to be remembered, as spacious and varied a one as possible, the forecourt, the living room, bedrooms, and parlours, not omitting statues and other ornaments with which the rooms are decorated. The images by which the speech is to be remembered...are then placed in imagination on the places which have been memorised in the building. This done, as soon as the memory of the facts requires to be revived, all these places are visited in turn and the various deposits demanded of their custodians.... The method insures that the points are remembered in the right order, since the order is fixed by the sequence of places in the building. (Yates, 1966, p. 3)

For example, a student might imagine a systematic walk from one campus landmark to another: dormitories to dining hall to library to computer center to gymnasium and so on. To learn the items of some ordered series, the student imagines each item at each successive location. To recall the series, the student then takes the imaginary walk again, remembering each item in its appropriate place. To learn a new sequence later, the student repeats the imaginary walk in the same order, this time visualizing the new items in their places. The new series will be learned with relatively little interference from the first one (e.g., Bellezza, 1982), but the first series may no longer be well remembered. Thus, the method is useful mainly for series that only need to be remembered temporarily (e.g., shopping lists).

These mnemonic techniques work best with easily visualized or imagined sequences of concrete items. Unusual or bizarre items or combinations may have advantages over common ones, but even more important is the spatial closeness or connection of the items and places. For example, in learning pairs of objects, learners more accurately remembered objects standing in some relation to each other (e.g., one sitting on top of the other) than if they were simply paired side by side (Wollen, Weber, & Lowry, 1972).

Various mnemonic devices were developed during Greek and Roman times. In the Middle Ages they became methods for remembering details of religious matters (e.g., stations of the cross: Yates, 1966). They gradually became the basis for ritualized forms of religious art, as when particular figures were used to represent vices and virtues. In the course of this evolution, the mnemonic origins of these art forms were gradually (and ironically) forgotten. Mnemonic techniques were given relatively little attention throughout much of the history of psychology; interest in them is a fairly recent development (e.g., Bower, 1970).

Some techniques are designed so that abstract items can be converted into concrete ones. The *peg-word* technique translates numbers into a sequence that can be visualized, as in *one is a bun, two is a shoe, three is a tree*, etc. An ordered list can then be learned by imagining each item together with the object corresponding to its numerical position in the list. These systems improve on the method of loci, in that the learner can recall the item in any position without having always to start from the beginning (e.g., to recall the third item, the learner has only to remember what was imagined with *tree*).

A more elaborate system provides a code for translating numerals into consonants. This system has been a part of popular mnemonic techniques for more than a century (e.g., Loisette, 1899). One version is the following:

Number	*Consonants*	*Rationale*
1	*t, d*	*t* has one downstroke
2	*n*	*n* has two downstrokes
3	*m*	*m* has three downstrokes
4	*r*	*r* is the fourth letter of *four*
5	*l*	*l* is the Roman numeral 50
6	soft *g, j. ch*	in script, g is an upside-down 6
7	*k*, hard *c*	*k* can be combined with 7: *ꝁ*
8	*f, v*	both 8 and script *f* have two loops
9	*p, b*	backwards *p* or upside-down *b* is 9
0	*z, s, sh*	*z* is the first letter of zero

I've used this version for classroom demonstrations over many years of teaching. My standard arrangement was to have students contribute twenty two-digit numbers between 21 and 99; I excluded the numbers from 1 to 20 because I wanted to use them for the twenty positions of the contributed numbers. While coming up with my own codings, I also had students independently contribute their own. For example, if the second contributed number was 53, I suggested Noah (2 = n) for the position code and lamb (53 = l + m) for the number code (notice that the coding is based on sound, not spelling). Noah-lamb is much easier to remember than 2-53, especially if you visualize Noah on the ark waving goodbye to a lamb who has been left behind. By the time we had gone through the first ten contributed items, I demonstrated that based on my codings I could recall virtually all twenty items in any order called for. In one semester I successfully learned twenty items in a late morning undergraduate class and twenty different items in a late afternoon graduate class on the same day. More important, most students found they could recall all or almost all of the ten numbers they had coded, so I had successfully made the point that remembering is a skill that we can get better at.

I must confess that I never got as good with student names as I did with arbitrary numbers. Nevertheless, with a code like this, any number can be converted into a word or a sequence of words in which consonant sounds correspond to successive digits. With a little practice, learners can quickly translate numerical information, such as dates or telephone numbers, into a form that is easy to remember. The potential applications are limited only by the learner's ingenuity.

Consider an example. I want to talk about the visual spectrum in a lecture on generalization gradients, but sometimes I forget which end is infrared and which is ultraviolet. I therefore construct a colorful scene from early in the American Civil War. A Union soldier, in his **blue** uniform and with his steel-**blue** *rifle* over his shoulder, is in a country kitchen taking leave of his sweetheart. Golden **yellow** *loaves* of fresh-baked bread are sitting on the counter as he *kisses* her on her **red** lips. *Rifle* translates to 485, roughly the wavelength of blue in millimicrons; *loaves* translates to 580, or yellow; and *kisses* to 700, or red. Infrared, then, is at the long wavelength end of the spectrum. It's also easy now roughly to determine the wavelengths of other colors (e.g., 530 is a good guess for green, which is located between blue and yellow).

This procedure may seem like overkill, but with a little practice you should find that you can code and learn arbitrary sequences with relatively little effort. There is no reason not to have the letter-to-number code handy while you are learning it. Arbitrary memory feats may impress your friends, but whenever you encounter sequential information worth remembering (e.g., combination lock settings, computer passwords), you can put mnemonics to practical use. More important, they demonstrate the flexibility and capacity of human memory. Educational systems have tended to emphasize learning through understanding and have correspondingly de-emphasized or even discouraged memorizing. It's unlikely, however, that a learner will be disadvantaged by learning in more than one way, so mnemonics can effectively supplement other methods of study. Mnemonics, a far cry from the pigeon example discussed earlier, illustrate how varied the phenomena of remembering are. They are classes of behavior that can be learned, and they show that **what the learner remembers depends on what the learner does.**

THE METAPHOR OF STORAGE, RETENTION AND RETRIEVAL

An episode of remembering is defined by three components: the initial learning of an item, the passage of time, and then an opportunity for recall. In the research literature on memory, a metaphorical treatment of these three components has gradually evolved into a technical language. Initial learning is said to result in *storage* of the item, which determines how the item is retained over time; a period of *retention* is followed by the opportunity for recall; recall of the item is then called *retrieval* from storage. Other memory metaphors exist (cf. Roediger, 1980). For example, some theories have appealed to the metaphor of resonance, as when something reminds us of something else and we say "that rings a bell." The metaphor of storage and retrieval, however, has been the most influential in determining the direction of memory research.

The stages of storage and retrieval are analogous to storing information in computer memory and retrieving it later. The language of storage and retrieval is an effective analogy for what happens in remembering, but it remains a metaphor. In fact, one objective of some memory research is to explore the limits of the metaphor. One way to describe the functional properties of remembering is to determine the range of conditions over which the metaphor holds. This is a matter of function not so much because it shows us what remembering does (the functions of remembering facts for an exam or how much money someone owes you are pretty obvious), but rather because it shows us how remembering works.

According to this metaphor, a remembered item is one that's been stored and retained and retrieved. The failure to remember an item may occur because the item wasn't stored in the first place, or because the item was stored and then lost from storage during retention, or because the item wasn't retrievable at the opportunity for recall (Watkins, 1990, 1996). As we examine how remembering can be affected by events during these three phases, we may have something to say about what it is that metaphorically gets stored and retained and retrieved.

Storage: Encoding and Levels of Processing

When you remember a stimulus, what you remember isn't so much the stimulus itself as your response to it. That response inevitably differs from the stimulus even when the two are in the same modality (as when you repeat a spoken verbal item aloud). Remembering, in other words, isn't just reproducing the stimulus. Even the immediate response to a stimulus can't be interpreted that way. The issue is long-standing. For example, this passage from the Greek philosopher Theophrastus dates from about 300 B.C.:

> with regard to hearing, it is strange of him [Empedocles] to imagine that he has really explained how creatures hear, when he has ascribed the process to internal sounds and assumed that the ear produces a sound within, like a bell. By means of this internal sound we might hear sounds without, but how should we hear this internal sound itself? The old problem would still confront us. (Stratton, 1917, p. 85)

A more contemporary version is the following:

> Suppose someone were to coat the occipital lobes of the brain with a special photographic emulsion which, when developed, yielded a reasonable copy of a current visual stimulus. In many quarters this would be regarded as a triumph in the physiology of vision. Yet nothing could be more disastrous, for we should have to start all over again and ask how the organism sees a picture in its occipital cortex. (Skinner, 1963; see also Skinner, 1976, p. 74)

As in the analysis of stimulus control, the problem of remembering isn't to be solved by trying to follow the stimulus into the organism; rather, we must discover how to characterize the organism's behavior with respect to the stimulus; "it is not sensible to inquire about the characteristics of the memory trace when remembering is not occurring" (cf. Craik, 1985, p. 200).

The learner's behavior with respect to the stimulus to be remembered is called *encoding* (Melton & Martin, 1972). Consider the following experiment (Conrad, 1964). In one part, learners named spoken letters presented in background noise; in another, they saw a sequence of six letters and wrote them down in order. When learners erred in the first task, they did so along dimensions of common acoustic properties: they were more likely to confuse *V* with rhyming letters such as *B* or *C* than with letters such as *N* or *X* that have straight lines in common with *V*. The second task used visually presented letters but produced the same kinds of acoustic errors as the first: they were more likely to confuse *D* with the similar-sounding *T* than with letters such as *O*, which has curvature in common with *D*, or *P* or *B* or *R*, all of which share with *D* a vertical on the left and an area enclosed by a curve on the right. Whether the stimuli were auditory or visual, learners erred more often along dimensions of common acoustic than common visual properties. These errors must have occurred because the learners encoded the stimuli acoustically rather than visually. In some form or other, the learners were saying the letters to themselves.

Remembering depends on how the items to be remembered are encoded. For example, some tasks favor encoding based on semantic properties (defining technical terms); others favor encoding based on visual or phonological properties (learning spellings or pronunciations); still others favor encoding based on tonal or temporal properties (following the score of a musical work). Encoding can vary from time to time and within or across tasks, and it can be based on combinations of properties as well as on single dimensions. It can be as simple as the repetition of the item (sometimes called *maintenance rehearsal*) or as complex as an extensive mnemonic system (sometimes called *coding rehearsal*).

Two major classes of encoding are called *substitution* and *elaboration*. Substitution corresponds to familiar examples of codes, as in Morse-code learning (e.g., Keller, 1958). Consider learning a sequence of the binary digits, 0 and 1 (Miller, 1956). The sequence 010001101011101001 substantially exceeds the number of digits that can be remembered after a single presentation. Each group of three binary digits, however, can be replaced by a single octal digit, according to the following list:

000=0 001=1 010=2 011=3
100=4 101=5 110=6 111=7

The binary sequence then can be coded as the octal sequence 215351 (= 010 001 101 011 101 001), which can be remembered after just one presentation. Reducing the number of items to be remembered by encoding groups of items is called *chunking*, much as intraverbal sequences can become integrated units. The acoustic encoding of written letters can also be regarded as substitution in that it involves unique correspondences between spoken and written letters.

In elaborative encoding, an item to be remembered is transformed into a word or phrase or other complex unit. For example, a learner might rehearse the consonant sequence *QBF* as *Quick Brown Fox*. We've already discussed visual elaborative encoding in the context of mnemonics (cf. Paivio, 1971). Elaborative encoding doesn't guarantee the unique correspondences that characterize substitution, so encoded items may be more likely to occasion inappropriate responses. For example, the learner using the method of loci who encodes the word *baggage* by imagining a tower of suitcases standing in a parking lot might instead say *luggage* at recall.

Encoding, whether by substitution or elaboration, is inevitably selective. Some stimulus properties are more likely to occasion responses than others. For example, written words are very likely to produce the kind of encoding called reading; it isn't easy to look at a word without reading it (cf. Stroop, 1935). Can you obey the instruction "Don't read this sentence"?

Simply naming something may qualify as encoding. Two experiments involved recognition tasks with rhesus monkeys and with humans (Cook, Wright, & Sands, 1991; Wright *et al.*, 1990). The stimuli included kaleidoscope pictures or travel slides. Each stimulus was presented briefly on a screen; responses were the movement of a lever to the left or to the right. The monkey behavior was maintained by food reinforcers; the human behavior was maintained by a tone produced by correct responses (for convenience, we'll treat it as a reinforcer, though instructions and feedback presumably were also important variables). For both monkeys and humans, the task involved successive presentations of six pictures at the top of the screen and then a single picture at the bottom. If the final picture was different from the previous six, a lever movement to the left was reinforced; if it was the same as one of them, a movement to the right was reinforced. In other words, a correct response to the right corresponded to recognizing the final picture as one that had appeared among the six.

The six pictures were presented at interstimulus intervals of 0.08, 1 or 4 s. With either the kaleidoscope or the travel pictures, the monkeys were most accurate at the shorter intervals; the more slowly the pictures appeared, the more likely the monkeys were to forget earlier ones. The human performances were similar with the kaleidoscope pictures, but with the travel pictures the humans became more rather than less accurate as the pictures were presented more slowly. One interpretation was that slower presentations allowed verbal encoding and rehearsal of the travel pictures but weren't helpful with the kaleidoscope pictures because they had no names (cf. Intraub, 1979). Thus, the next stage of the experiment with humans was to teach them arbitrary names for each of the kaleidoscope pictures. The recognition task with the kaleidoscope pictures was then repeated, and the human performances became similar to those with the travel slides: The slower the presentation rate, the more accurate they became. The implication is that the names allowed them to encode and rehearse the kaleidoscope pictures.

The results are of interest both for the species difference they demonstrate and for their relevance to the role of naming in remembering (see also Wright, 2012). For the present purposes, the main point is that naming itself functions as a variety of encoding. That's probably why we are less likely to recall events from our infancy than from later times in our lives; if we did engage in encoding before we became verbal, that coding must have been very different from our later verbal encoding. Remembering in the very young is more easily demonstrated in the recall of the child's own actions than in other ways (cf. Bauer, 1996).

The several categories of encoding we've considered are neither exhaustive nor mutually exclusive. Mnemonic systems, for example, can combine substitution, as in the number-to-consonant code, with elaboration, as in visualizing objects according to the method of loci. We shouldn't expect an exhaustive list, because types of encoding are as unlimited as the different ways in which we can respond to events in the world. For the same reason, we shouldn't expect any one type of encoding to be invariant across different learners or different tasks. Different learning histories and different contingencies upon remembering guarantee variability in the ways that we encode stimuli.

Some kinds of encoding seem more superficial and therefore less memorable than others; this dimension of encoding has been called *level of processing* or *depth of processing* (Cermak & Craik, 1979). Fewer words from a list are recalled after tasks requiring responses to formal properties of words (e.g., crossing out vowels) than after tasks requiring semantic responses (e.g., assigning words to categories). The implication is that the deeper the level of processing, the more likely an item is to be remembered. We're certainly more likely to remember items to which we've responded in rich and novel ways, but as with so many other concepts we've encountered, level of processing is descriptive and not explanatory. It describes relations between types of encoding and the likelihood of remembering; it doesn't explain the remembering.

Retention: The Question of Memory Reorganization

According to the metaphor of storage and retrieval, after an item has been encoded it may be stored. But how does it get stored and what happens to it after it has been stored? At this point the metaphor of storage and retrieval doesn't help much. Retention is implicit in the concept of storage, but we can't expect to find the to-be-remembered item inside the learner. It is only because the item emerges at recall that we assume that it somehow existed throughout retention. That may be a bit like assuming that a pianist at a keyboard releases Beethoven's sonatas from their storage in the piano or even, more simply, that the pianist releases sounds that have been stored in the strings; whoever searches for the sonatas or the sounds by disassembling the piano will be disappointed.

Contemporary accounts of remembering deal with it not in terms of *reproductive* processes, in which events are recalled directly or reproduced, but rather in terms of *reconstructive* processes, in which aspects of past events are derived or reconstructed from what had been encoded (e.g., Hasher & Griffin, 1978). For example, we may construct a narrative about some family event based not on specific details that we remembered but rather based on the consistent features of different stories about the event that we've heard from family members. One crucial thing to remember is that **remembering is reconstruction, not reproduction.**

Accounts of remembering have included theories of *consolidation* , which argue that what's learned becomes fixed or consolidated in memory over some time following learning, and theories of *incubation*, which argue that remembered events and relations are spontaneously reorganized over time (perhaps especially during sleep), sometimes so that their combination constitutes the solution to a long-standing problem, as in some anecdotal examples of scientific creativity (Hadamard, 1949). On the other hand, the evidence for subliminal learning or learning without awareness during sleep isn't persuasive. For that matter, neither is the evidence for learning without awareness during waking (e.g., Pratkanis, 1992).

Studies of memory change during retention have used reproductions of remembered material over time (e.g., successive retellings of stories; (Bartlett, 1932). Figure 26-1 provides an example using visual figures (Carmichael, Hogan, & Walter, 1932). The remembered figure changed over retention as a function of what it had been called. But these changes might have occurred during encoding or retrieval instead of during storage.

Events that occur between storage and retrieval, however, can affect remembering. For example, consider the following simulation of eyewitness testimony in a courtroom (Loftus & Palmer, 1974). Observers who viewed a film of an automobile accident were later asked to estimate the speed of the cars when they collided. In the wording for one group, the question was how fast the cars were going when they *hit* each other; for the other, it was how fast they were going when they *smashed into* each other. The wording made a difference. Speed estimates in the first group averaged about 8 miles per hour whereas those in the second averaged more than 10. More important, when the observers were asked a week later whether there had been broken glass at the scene of the accident, those from the second group were more likely to say, incorrectly, that there was. Events during retention affect subsequent remembering (see also Loftus, 1993).

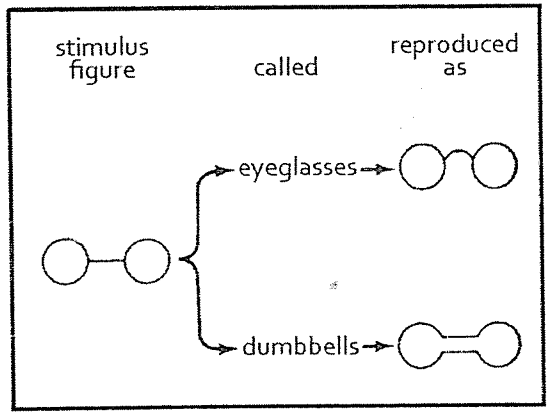

Figure 26-1 Examples of changes in remembered visual forms as a function of verbal labels. (Adapted from Carmichael *et al.*, 1932, Figure 1)

The effects of the wording of questions on recall raise questions about the reliability of eyewitness testimony. For example, everything else being equal, an eyewitness asked "Did you see the broken headlight?" is more likely to say yes than one asked "Did you see a broken headlight?" (Loftus & Zanni, 1975). The observer's verbal behavior under such questioning is rehearsal of sorts, but has the disadvantage that it is initiated long after the event. Given that we can affect remembering by interrupting or otherwise interfering with rehearsal shortly after an event, we should not be surprised that we can also do so by distorting or interfering with it later on. Some distortions of remembering may be based on the timing of the remembered event and a later recall that introduces something new; the domination of one over the other may reverse depending on their separation in time (e.g., Riccio, Rabinowitz, & Axelrod, 1994).

These cases are not so much concerned with retention as with successive retrievals. Perhaps that's why we usually leave out retention when we speak of the metaphor of storage and retrieval. However we look at these phenomena, they raise significant issues. For example, asking children about past events can create false memories that are more enduring than real ones (Brainerd, Reyna, & Brandse, 1995). If so, how should we interpret reports of other kinds of memories, which might range from sexual abuse in childhood to kidnapping by the alien crews of UFOs? The reality of such repressed memories has raised troublesome legal and ethical issues (Loftus, 1993). If someone claims to have just recalled an unsolved murder that occurred two decades ago, the importance of closing the case, in the absence of other kinds of evidence, must be weighed against what we know about remembering. How seriously can we take such testimony? If the sudden recall by an adult of sexual abuse in childhood is based on real events, it would be tragic to ignore it. But it would also be tragic to put innocents in prison based on testimony that is unsupported by such events.

This is far from hypothetical. My wife was once reading a magazine article about people who were exonerated of child abuse after years in prison (Jones, 2004) when she recognized the name of a childhood family friend. She learned some of his story after contacting him. With the help of the Innocence Project he was eventually cleared after many years in prison. When he had become a parent, his son and some of his son's friends were repeatedly interviewed about child abuse in their Bakersfield, California, community, in what eventually came to be recognized as a notorious example of a child abuse witch-hunt initiated by overzealous prosecutors (cf. Bottoms, Najdowski, & Goodman, 2009). He and several other parents were tried for and convicted of child abuse. The convictions were eventually overturned but only after extensive legal proceedings. It became evident that false testimony had been coerced from the very young child witnesses. The contingencies included keeping them from their parents until their stories changed so as to be consistent with scenarios created by the prosecution. As young adults long after, some now with families of their own, almost all of those children recanted their testimony, but for a few of them the line between what they had learned to say about their parents and what they thought they remembered remained a difficult one to draw (Coleman & Clancy, 1999; Goodman, 2006).

An interview, whether by attorney, reporter or psychiatrist, too often lends itself to the shaping of verbal behavior by the interviewer, perhaps deliberately or perhaps inadvertently (cf. Chapter 24). As we have seen, remembering by children may be especially susceptible to such shaping. With some verbal behavior to start with, such as an account of strange lights or some other unusual sighting, an interviewer, simply by reacting more enthusiastically to some parts of the descriptions than others, can shape consistent stories across individuals who have never met each other. Over time, the descriptions may come to include details of flying saucers and their occupants. Once that has happened, plenty of social consequences (some monetary) can keep this verbal behavior going and discourage reports of where it came from. If such distortions of remembering seem far-fetched, we need to be reminded that they are only more extreme cases of robust and well-documented phenomena of remembering. As such, they are surely not as far-fetched as UFO abduction stories. It is easy to make people think they remember words that never actually appeared on a list (e.g., intrusions in free recall: Chapter 21). Why should we be surprised if we sometimes encounter cases that are far more extreme?

Retrieval: Cue-dependency and Accessibility

False remembering isn't really remembering, so let's now consider remembering when it actually works. We introduce retrieval with a quotation:

> In order to understand retrieval processes, some basic principles must first be accepted. One of the most important of these was formulated by St. Augustine more than 1500 years ago; we cannot seek in our memory for anything of which we have no sort of recollection; by seeking something in our memory, "we declare, by that very act, that we have not altogether forgotten it; we still hold of it, as it were, a part, and by this part, which we hold, we seek that which we do not hold." (Tulving & Madigan, 1970)

In other words, we aren't likely to remember an item or event in the absence of discriminative stimuli correlated with some properties of the item or event to be remembered. Sometimes these properties are specified by instructions, as when we're asked which of two items came earlier in a list or whether we've met someone before or where we were last year on the night of November 5. In these cases, the "parts we hold" are the items or a face or a date. At other times, circumstances define these properties, as when we can't remember the article we were supposed to buy at the store or where we put our keys. Here, the "parts" are our presence in the store or the activity that requires the keys.

An item that's stored is said to be *available* in memory. The item, however, may or may not be remembered. When you can remember, we say it is *accessible*; when you can't, we say it's *inaccessible* (Tulving & Pearlstone, 1966). The trouble is that if you remember an item we know that it was both available and accessible, but if you can't we don't know whether it was unavailable or available but inaccessible (cf. Watkins, 1990, 1996).

The accessibility of an item depends on the stimuli or cues present at the moment of recall. Their influence on remembering is called *cue-dependency*. For example, visualizing a place used in the method of loci reinstates one condition that existed when you encoded the item. Similarly, reciting the alphabet may help you to remember a forgotten name, because producing the person's initial may reinstate one part of saying the name. When we use such techniques, we generate our own retrieval cues. The most critical feature of retrieval is producing conditions similar to those during encoding and storage.

In a special case of cue-dependency called *state-dependent learning*, recall is affected by the similarity of the learner's condition at retrieval to that at storage. Conditions studied include drugs and physiological states (e.g., electroconvulsive shock) as well as experimental settings (e.g., Bower, 1981; Overton, 1964). For example, the sober learner may be less likely to remember something learned while drunk than the learner who has become drunk again.

Cue-dependencies and state-dependencies may have important practical implications. But we might not want to advise a student who's preparing for an exam to study in the room in which the exam will be held. For one thing, what's been learned will be useless after the course is over if the student can't remember it anywhere else. More important, what's learned will be best remembered if many of the conditions for remembering arise out of the subject matter itself (cf. Marholin & Steinman, 1977). A systematic subject matter is one in which each component prompts and is prompted by other components. To outline the subject matter, therefore, is to create an encoding that determines not only how well the various parts are remembered but also how the recall of one part may aid in the recall of another.

METAMEMORY

Remembering is behavior and remembering can be learned. Not only do we learn patterns of rehearsal and mnemonic techniques, but we also learn to judge the properties of our own remembering (e.g., Flavell, Friedrichs, & Hoyt, 1970; Nelson, 1992). What we do when we remember depends at least partly on the past consequences of our remembering. We can therefore define *metamemory* as the differentiation and discrimination of our own remembering.

Running or Working Memory

Continually updating what's remembered by dropping some items and adding others is sometimes called *running memory* or *working memory*. The short-order breakfast cook working on three orders of waffles, a stack of pancakes, two eggs over easy and two scrambled can forget some items as each dish is handed out to be served and must add some items as each new order is called in. Earlier breakfast orders over the course of the morning are potentially a major source of proactive interference. Nevertheless, this task is often performed with considerable skill. It can't be done well simply through undifferentiated remembering. The updating of what can be forgotten and what must still be remembered is crucial to doing the job successfully (e.g., Bjork, 1978). Other examples are keeping track of cards played in a card game or following the changing statistics during an athletic event (most appropriately, running memory of races run, laps completed, relative positions of competitors and other information at a track meet). In such circumstances, what can be forgotten may be as important as what must be remembered.

Discriminated Remembering

Not only can learners learn to remember and forget differentially, as when they're more likely to remember unfinished than finished tasks, as in the Zeigarnik effect (Zeigarnik, 1927); they can also discriminate among properties of their own remembering. The "tip-of-the-tongue" phenomenon (Brown & McNeill, 1966; cf. Chapter 22) is an example of a discrimination based upon the likelihood of remembering. Partial reports of the word to be remembered can help us to judge the accuracy of this discrimination. When a word is "on the tip of my tongue," I can often report its initial letter or the number of syllables or some other property, and I'll probably recognize it when I see it even if I can't recall it right now.

Both at storage and retrieval, learners can estimate the likelihood of remembering; Having remembered something, we often discriminate among its sources. In recalling some incident, for example, I may be able to report that some facts are based on my own experience whereas I deduced others (e.g., Johnson & Raye, 1981). Another instance of discriminating my remembering is when I describe my confidence about something I've recalled.

We can usually discriminate between never having learned something and having learned and then forgotten it. But how can the metaphor of memory search deal with how rapidly and accurately we can discriminate between knowing or not knowing something (e.g., Kolers & Palef, 1976)? If we must examine every item in metaphorical storage to be sure we don't know something, why is that search usually slower when an appropriate item exists to be found than when no such item exists?

Our capacity to make such judgments changes over time, as shown in metamemory through the life span (cf. Lachman, Lachman, & Thronesbery, 1979) and in intellectual self-management in old age (Skinner, 1983). Such judgments, closely related to the autoclitic processes of Chapter 22, aren't memory but probably are important components of it. For example, the student who can't distinguish between having learned something well and having learned it inadequately may have trouble studying effectively and raising appropriate questions (e.g., Bisanz, Vesonder, & Voss, 1978; Miyake & Norman, 1979). How can that student judge which parts of a text to spend more time on or which points need to be worked through more thoroughly?

Remembering is a higher-order class, and many aspects of it are presumably shaped by natural contingencies. But the various dimensions of metamemory should remind us of the possibilities for differentiating it further. In other words, to propose a metaphor perhaps worth acting upon, **remembering is a skill that can be sharpened.**

Chapter 27
The Structure of Remembering

> Memory *has its source in an Indo-European root* smer- *or* mer-, *to remember, through which it is related to* remember *and* mourn. *The root does not seem to be linked to the Indo-European* men-, *to think, which is the source of* mnemonic, amnesia, memento, reminiscence, automatic, *and, perhaps most interesting,* mind *and* mental. Forget *has a source in the Indo-European root* ghend-, *to seize or take. Through the Middle English* gessen, *to try to get, it is related to* guess, *and through the Latin* prehendere, *to hold before, it is related to* apprehend *and* comprehend.

Many theories of remembering have been concerned with how it works, but a dominant theme in memory experiments is the study of what we remember. Analyses of what is remembered, like analyses of the effects of reinforcement, tell us about behavior classes. Our example of the pigeon's remembering at the start of Chapter 26 required continued behavior that bridged the temporal gap between the stimulus and the later response, but the human mnemonic examples did not include uninterrupted intervening behavior. Clearly the cases differ. We'd similarly distinguish between remembering an appointment after seeing it on a calendar and remembering without checking the calendar. These are different varieties of remembering.

One criterion for organizing kinds of remembering is the time period over which something is remembered; another is what is remembered. We'll sample both classifications. First we'll treat the duration of remembering by examining three phenomena: (1) the relatively brief persistence of stimulus effects; (2) the maintenance of responding occasioned by a stimulus, as in rehearsal; and (3) remembering after some time elapses without rehearsal. These categories have been respectively called (1) *iconic memory*, (2) *short-term memory*, and (3) *long-term memory*. Then we'll treat categories of remembering defined by what's remembered, in *autobiographical* or *episodic* memory and *semantic* memory; autobiographical memory involves remembering events in the learner's past, and semantic memory involves remembering properties of the learner's language, such as word meanings. The structure of what's remembered will be relevant to the distinctions among these types of remembering, but as in Chapter 26 we'll discover that structure and function often interact.

KEY TERMS: Iconic Memory, Short-Term Memory, Long-Term Memory; Interference, Forgetting; Procedural, Declarative, Autobiographical and Semantic Memory.

ICONIC MEMORY: THE PERSISTING EFFECTS OF STIMULI

The effects of a stimulus may continue even after the stimulus is gone. Persisting aftereffects of visual stimuli are called *icons* (Long, 1980), and the topic of *iconic memory* is concerned with their time course. An afterimage is one persisting effect of a visual stimulus. How do we measure that time course? If an observer reports some letters that have been briefly presented, how do we tell the difference between a report based on the sensory aftereffect and one based on the observer's rehearsal of the letters?

We can solve the problem by showing more items in the stimulus display than the observer can remember in a single trial and asking the observer for some sample of the items at different times after the display. We are limited in the number of items we can rehearse or remember after a single brief presentation. This limit, typically within the range of seven plus or minus two items (Miller, 1956), is called the *span of immediate memory*. Here it serves as a tool for studying iconic memory. The range of the span of immediate memory is to some extent independent of the nature of the items. For example, you can remember roughly five to nine words almost as easily as five to nine letters, even though the words themselves include many more than nine letters (cf. chunking in Chapter 26).

An exam tests only part of the material that a student is supposed to have learned during a course and assumes that the score represents what proportion of the entire course has actually been learned; similarly a sampling of what the observer reports at different times after a display assumes that the observer's responses represent the proportion of reportable items. The time course of iconic memory is measured over fractions of a second, but during this time the observer can in effect still read items in the display even though the display is no longer there.

Figure 27-1 shows data from one experiment based on this rationale (Sperling, 1960). A 4-by-3 matrix of letters and numbers was presented to observers for 50 ms (0.05 s). An example is the following:

7	I	V	F
X	L	5	3
B	4	W	7

Observers could not report this many items, so a tone of high or medium or low frequency served as an instruction to report the characters in just the top or middle or bottom line of the matrix. The tone sounded either before the display, at the moment it ended or at some time after it was gone (x-axis in Figure 27-1). The mean number of characters correctly reported from the single line was multiplied by three and taken as an estimate of the total reportable characters in the matrix (y-axis in Figure 27-1).

About 10 characters from the matrix were reportable when the instruction tone preceded the display or came just as it ended (delays of -0.10 and 0 s); as the instruction tone followed the display with delays of 0.15 s or longer, the reportable characters decreased with increasing delay. When it was delayed by 1 s, the reportable characters were about equal in number to the span of immediate memory in this task (this span of about 4 to 5 items, at the extreme low end of the usual 5- to 9-item range, probably depended on the relative complexity of the observer's task, which included

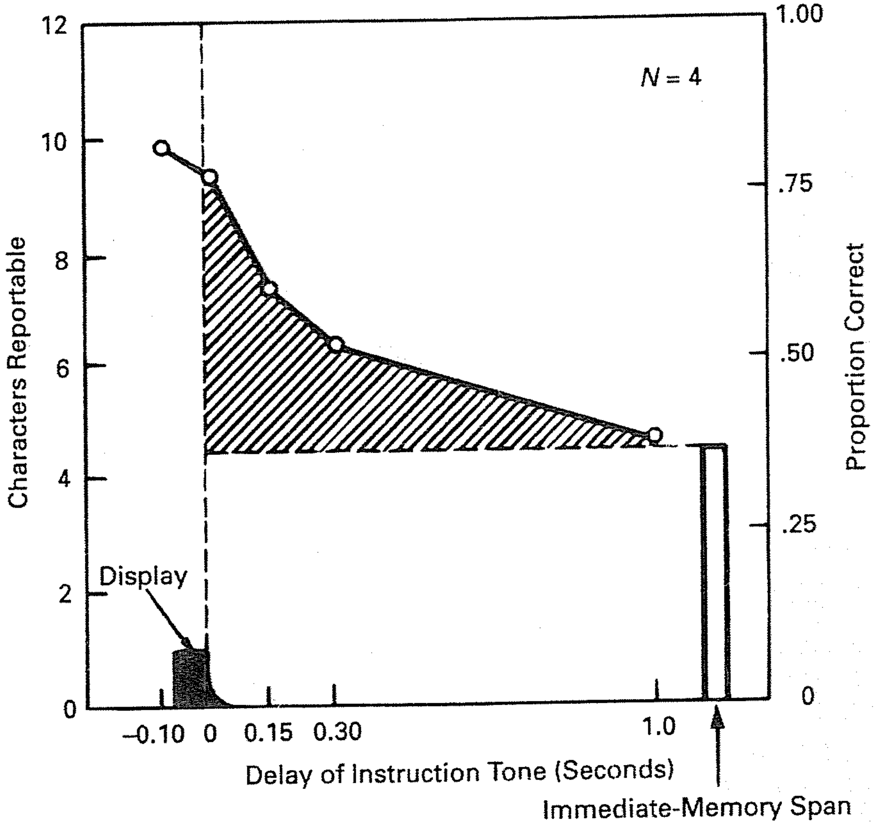

Figure 27-1 Characters reportable from a 3-by-4 matrix of letters and numbers when the signal for reporting occurs at various times relative to the display. The signal was a tone of high, medium or low frequency that instructed the observer to report the characters in the top, middle or bottom row of the matrix. With increasing delays of the instruction tone, accuracy approached the immediate-memory span. Data are means across four observers. (Adapted from Sperling, 1960, Figure 7)

mixed letters and numbers and the matrix format). The shaded area between the data points and the immediate-memory span represents the persistence of the sensory effects of the display,, in the sense that this is the extent to which the observer can still read some of the matrix even though it is no longer there.

The sensory effects of a stimulus continue briefly after the stimulus has been removed. Similar effects have been demonstrated with auditory presentations of verbal stimuli, as when different sequences of spoken letters were simultaneously presented in three auditory channels (left ear, right ear or both ears; Darwin, Turvey, & Crowder, 1972). Effects of auditory stimuli appear to decrease more slowly, perhaps over several seconds, than those of visual stimuli. The auditory case is called *echoic memory*, but should not be confused with echoic verbal behavior as in Chapter 20. Echoic memory involves the persisting effects of auditory stimuli, without regard to the nature of the listener's response to the stimuli, whereas echoic verbal behavior is defined by correspondences between auditory stimuli and the listener's vocal responses.

Obviously the continuing sensory effects of stimuli for a brief time after they end has only a little to do with remembering over longer periods.

Visual stimuli must be seen and auditory stimuli must be heard to be remembered, but remembering over minutes or days cannot be attributed to their persisting but brief sensory effects.

SHORT-TERM MEMORY: THE ROLE OF REHEARSAL

The recall of a human learner immediately after presentation of a verbal sequence, usually limited to about five to nine items, is the span of immediate memory (Jacobs, 1887; Miller, 1956). For example, you can probably correctly repeat the sequence 706294 after a single hearing, but you're not likely to manage the longer 549628367999102.

The limit on the span of immediate memory provided the historical basis for studies of what has come to be called *short-term memory* (STM or, in some usages, *primary memory* (cf. Daniels, 1895; Smith, 1895; Waugh & Norman, 1965). A major question is whether the response is occasioned directly by prior stimuli or is instead occasioned indirectly by intervening behavior such as rehearsal. The question can be addressed with a task that prevents the learner from rehearsing items between their presentation and the opportunity for recall. The tasks used to measure the span of immediate memory create such conditions incidentally, because a learner cannot simultaneously rehearse the early items and listen to the later ones if a long sequence of items is presented. Later items therefore prevent the behavior on which the recall of earlier items may depend.

Now consider the short-term memory experiment of Peterson and Peterson (1959). Vocal stimuli consisted of three consonants and then a three-digit number. The learner was instructed to count backward by threes from the number and then, when a signal light flashed, to recall the consonants. A trial might start with the spoken items, CHJ 503; the learner then counted backward: 503, 500, 497 and so on until the signal. Recall accuracy decreased with delay until, at delays of 15 or 18 s, the proportion correct was less than 10%. In two other conditions, learners were allowed vocal or silent rehearsal before they were given the number from which to count backward. Figure 27-2 shows data from the main procedure (open circles) and with 3 s of silent rehearsal (filled triangles) or vocal rehearsal (filled circles) before the backward-counting task. The more overtly the learners could rehearse, the better their recall. Recall is determined more by our behavior with respect to past stimuli than by the stimuli themselves.

These and related findings led to proposals for two distinct types of memory: short-term memory and long-term memory (e.g., Shiffrin & Atkinson, 1969). Items were said to be rapidly lost from short-term memory unless maintained there through rehearsal; by some means, perhaps through rehearsal itself, items in short-term memory were transferred to a more permanent long-term memory. Some accounts distinguish between two types of rehearsal (e.g., Craik & Lockhart, 1972; Rundus, 1977, 1980). In *maintenance* rehearsal, an item is simply repeated (as in reciting something to be memorized); in *coding* rehearsal, the item is transformed or elaborated in some way (as in mnemonic encoding). Note that both types of rehearsal depend on encoding; even naming a stimulus is a type of encoding.

Only a few items at a time can be held in short-term memory, but the capacity of long-term memory is virtually unlimited. More detailed analyses of data from short-term memory procedures complicated this view, however. Without rehearsal (Figure 27-2, open circles), the data at first seem to represent the gradual fading or decay of items to be remembered. But several lines of evidence suggest this isn't so.

Consider the effects of varying the number of letters to be remembered, as in Figure 27-3 (Melton, 1963). A stimulus item containing one to five consonants and then a three-digit number were presented visually; the learner read the consonants aloud and then began counting backward by threes or fours until a visual signal set the occasion for recall. When the item con-

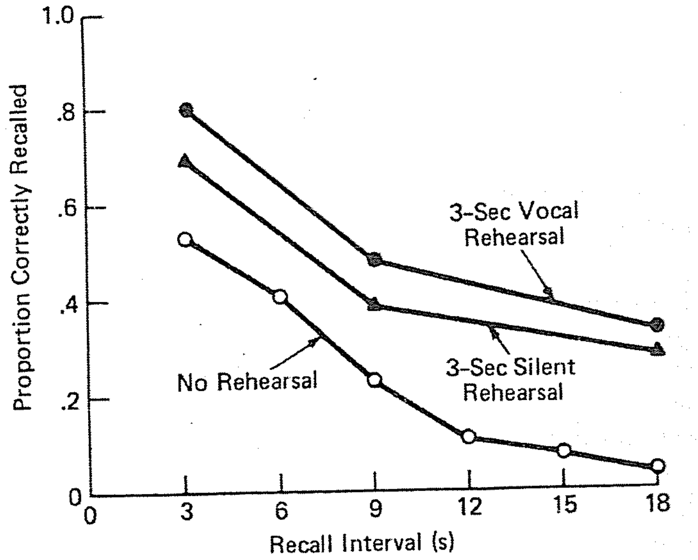

Figure 27-2 Proportion of three-consonant items correctly recalled as a function of delay between presentation and opportunity for recall. One condition (open circles) prevented rehearsal by having the learner count backwards by threes from a number immediately following each item. In two other conditions, the learner was allowed 3 s of silent rehearsal (filled triangles) or vocal rehearsal (filled circles) before starting to count backwards. (Adapted from Peterson & Peterson, 1959, Figure 3 and Table 1)

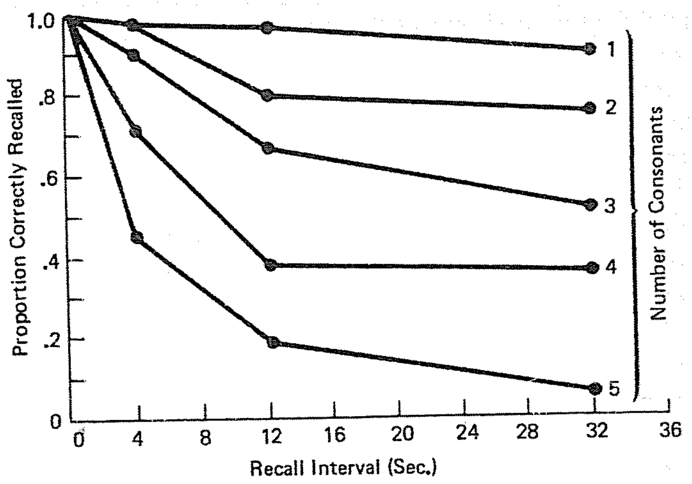

Figure 27-3 Proportion of consonants correctly recalled as a function of delay between presentation and recall, with number of consonants as a parameter. The consonants and a three-digit number were presented visually; the learner read the consonants aloud and then began counting backwards by threes or fours until a visual signal for recall. (From Melton, 1963, Figure 2)

sisted of just a single consonant, recall accuracy remained high even over 32 s. The proportion of items recalled decreased more rapidly over time as the number of consonants increased. But now which set of data represents the time course of short-term memory? If recall of a single consonant is nearly perfect, then adding new consonants may simply interfere to some extent with recall of the old.

The two extreme possibilities in these procedures are perfect recall and no recall at all. Perhaps the backward-counting task prevented rehearsal only partially, and intermediate outcomes occurred because learners had some opportunities for rehearsal even while engaging in that task. If more rehearsal is necessary for recall of more consonants, delayed recall accuracy should decrease with the number of consonants. Thus, the form of the short-term memory function may depend mainly on the learner's opportunity to respond to the stimulus items (cf. Crowder, 1976, p. 196).

Once we recognize that each item can affect the recall of others, we can ask whether the numbers in backward counting prevent rehearsal or interfere with recall in some other way. If each number produced by a learner during backward counting is an item, does recall vary with time or with number of items? Surveys of memory experiments in which delays between items and recall included various numbers and rates of intervening items (Waugh & Norman, 1965) indicated that the number of intervening items during the delay was more important than the delay duration. For example, learners were given numbers at rates of 1 or 4 per s and then were asked to name the number at a particular place in the sequence. Recall probability after 4 items in 1 s was about equal to that after 4 items in 4 s despite the fourfold difference in delay; and recall probability was greater after a delay with 2 intervening items in 2 s than after one with 8 items in 2 s. In both cases, behavior mattered more than mere passage of time. If any rehearsal occurs while items are presented, items presented slowly should permit more rehearsal than items presented rapidly. Given that short-term memory procedures are usually designed to prevent rather than encourage rehearsal, it is curious that rehearsal has played such a substantial role in their interpretation.

Still another problem is that recall probability changes over the first few trials of a short-term memory procedure (e.g., Keppel & Underwood, 1962). For example, recall of a three-consonant item on the first trial of a short-term memory session is nearly perfect at a delay of either 3 or 18 s; recall probability decreases over the next three to six trials, but the decrease is much greater with a delay of 18 s than with one of 3 s. In other words, the relation between recall probability and delay, as in Figures 27-2 and 27-3, is absent at the beginning of sessions and builds up over several trials. This is an example of proactive inhibition: Learning consonants on the first trial interferes with the recall of other consonants on later trials.

LONG-TERM MEMORY: INTERFERENCE AND FORGETTING

In studies of short-term memory, the time from item presentation to recall is usually a matter of seconds, whereas in studies of long-term memory (sometimes abbreviated LTM) it may be minutes, hours, days or even years. It has been argued on the one hand that short-term and long-term memory are separate types of remembering with different properties, and on the other that short-term and long-term memory are merely extremes on a single continuum determined by variables that affect remembering (Melton, 1963; Tulving & Madigan, 1970). The account here will mainly elaborate on a simple procedural consequence of the existence of a span of immediate memory.

When the number of items in a list exceeds the span of immediate memory, a single presentation isn't enough for the learner's recall, even if recall follows immediately after the list ends. The only way to study the remembering of such lists is to arrange repeated presentations of some or all of the items to be learned. In other words, the number of items to be remembered determines whether repeated presentations are needed; that's why they're more important in distinguishing between short-term and long-term memory than the period of time over which remembering is measured. Long-term memory includes all cases in which the items to be remembered exceed the immediate memory span. Such remembering doesn't allow uninterrupted rehearsal between presentation and recall; such remembering requires repeated presentations of the to-be-remembered items; so, inevitably, such remembering is likely to be studied over relatively longer time periods than those common to research on short-term memory.

Data from the classic study of long-term memory by Ebbinghaus (1885/1964) are shown in Figure 27-4. The remembering of nonsense-syllable lists was assessed by the method of savings (the difference between original learning and relearning) at 20 min, 1 and 8.8 hr, and 1, 2, 6 and 31 days after original learning. The largest decrease in the proportion remembered (in other words, the most forgetting) occurred shortly after original learning, but substantial savings remained even after 31 days. Ebbinghaus entertained the alternative possibilities that memories deteriorated over time or that they remained intact but were gradually overlaid by or hidden beneath other memories.

These two views were the precursors of many theories of forgetting. Those based upon the passive decay of memories, sometimes called *trace* theories, usually assumed correspondences between remembering and hypothetical processes in the nervous system. Those based upon competition among different memories, sometimes called *interference* theories, were more likely to rely on stimulus and response variables. A third kind of theory, memory *consolidation*, argued that memory is impermanent immediately after learning and that time is required for it to become fixed or consolidated; during that time, various events can disrupt it (e.g., trauma). Consolidation theories were influenced by the phenomenon of *reminiscence* (Ballard, 1913; Kamin, 1957). Reminiscence, most likely to occur with incompletely learned materials, is an increase in recall probability as time passes since the end of learning; it only occurs under certain conditions. For example, it is more likely with pictures than with verbal material (Erdelyi & Kleinbard, 1978).

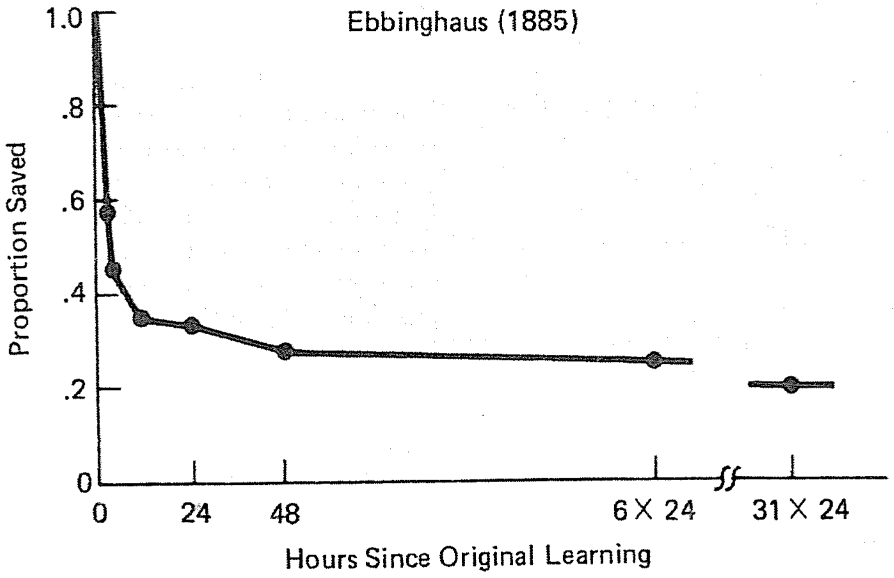

Figure 27-4 Ebbinghaus's forgetting curve. Ebbinghaus learned and relearned 13-syllable lists. With different times between original learning and relearning, he measured how much was saved from the original learning. Forgetting was substantial even soon after original learning (the first point is at 20 min), but after 31 days savings were still about 20%. (Adapted from Ebbinghaus, 1885/1964, Chapter VII)

The notion of memories as traces that fade or decay over time was implicit in many early generalizations, such as Jost's law: If two associations are of equal strength but different ages, the older will lose strength or be forgotten more slowly than the newer (Jost, 1897). The problem with such generalizations is that what is remembered varies with how we measure it: Forgetting as measured by savings on relearning a list differs from that measured by recall or recognition of the items; an easily learned list isn't necessarily better remembered than a list learned with difficulty; two lists that differ in difficulty or in the time to meet learning criteria may be forgotten at approximately equal rates if mastery of the items on the two lists is equated (Underwood, 1964).

In one form or another, accounts in terms of interference have come to dominate analyses of long-term memory. For example, interference theories were consistent with effects of sleep. Although we've already seen that there's no convincing evidence for learning during sleep (Chapter 26), less forgetting of what's learned while awake occurs during sleep than during equal waking periods (Jenkins & Dallenbach, 1924). One interpretation is that events during waking are more likely than events during sleep to interfere with what's been learned.

Early treatments of forgetting in terms of interference assumed that events between learning and recall were the major source of interference. They were assumed to act retroactively, so that recent events affected what was learned earlier.

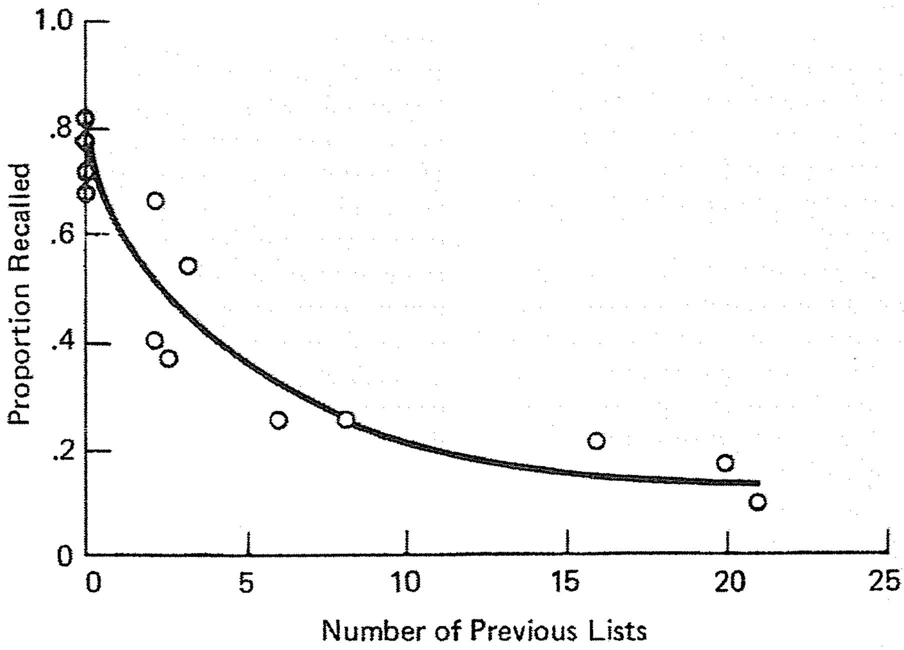

Figure 27-5 Recall of items on a list as a function of the number of lists previously learned. Each circle represents data from a different study. Forgetting was greater as the number of previous lists increased; in other words, forgetting depended on proactive interference from the earlier lists. (Adapted from Underwood, 1957, Figure 3)

But that assumption was wrong. The critical finding for analyses in terms of interference was that interference worked the other way around: Earlier learning had substantial proactive effects, influencing the forgetting of more recently learned material (Underwood, 1957; Underwood & Postman, 1960). Figure 27-5 summarizes data from several experiments and shows how the forgetting of a list varies as a function of the number of lists learned earlier (Underwood, 1957).

The demonstration of proactive interference was surprising and significant. Most data on human memory over more than half a century had been obtained from practiced learners who had served in experiments involving the learning of many lists. Ebbinghaus himself typified that circumstance, and even learners who didn't participate in experiments with many lists under many conditions were usually given practice lists before the experiment began. From those studies, the estimate of forgetting over 24 hours had been about 75%; Underwood's analysis showed that most of this forgetting is produced by proactive interference from the earlier lists, and that without such interference forgetting is only about 25% (zero lists, Figure 27-5). In other words, the major cause of forgetting is that **older learning interferes with the remembering of what has been learned more recently**.

If so much forgetting was produced by proactive interference from what was learned in the laboratory, how much of the forgetting that was left was produced by proactive interference from sources outside the laboratory (Underwood & Postman, 1960)? We should not be surprised that older people often remember events from long ago while forgetting recent details like what they had for breakfast. The older you get the more there is to interfere with your remembering of the most recent events. Even someone in an experiment on verbal learning for the first time enters with an extensive verbal history. If learning a list in the laboratory can have some effect, then learning outside it should act the same way. Might there be little forgetting or even none at all if it were not for proactive interference ?

THE ROLES OF HEARSAL AND REHEARSAL IN REMEMBERING

The most important difference between experiments on iconic memory and those on short-term memory is whether the learner is given an opportunity to respond to the to-be-remembered stimulus. The most important difference between experiments on short-term memory and those on long-term memory is whether the learner is given an opportunity to respond more than once to the to-be-remembered stimulus. A term already exists for the learner's repetitions of a stimulus; we call it *rehearsal*. But the very first response of those repetitions is special. The difference between iconic memory and short-term memory depends on it, and as the initial response it doesn't really count as a repetition. The term *rehearsal* is well established, but it doesn't sit well as applied to the learner's initial response. Let us therefore drop the prefix and apply the term *hearsal* to that first response.

This also lets us add a fourth memory category to our list, because hearsal and rehearsal are occasioned by the initial encounter with something to be remembered (i.e., during storage), but rehearsal can also occur during other stages of an episode of remembering, as when retention is probed by asking a learner about some remembered item. In other words, retention and recall involve rehearsal at a time when hearsal is already long past. We can now review these main memory categories in terms of the distinction between hearsal and rehearsal.

Table 27-1 summarizes a few properties of our major categories of remembering. We considered the three main varieties (iconic, short-term and long-term) in terms of what is remembered, but we also appealed to functional differences among these classes, in the sense that what is remembered depends on what the learner does at various stages of an episode of remembering.

In experiments on iconic remembering, the learner sees a very brief presentation of an array of items, too many to read off and recite at one time (i.e., the number of items in the set exceeds the Span of Immediate Memory). Only after this array has vanished does the learner receive the signal that says which items to report. The learner could respond to a few items by reciting them, but that will mean hardly ever doing well when the signal calls for other items. Thus, the learner will usually drop that kind of responding with increased practice. Once that happens, the best the learner can do is to respond to the persistent effects of the stimulus, as in seeing an afterimage. In other words, this procedure denies the learner an opportunity even for a hearsal, and rehearsal is hardly feasible if the learner hasn't respond to the stimulus at all.

Table 27-1 Some Properties of Some Classes of Remembering

Classes of remembering	*Content criteria*	*Response criteria*	*How long they last*
Iconic Memory	Stimulus arrays presented so briefly that the learner cannot respond to all items	Neither Hearsal nor Rehearsal	Fractions of a second
Short-Term Memory	Item sequences shorter than the immediate-memory span and therefore manageable on a single presentation	Hearsal only	Seconds
Long-Term Memory	Item sequences longer than the immediate-memory span and therefore requiring more than one presentation	Hearsal plus Rehearsal	Seconds, hours, years
Retention and Retrieval	Items recalled some time after they were first encountered	Rehearsal only	(not applicable)

Experiments on short-term remembering briefly present verbal sequences in which the number of items is within the span of immediate memory and then engage the learner in a task that makes it difficult to rehearse the stimuli (e.g., counting backward by threes). But that task does not prevent the learner from responding to the initial presentation of the sequence. In other words, this procedure allows hearsal but prevents rehearsal.

Experiments on long-term remembering inevitably allow both hearsal and rehearsal, because they involve sequences in which the number of items exceeds the span of immediate memory. The learner cannot master such sequences on a single presentation, so repeated presentations are required. Under these circumstances, learners will probably respond to some or all items as they are presented and also rehearse items between presentations. In those rehearsals, the learners present the items to themselves again, and undoubtedly these re-presentations contribute to the learners' long-term remembering. In any case, the procedures obviously allow both initial hearsal and subsequent rehearsal.

In studies of retention and retrieval, the items to be remembered are no longer present and so do not allow hearsal. The hearsal occurred earlier when the items were first learned. The repetition does constitute rehearsal, however. Furthermore, in the absence of the stimulus the repetitions may be altered by other variables, as when descriptions of remembered events become elaborated in different ways depending on the reactions of an audience. In fact, studies of retention typically become studies of retrieval, because it is only possible to determine what a learner has retained by requiring the learner to recall it.

KINDS OF REMEMBERING

Many kinds of events and relations can be remembered. For example, the pigeon that pecks a key today because pecks yesterday occasionally produced food can be said to be remembering the past contingency between pecks and food. We've already asked whether remembering of this sort is related in any simple way to remembering in human verbal learning. The vocabulary of memory may be occasioned by almost any situation in which current behavior is influenced by past events. We remember words, particular incidents, contingencies, stimuli, definitions, syntax and our own behavior, among other things. Structure exists in what is remembered in specific instances. It also exists at the higher-order level of relations among different kinds of remembering.

Table 27-2 summarizes a few classes of remembering defined in terms of what is remembered. With regard to many of these distinctions, functional considerations cannot be separated from structural ones. In comparing autobiographical and semantic memory, for example, your recall of a particular event, whether a personal incident or your participation in an experiment on verbal learning, typically involves just one exposure, whereas experience with your language involves countless verbal contingencies operating over a lifetime. Remembering incidents and remembering words obviously differ, but how could we ever equate them in terms of number of exposures or in terms of rehearsal? Some things are more easily remembered than others, such as major life events as opposed to routine daily activities, but we must remain alert to the complexity of teasing apart the relative contributions of such content variables from those of variables like repetition that can be more explicitly measured.

Autobiographical and Semantic Memory

When we speak of everyday remembering, we're usually concerned with particular incidents that occurred at certain times and places. This kind of remembering is called *autobiographical* or *episodic* memory, in that it involves the recall of episodes from our own lives. Even this class of remembering may include subclasses, such as differences between ordinary day-to-day memory and memory of emotionally charged events. One special case of autobiographical memory, sometimes called *flashbulb* memory, involves remembering the details of where we were when news came about a significant event such as a major disaster (Neisser, 1992). Despite our intuitions about it, even this type of remembering can be modified substantially by the conditions under which we retell the story.

Autobiographical memory is often contrasted with *semantic* memory (or the related lexical memory), or remembering our language: word usages, idioms and the functional aspects of grammar (Tulving, 1972). You can't understand a sentence if you can't remember what the words mean. But there is a difference between remembering what a word means and remembering where you saw it last. Only the former is semantic memory. In serial learning, for example, the learner recalls the items on a particular list in a particular setting. That list constituted an episode in the learner's life and therefore qualifies as autobiographical memory.

Some classes of remembering are categorized on the basis of selective loss. For example, some cases of amnesia involve the forgetting of personal history without loss of language. But the details of the two kinds of learning are rarely comparable in terms of frequency and recency, so in many other respects autobiographical and semantic memory have similar functional properties (e.g., Anderson & Ross, 1980; McCloskey & Santee, 1981).

Table 27-2 Some Classes of Remembering*

Classes of Remembering	*Content*
Autobiographical or Episodic	Life events (including the learning of particular lists)
Semantic	Vocabulary (especially of one's native language)
Procedural	Actions (especially motor skills; e.g., bicycling, driving, typing)
Declarative	Names, definitions, and other verbal material
Cue-Dependent or State-dependent	Events the recall of which varies with the context in which the recall occurs

*These are not necessarily mutually exclusive, and this is a partial list (e.g., it does not include working memory and other metamemory derivatives).

Procedural and Declarative Memory

Given the variety of kinds of remembering, we might expect their properties to depend on what is remembered (cf. Roediger & Craik, 1989). Remembering how things are done is called *procedural* memory. Motor memory is an example (e.g., Baddeley, 1976). Two classes of motor memory are discrete skills, such as typing or changing gear in a manual-shift car, and continuous skills, such as tracking a continuously moving target or steering a car. The distinction is important mainly because discrete skills are more likely to be forgotten than continuous skills (one doesn't forget how to swim or to ride a bicycle). The two skills seem otherwise closely related. If a memory difference can be observed even in this comparison, we must be alert for others.

Procedural memory is usually contrasted with *declarative* memory, which is verbal remembering or the remembering of facts (e.g., Tulving, 1985). This is similar to the distinction between knowing how and knowing that (cf. Chapter 1). Most memory examples considered here and in Chapter 26 have involved declarative memory.

Remembering and Cue-Dependency

In cue-dependent remembering, the likelihood of recall becomes greater as the conditions of recall more and more closely resemble the conditions of learning. The phenomenon is evident not only with respect to the environmental contexts within which learning takes place but also with respect to conditions of the organism, such as mood and emotion (when it is sometimes called state-dependency rather than cue-dependency). It is reasonable that response probabilities at the time of recall should vary depending on whether stimuli present during learning are also present during recall.

Remembering may also be discriminated responding, however, not only in the sense discussed in connection with metamemory in Chapter 26, but also in the sense that it may become more likely in some conditions than in others as a result of contingencies of reinforcement. For example, consider a pigeon responding in a delayed matching-to-sample task, in which a peck on a sample turns it off and some time later comparisons appear. In green versus red matching, only pecks on the green key of a green-red pair are reinforced given a green sample and only those on the red key of the pair are reinforced given a red sample. The pigeon's accuracy typically decreases as the time between sample and comparisons increases; we might say that the sample stimulus is remembered less and less well as time passes.

But suppose we compare the relation between accuracy and delay in two groups of pigeons, one of which was trained with no delay between sample and comparison and the other with a 4-second delay between sample and comparison (White, 2001, 2002; White & Wixted, 1999). In this case, we see delay-specific remembering. The birds with 0-second delay training become less accurate as delay increases, whereas the birds with 4-second delay training are most accurate with 4-second delays and show lower accuracies at 0-second and 2-second delays as well as at delays longer than 4 seconds. We probably do not want to say that the birds in the latter group remembered less after 0 or 2 seconds than they did at 4 seconds, but it is reasonable to say that they showed discriminated remembering, in that their remembering varied as a function of the temporal dimensions of the training stimuli.

The functional significance of delay-specific remembering in natural environments is illustrated by the storage of food in the wild by squirrels and birds and other animals. If some foods are more perishable than others, the times when it is most important to remember where they were buried will vary from one food to another. This also implies that as time passes contingencies may become relatively more potent in altering what is remembered.

A NON-SUMMARY

It would be redundant to summarize the treatment of remembering in Chapters 26 and 27. The essential vocabulary of memory has been included in their section headings. These headings can be listed and remembered in a variety of ways: by rote rehearsal, by mnemonic techniques, by reviewing the relations among particular experiments and concepts, and so on. The reader who learns this material merely as a sequence of words won't remember it as well as the one who has responded to it in several other ways. The reader who has explored all of the resources—the headings, the figures, the glossary, the index, and yes, even the etymologies—will remember more than the reader who has more narrowly used this book.

We haven't explained memory for words, texts, specific events, structures or our own behavior, but we've seen that what is remembered defines significant units of behavior. Just as the relation between contingencies and their behavioral products defines functional response classes, the relation between what has happened and what is remembered defines the structure of memory.

ADDENDUM 27A: PROACTION AND RETROACTION

We've seen that learning list 1 can affect learning list 2. This effect is called *proactive*, because the direction is from an earlier to a later task. But list-2 learning can also affect later list-1 performance. For example, suppose a learner who has mastered an A-B list is switched to learning an A-C list, in that all of the stimulus items remain the same but all of the response items are changed. Once the A-C list has been mastered, the learner who had given B responses to A stimulus items may no longer be able to do so. The A stimulus items now occasion C responses. This effect is called *retroactive*, because the direction of effect is from a later task to performance on an earlier one. Most research on such effects has concentrated on cases in which learning one list interferes with rather than enhances performance on other lists (Müller & Pilzecker, 1900); such effects are called proactive and retroactive *interference*.

Whether transfer from one list to another is positive or negative depends on relations among the items in the two lists, such as similarities between corresponding first-list and second-list items (Osgood, 1949). Positive transfer is maximal when stimulus and response items in the two lists are identical (A-B to A-B transfer); negative transfer is maximal when stimulus items are identical but second-list response items are incompatible with first-list response items (A-B to A-C transfer). Transfer approaches zero as second-list stimulus items become more different from first-list stimulus items (A-B to C-B or A-B to C-D transfer).

The significance of this summary of transfer, however, depends largely on the definition of similarity. Consider the learner who has learned *LOUD* as the response to a stimulus item. New lists might substitute new response items related to *LOUD* by spelling (*CLOUD*), as synonyms (*NOISY*) or as antonyms (*SOFT*). How can we judge the similarity between list-1 and list-2 response items?

In fact, transfer will probably be positive in each case. Having learned *LOUD* in response to a stimulus item in the first list, the learner would readily learn *CLOUD* in response to the same item in a new list. But that would also be true for transfer from *LOUD* to *NOISY* or *LOUD* to *SOFT*. Transfer might even be greater with antonyms than with synonyms or spelling relations. Such features of transfer data are probably best regarded as part of the definition of similarity. Similarity isn't a stimulus property; rather, it is derived from behavior with respect to stimuli. It follows that doing the experiment remains the most appropriate way to judge the amount and direction of transfer from one task to another.

ADDENDUM 27B: OTHER KINDS OF MEMORY

Procedural memory doesn't necessarily involve words. Some types of remembering can't be assessed simply by asking what someone remembers. We may not know how much we know. This is as true for the remembering of verbal material as for the remembering of ways of doing things, and it is the basis for distinguishing between *implicit* and *explicit* memory (Johnson & Hasher, 1987; Roediger, 1990). These two types of remembering are usually well integrated in normal adults, but in individuals with certain types of verbal deficits called *aphasias* they can be dissociated. For example, suppose an observer has the task of reporting tachistoscopically presented words (written words that are presented very briefly); one measure of the threshold for seeing the word is the duration at which the observer can reliably report it. This threshold can be reduced if a priming word, such as the same word or a semantically related word, is shown at some time before the test word. With aphasic adults, thresholds can be reduced by priming words even though these aphasics cannot remember what the priming words were. In other words, verbal stimuli can have effects on later verbal behavior (they are implicitly remembered) even though they cannot be reported later (they are not explicitly remembered). Similar effects can occur in normal remembering (cf. Watkins, 1989).

Some accounts of kinds of remembering have attempted to organize them into a hierarchy of memory systems. Tulving (1985), for example, suggests that procedural memory comes first: It's shared by both verbal and nonverbal organisms. Semantic memory is then built upon procedural memory: Until a language exists, it isn't possible to talk about remembered events. Finally comes autobiographical memory. It is implicit in such a hierarchy that discriminations relevant to one's own remembering emerge last (cf. metamemory in Chapter 26). Others have extended memory classes. For example, analyses of the remembering of words and texts in different languages have demonstrated that the structure of the semantic classes of bilingual speakers extends across language boundaries, as when the learner's remembering of meaning is independent of the language in which items are presented (MacLeod, 1976).

We have not exhausted the list of kinds of remembering. For example, we haven't considered *spatial* memory (remembering paths and things located on them), though such remembering was involved in the treatment of cognitive maps and in the mnemonic method of loci. Some accounts deal with specific sensory systems (e.g., smell: Herz & Engen, 1996). Another memory distinction is that between *retrospective* and *prospective* memory, which is concerned with the difference between remembering on the basis of either past events or what is to be done in the future (e.g., Urcuioli & Zentall, 1986; cf. Wixted, 1989); for example, you might arrive at a hardware store remembering what you wanted to repair or remembering the tool you came there to buy.

Chapter 28
Knowing: Cognitive Processes

The Indo-European root gno- *is the source through Old English of* know *and* knowledge *and through Germanic of* cunning *and* can, *in the sense of being able to. Its Latin derivative is* gnoscere, *to know or to get acquainted with, and this is a root of* cognition, ignorant *and* recognize. *Synonyms of* knowledge *are often related to sensory language. For example, the Indo-European* weid-, *to see, leads to* view *and* vision *through the Latin* videre, *to* idea *and* history *through the Greek* eidos, *and to* guidance *and* wisdom *through various Old English and Germanic forms. The word* see *in itself implies understanding, as in the phrase* I see, *and it has such relatives as the word* insight. Understand *is a curious term, with a literal English derivation of* standing under; *few European translations have components that correspond to either part (e.g.,* comprendre *in French,* capire *in Italian,* entender *in Spanish) except for some related to* stand *(e.g.,* verstehen *in German).*

Cognitive Processes
  Visual Imagery
  Mental Representations
Addendum A: Processing Stages

Early in this text, before embarking on our treatment of the effects of stimuli and contingencies, we distinguished between structural and functional problems in behavior analysis. Structural problems are those concerned with properties of response and stimuli classes; these are the fundamental units of behavior. We've now considered many response classes and stimulus classes, some defined by contingencies (reinforcers, operants and discriminated operants) and others defined by what is occasioned or what is remembered (semantic and syntactic classes).

We noted in Chapter 1 that concern with functional problems is often correlated with behavioral vocabularies whereas concern with structural problems is often correlated with cognitive vocabularies. This text began by emphasizing experiments and vocabulary that grew out of a behavioral tradition. That treatment was extended to verbal behavior, which began the transition to topics that are a primary concern of contemporary cognitive psychology. We examined those topics in the context of verbal learning, psycholinguistics, and memory. This chapter further illustrates research methods and issues in cognitive psychology.

COGNITIVE PROCESSES

We can't see what someone else is thinking or imagining, but like walking or talking, imagining is something we do. Some might argue that events like these shouldn't be called behavior. Yet we haven't limited our treatment of behavior to movements; for example, we don't define operants by response topography. Certainly it is hard to say just what a person does when imagining (cf. Chapter 14 on private events). Nevertheless, it is plausible to assume that imagining shares something with the behavior of looking at things. We can discriminate imagining from seeing; when we fail to do so, we're said to be hallucinating (cf. Skinner, 1953). The difficulty is that such behavior is relatively inaccessible to anyone but the one who engages in it. We may recall the recommendation to convert psychological nouns to verbs: Instead of "cognition and thought" we should say "knowing and thinking" (Kolers & Roediger, 1984; Malcolm, 1971; Woodworth, 1929).

KEY TERMS: Cognitive Processes; Visual Imagery; Processing Stages; Representations versus Copies.

In Chapter 11, we treated attention as the control over responding by some stimulus features but not others. We converted the word from its everyday usage into a technical term, but that didn't remove it from our everyday language. Attending is something we do. How can we identify it and how can we measure its properties?

It isn't good enough just to ask whether someone is attending. The problems of introspection (the reporting of private events) were amply demonstrated in the history of psychology. Introspection wasn't reliable. Different individuals and even the same individual at different times reported events in different ways. Attending presented particular difficulties: Can you attend to your own attending? The difficulties were inevitable. A private event is by definition available only to the person behaving, but a language of private events can be created and maintained only if it is linked to public events accessible to the verbal community (cf. Skinner, 1945, and Chapter 22). Thus, consistencies in reports of private events must depend on consistencies across public and private vocabularies. Imagining an object, for example, must share something with seeing the object.

Attending isn't necessarily movement, but we'd probably want to distinguish between sustained attention, perhaps analogous to maintaining a posture, and switching attention from one thing to another. For visual stimuli, attention may seem superficially like looking toward or even pointing at something. This kind of extension was implicit when we gave a pigeon's key pecks functions similar to those of attending and called them observing responses (Chapter 16). We can create contingencies for discrete responses such as pecks that approximate those for other responses we can't count so easily. But what about looking without seeing, as when we daydream or are "lost in thought" (we then say we weren't paying attention). If we treat attending as mere eye movement, we make looking without seeing equivalent to looking and seeing. That would be a mistake. Attending shares something with what we do when we say we're watching for or looking for something, but we need to distinguish among separate acts of looking somewhere if what we look for differs each time. For example, what we notice when we scan a page for a name isn't the same as when we scan it for a definition.

When behavior doesn't involve movement, we may still be able to record duration, latency or other properties (cf. Posner, 1982). Shifts in attention are particularly well suited for such measurement, as illustrated by Sperling and Reeves (1980). Observers were instructed to look at a visual fixation point. Just to the left of this point, letters were presented one at a time in rapid succession; just to the right, numerals was similarly presented. Letters and numerals were close together, like two adjacent letters in a word on this page, so observers could see both without moving their eyes (an infrared beam reflected off their corneas detected eye movements and data from trials with eye movements were discarded). Their instructions were to attend to the letters until seeing a particular one (e.g., *B*), and then to shift attention to the numerals and report the one they saw. The observers could perform this task with rates of stimulus projection in excess of 20/s. The time between the critical letter and each numeral was known, so the numeral reported made it possible to derive the time taken to switch attention from the letters to the numerals. The observers couldn't report the one that appeared at the same time as the critical letter; instead, they reported one that appeared some fraction of a second later. In other words, shifting attention took some time; these latencies or reaction times showed how much.

Visual Imagery

> "Think of a cube, all six surfaces of which are painted red. Divide the cube into twenty-seven equal cubes by making two horizontal cuts and two sets of two vertical cuts each. How many of the resulting cubes will have three faces painted red, how many two, how many one, and how many none?" It is possible to solve this without seeing the cubes.... But the solution is easier if one can actually see the twenty-seven small cubes and count those of each kind. This is easiest in the presence of actual cubes, of course, and even a sketchy drawing will provide useful support, but many people solve the problem visually without visual stimulation. (Skinner, 1953, p. 273)

Even if a shift of attention is a response, what about more complex private events such as visualizing or imagining? We considered imagery in connection with mnemonic techniques in Chapter 26. The methods for measuring imagery are necessarily indirect, but experimental techniques have been developed to study it (Paivio, 1971). Here again, temporal measures have been particularly effective.

For example, observers were shown pictures of pairs of three-dimensional figures, with instructions to report whether the figures were the same or different by pressing a left or right button (Shepard & Metzler, 1971). When they were the same, one was rotated relative to the other, as illustrated in Figure 28-1. For such figures, the latency of the report that they were the same was linearly related to the difference in orientation, in degrees of rotation. In other words, mental rotation, or rotation of an image, has some of the same properties as rotation of an actual object: In both cases, the time taken to do the rotation is proportional to the angle through which the object rotates. The behavior of pigeons is different; when they're taught to discriminate same versus mirror-image rotated figures, their reaction times don't increase with the amount of rotation of one figure relative to the other (Hollard & Delius, 1982).

By treating the image as a thing we mustn't be misled into using it to explain behavior; we shouldn't seek a screen somewhere in the observer's head on which the image is projected. The controversy over what the image is (pictorial versus propositional or verbal) is less important than acknowledging that imagining the object, like the rotation itself, is something the observer does (cf. Kolers & Smythe, 1979). The temporal properties of rotating real and imagined objects are similar because of the correspondence between what observers do when seeing objects and what they do when imagining them. This correspondence demonstrates the functional equivalence of imagery and movement (e.g., Johnson, 1982). **Imagining is visual behavior in the absence of a visual stimulus.**

It has been suggested that John B. Watson, the founder of behaviorism, denied the existence of images because he himself was incapable of visual imagery (Skinner, 1959b). His denial, unfortunately, retarded the behavioral analysis of imagery. What is at issue is the role of private events in accounts of behavior. Our primary concern must be to determine their properties rather than to base accounts of other kinds of behavior upon them. A behavioral account doesn't deny the existence of such events, but it limits the conditions under which they can properly serve as explanations of behavior (cf. Paivio, 1975, p. 287).

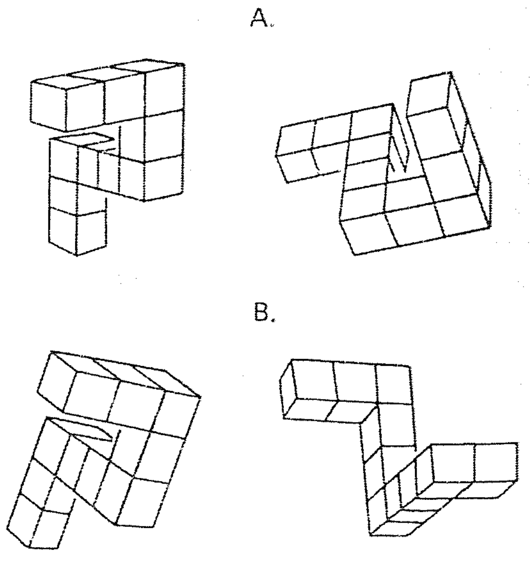

Figure 28-1 Pairs of figures from a study of mental rotation in which observers were instructed to report whether two figures were the same or different. In A, they're the same but the right figure is rotated 80 degrees from the left one in the picture plane; in B, they're different and can't be matched by any rotation. (From Shepard & Metzler, 1971, Figure 1)

The example of mental rotation used visual stimuli. The response of visualizing also occurs in the absence of such stimuli, as we saw in the case of visual encoding as a mnemonic technique. Such cases tempt us to speak of images as things seen rather than of imagining as behavior. In the photographic memory of someone with eidetic imagery, for example, reports of the details of past scenes are accompanied by eye movements similar to those when scanning current visual scenes (e.g., Haber, 1969). Yet to say that the reports or the eye movements are caused by the remembered image doesn't explain anything; an account of visualizing must be derived from the behavior of seeing and not from hypothetical stimulus properties of private events.

Consider more research on visual imagery (Moyer & Dumais, 1978). Observers were given pairs of words (e.g., *mouse* and *elephant*) and were instructed to choose the larger. The greater the difference in size between the two named classes, the shorter the reaction time (e.g., the response to *truck-cat* was faster than that to *book-chair*). The words themselves didn't differ systematically in size, but the relation between size difference and reaction time was similar to that with pictures instead of words as stimuli. These and related experiments (e.g., Paivio, 1975) suggest that the task involved imagery occasioned by words, presumably analogous to the visual encoding of Chapter 26.

In another experiment, learners are given pairs of statements such as *A is taller than B* and *A is shorter than C*, or *X is to the left of Y* and *Z is to the right of Y*, and are instructed to indicate whether such statements as *B is taller than C* or *Y is to the left of Z* are true or false (e.g., Huttenlocher, 1968). Latency measures in this verbal task and analogous visual tasks suggest that performance in the verbal task is better described as visualizing spatial relations than as verbal responding occasioned by the statements. The private construction of a spatial array in this task seems to parallel the behavior of moving objects into various spatial arrangements. Such experiments identify visualizing as a behavior class and then demonstrate that it enters into what the learner does.

Treating imagery as a kind of behavior as opposed to something an observer "has" or "doesn't have" raises the possibility that visualizing can be taught (cf. rehearsal of pictures: Graefe & Watkins, 1980). An artist presumably learns some of this behavior in progressing from sketching live models to sketching without a model. Visualizing might be shaped by gradually dimming a scene as an observer describes or sketches it or gradually increasing the time between presentation of a scene and the observer's description or sketch. There are few systematic studies of such phenomena.

Mental Representations

The temporal properties of attending and visualizing and searching help us to determine the structure of these behavior classes. Other structural aspects of behavior are its sequential and hierarchical organization. For example, consider the sequential and hierarchical components of completing a college major. The major consists of individual courses the order of which may be constrained by prerequisites; the courses may be decomposed into smaller units consisting of assignments and examinations; these in turn may be further reduced to reading particular passages or answering certain questions, again perhaps in specified orders; and so on. The completion of the major corresponds in structure to the contingencies established by the educational environment. The development of cognitive structure can be regarded as the development of correspondences between the structure of the environment and the structure of behavior. Analyses of concept learning or strategies in problem solving are sometimes based on demonstrations of such correspondences (e.g., Garner, 1974; Markman, Horton, & McLanahan, 1980).

The development of cognitive structure has been a central theme in the work of Piaget (Fischer, 1980; Piaget & Inhelder, 1969). Only a brief and inevitably oversimplified discussion is possible here. The several phases of child development in Piaget's system (sensory-motor, preoperational, concrete operational and formal operational) correspond to a progression from relatively simple relations between motor responses and their consequences (as in reaching for and manipulating objects) to complex relations that depend on correspondences between both verbal and nonverbal behavior and environmental and behavioral structure.

Piaget speaks of the development of such correspondences in terms of *accommodation* and *assimilation*; children must accommodate to the constraints that environmental structures and contingencies impose upon their behavior, but these structures and contingencies are assimilated to the extent that they become incorporated into their behavior. What is assimilated, however, is sometimes said to be a structure, often called a mental representation or a schema. Cognitive development, therefore, is assumed to be the gradual enrichment of representations of the world; these representations then become the basis for behavior.

Piaget studied biology and evolution, but with orthogenesis rather than Darwinian selection as the then dominant evolutionary theory; thus, it is probably fitting that his treatment of development often seems more like the unfolding of a progression of competencies than like their selection by contingencies (cf. Chapter 3). Unfortunately, some arguments against using contingencies in teaching have been based on an improper implication of that unfolding: if children haven't learned something yet it's only because they are not ready.

The phenomenon called *conservation* illustrates some aspects of cognitive structure dealt with in terms of representations. If the water in a short wide beaker is poured into a tall narrow beaker, a preschool child is likely to report that there is more water in the second beaker than in the first. This report is based on the higher level at which the water stands in the tall narrow beaker. In Piaget's account of development, this type of response is characteristic of the preoperational phase; one property of this phase is that responding is based upon single properties of environmental events rather than upon relations among properties. Later, in the concrete operational phase, the child says that the amount of water is the same in one beaker as in the other. The child's reports are still restricted to specific cases, and only in the subsequent formal operational phase does the child begin to speak abstractly of these relations in terms of the conservation of matter. These properties of the child's verbal behavior are also consistent with what we know about the acquisition of naming and related higher-order classes (cf. Chapter 14). Conservation is assessed by verbal reports, and it would be useful to know how nonverbal responses such as choices among beakers differently filled with favorite drinks would be correlated with the verbal competencies that define the several developmental phases.

One way to deal with conservation and related phenomena is to attribute the increasing sophistication of the child's performance to successive refinements in the child's mental representations of events in the world. As the child moves from actually manipulating objects to imagining or talking about or thinking about manipulating those objects, the structure of the child's mental representations may be said to approximate more and more closely the structure of the contingencies that operate for manipulating objects in the world. But the structure of these mental representations is derived from correspondences between behavior and environmental contingencies, so it is inappropriate to use them to explain that behavior (Wixted & Gaitan, 2002). The representations exist in the child only in a metaphorical sense; their concrete existence is in our own discriminative behavior as we observe the child or, in other words, in the consistencies and correspondences that we observe in the child's behavior.

The issue is the same as the one raised in the context of remembering, and it may be useful to extend an analogy introduced then. If I take a piano apart to find the music, I'll be disappointed. I may have just listened to a performance of Beethoven's

Waldstein Sonata, but when I'm done with the piano I'll have only pieces of wire and wood and felt and so on. I'll also be disappointed, as will you, if we try to find behavior inside the organism. The environment plays upon organisms as the pianist plays upon the piano. Just as it may take different virtuosi to bring out what is special about pianos and what is special about violins, what an environment brings out may depend on which organism it plays upon. Or, to take the analogy further, just as some musicians may bring out some music more skillfully from some instruments than from others, different environments may bring out some sorts of behavior more effectively from some organisms than from others.

But wait, you say, what about the player piano? In that case you can find a representation of the music: a roll of paper with holes punched in it. No matter that we can substitute, in more contemporary versions, a magnetic tape or a computer chip (I anticipated these features of the analogy when I chose a piano rather than some instrument less compatible with representations, such as a violin or a trumpet). The holes in the paper aren't music, and even though they can be translated into the action of the keyboard, we can't complete the account unless we can say how the holes got punched in the first place. If we wish to deal with representations, we must start by putting them not inside the organism but rather inside the human observer who discovered the spatial or relational properties of the environment to which the organism was responding. In other words, the issue isn't so much about the existence of representations as it is about which organism they belong in.

In any case, the development of conservation in children is not incompatible with properties of behavior that we've already considered. Several progressions must occur in parallel, and each is part of a hierarchy in which more complex classes of behavior are based upon simpler classes. At first, the child's discriminations are based upon a single dimension of objects (in this instance, height); discriminations based upon two or three dimensions (area and volume) come later. At the same time, the child's vocabulary becomes more finely differentiated, as from *big-little* to *bigger-smaller* and then to *taller-shorter*, *fuller-emptier* and so on (cf. Ward, 1980). Verbal behavior occasioned by specific instances becomes the foundation for the more general classes we call abstractions. Saying whether or not the quantity of water in one container equals that in another is not the same as choosing the contents of one or the other container or filling two containers to equal volumes, and none of these is equivalent to a statement of the principle of conservation of matter.

Should we be surprised that these separate skills and the coordinations among them take time, or that they are more likely to be mastered in some orders than in others. On the one hand, it's important to recognize how very different these classes of behavior are; on the other, it's important to recognize that equivalences across these different classes are inevitable consequences of the way in which verbal and nonverbal behavior is related to events in the world.

The argument here has been that organisms don't have to produce copies of stimuli before they can respond to them. But we should also note that not all representational accounts are copy theories. As we saw with regard to remembering, someone who says T upon seeing the letter has encoded or created a representation of the letter but has not created a copy of it (Conrad, 1964). An organism that has responded to a stimulus is a changed organism. "Whether internal representations are copies or interpretations of images, something called 'seeing them' is still required. Notions such as 'convex edge,' 'concave edge,' and 'occluding edge' are a step in the right direction. They are the beginnings of an analysis of the stimulus rather than a replication" (Skinner, 1988, p. 337). Cognitive processing doesn't require copies. For example, no copies are involved in the simulations of complex behavior created by the variety of computer programming called parallel distributed processing (Donahoe & Palmer, 1989; cf. Chapter 3 on copy theories in phylogenic selection). "Organisms are changed by contingencies of selection, they do not store them" (Skinner, 1988, p. 472).

ADDENDUM 28A: PROCESSING STAGES

Temporal measures such as reaction times have been important features of analyses of cognitive processes. Besides showing that private events take time, they have also demonstrated other properties of these processes. One concern is how cognitive tasks can be decomposed into their separate components or stages (cf. Posner, 1978):

> One of the oldest ideas in experimental psychology is that the time between stimulus and response is occupied by a train of processes or *stages* - some being mental operations - which are so arranged that one process does not begin until the preceding one has ended. This *stage theory* implies that the reaction-time (RT) is a *sum*, composed of the durations of the stages in the series, and suggests that if one could determine the component times that add together to make up the RT, one might then be able to answer interesting questions about mental operations to which they correspond. (Sternberg, 1969, p. 421)

Early treatments attempted to identify particular stages. For example, it was assumed that a discriminative stage could be estimated by subtracting simple reaction times to a single visual stimulus from discriminated reaction times to one of two stimuli. One problem was that stage durations were highly variable, presumably because most tasks can be performed in varied ways. A second and more important problem was that stage durations were calculated by assuming that the presence or absence of one stage has no effect on the duration of others. But consider one task made up of stages A-B-D and another made up of stages A-B-C-D. If D is longer or shorter when following B than when following C, the actual duration of C won't equal the difference between the times to complete the first and the second task.

One solution is to devise tasks in which a stage is repeated several times (e.g., A-B-C-C-C-C-D). Here, the immediate contexts of stages B and D are unaltered by changing the repetitions of C. Thus, the increase in latency produced by adding a repetition of C may be taken as the duration of C. This was the rationale for experiments (Sternberg, 1969) in which an observer was first given a set of digits, the positive set (e.g., 1, 3, 4, 9). Then digits from 0 through 9 were presented visually, and instructions were to press one button if this new digit was from the positive set and another if it wasn't. Latencies from digit presentations to button presses were recorded with different numbers of digits in the positive set; data were discarded when responses were incorrect.

To speak of search here is again a metaphor, but what properties does such a metaphorical search have? The observer must decide whether the new digit matches one in the positive set, so let's assume that the observer searches in some sense by comparing the new digit with the positive-set digits. Can we tell whether the observer completes all the comparisons before dealing with a match (*exhaustive* search) or stops as soon as a match is found (*self-terminating* search)?

Each possibility has different quantitative implications. In exhaustive search the observer compares the new digit with all positive-set digits, even if it matches one of them. This means that each digit added to the positive set should add a fixed amount to every reaction time. With self-terminating search, however, a match will be found on the average after only half the comparisons for digits in the positive set (any comparison from first to last might yield a match), whereas all comparisons must be tried for digits with no match. Now digits in the positive set on the average add only half the time to each reaction time than digits not in the positive set.

Paradoxically, the data show that search in this task is exhaustive rather than self-terminating. Having found a match, why does the observer continue with comparisons through the end of the list? One possibility is that comparing and then dealing with a match are themselves separate stages; if switching between them takes time, then exhaustive comparisons may be more efficient with relatively short lists than individual comparisons each followed by dealing with a match. The task is highly specific, so changes in detail change the outcome (cf. Baddeley, 1976; Crowder, 1976). For example, some variations produce data characteristic of self-terminating rather than exhaustive search. The metaphor of search is strengthened by the quantitative detail of the data, but its range of application is limited.

Furthermore, the metaphor is based on the assumption of *serial search*, in which the observer makes comparisons one at a time, rather than *parallel search*, in which items are compared simultaneously instead of successively. Parallel search

seems to imply that search should take a fixed time regardless of how many items are in the positive set, but it is as reasonable to assume that it becomes slower as number of items increases. Thus, it is possible to develop accounts in terms of parallel search that predict outcomes equivalent to those for serial search (Townsend, 1971).

This is just one example of many lines of research concerned with cognitive processing. For example, *top-down* versus *bottom-up* processing are distinguished by whether we start with major categories and then move down to the finer details or start with the details and then work up to the whole from the parts (e.g., Kinchla & Wolfe, 1979). The distinction is perhaps akin to starting with a molar description of behavior and then moving down to its molecular constituents as opposed to starting with the molecular constituents as foundations and then moving up to the molar behavior.

Information processing is another common metaphor, but to the extent that it implies that the organism attends to and in some way interacts with informative stimuli without regard to their reinforcing or aversive properties, it is undercut by the literature on observing responses (cf. Watkins, 1981, and Chapters 11 and 16). Information has very specific mathematical properties in the analysis of information transmission (Shannon & Weaver, 1949), but the term is too often used in a general nonmathematical sense.

Chapter 29
Problem Solving

The concept of classes is implicit in the etymology of symbol, *which by combining the Greek* sym-, *together, and* ballein, *to throw, suggests the creation of a unit from separate parts. Combined with* dia-, *across,* ballein *provides via* diabolic *the root for* devil, *and combined with* pro-, *before, it provides the root for* problem.

The Indo-European root leu-, *to loosen or divide, is the source of* loss *and the suffix* -less. *Probably in combination with* se-, *apart, it provides* solve *and* solution*; in combination with* an-, *up or again, it provides* analysis. *The origins of synthesis can be found in the Greek roots,* syn-, *together, plus* tithenai, *to put.*

Simulate *shares its roots with* same *and* similar *in the Indo-European* sem-, *in the sense of* one, together, *and perhaps also* like.

We examined only a few samples of research in cognition in Chapter 28. Cognition is after all not learning. Yet many issues in cognition are closely related to problems in learning. As the last experimental topic in this book, we turn to problem solving, an interest in which can be traced back to the earliest days of the psychology of learning (Köhler, 1927; Thorndike, 1898, 1911). Its cognitive treatment often emphasizes the structure of problems. An analysis of problem solving also raises functional questions, as when we're concerned with conditions that make the solution of a problem more or less likely. Science itself is often a variety of problem solving, so we'll also consider some implications for a behavior analytic philosophy of science.

THE STRUCTURE OF PROBLEMS

Let's begin with a structural approach to problem solving in the Hobbits-and-Orcs problem (Thomas, 1974). Three hobbits and three orcs are traveling together. The orcs won't leave the hobbits behind, but they'll overpower any hobbits they outnumbered. The group reaches a river that must be crossed and finds a boat that can hold only one or two at a time. How do the hobbits organize the crossing so that no hobbits are ever outnumbered? This problem was presented to solvers to explore the effects of feedback and other variables.

KEY TERMS: Problems, Problem Solving; Conjunction Fallacy, Confirmatory Bias; Positive and Negative Transfer; Functional Fixity, Simulation.

The problem allows only a few possible moves. Its solution can be presented economically as in Figure 29-1. Except for an alternative first move not shown and the two branching moves at the beginning and end of the sequence, the only allowable alternatives to correct moves are ones that reverse direction through the sequence; solvers seldom recognized this property of the problem. For our purposes, it is enough to note that errors were more likely at some points in the sequence (states 321 and 110) than at others. These and other features of the data indicated that a solver's performance did not consist simply of the separate moves. Instead, the solution was based on larger units consisting of sequences of moves that led to some intermediate arrangement of hobbits and orcs on the two sides of the river. Thus, the analysis demonstrated correspondences between the structure of the problem and the structure of the solver's solution.

Another problem, the Tower of Hanoi, is illustrated in Figure 29-2. Disks of different sizes with holes in them can be stacked on any of three pegs, but disks can be placed only on empty pegs or on other disks that are larger. In the figure, three disks start on the left peg. By moving one disk at a time the solver must get them all to the right peg, the finish at lower right. Contingencies may also be placed on the solver doing so in the fewest possible moves. With three disks the problem can be solved in seven moves, as illustrated along the right side of the figure. The problem can still be solved if the number of disks increases, but the more the disks the longer the minimum sequence of moves (Anderson, 1993).

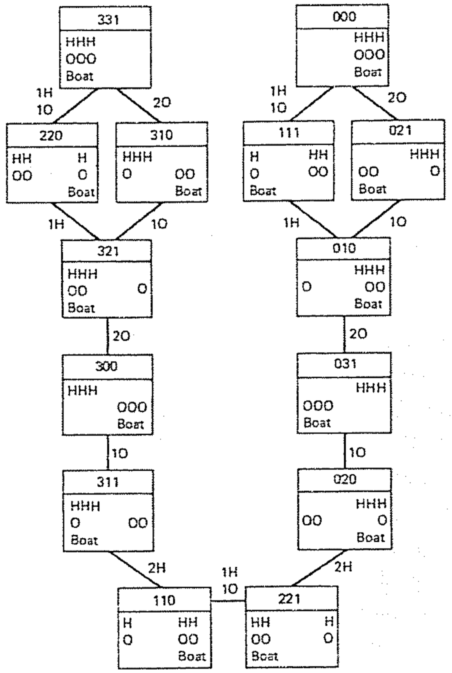

Figure 29-1 Successive states in solving the Hobbits-and-Orcs problem. Three hobbits and three orcs must cross a river. A single boat holds only one or two creatures and orcs may never outnumber hobbits. Three digits code each state: number of hobbits on the starting side, number of orcs on that side, and the boat on that side. The hobbit (*H*) and orc (*O*) and boat positions are shown below the code in each state; the transported creatures are shown in the transitions between states. Except for a transition to the dead-end state 320 from the first state (crossing of one orc and the boat, not shown), no other moves in which orcs do not outnumber hobbits are possible. (Adapted from Thomas, 1974, Figure 1)

Unlike the Hobbits and Orcs problem, the Tower of Hanoi problem can be solved by traveling different paths (Anderson, 1993). Adjacent pegboards in the figure show positions differing by one move, so all possible ways of tracing a path through this diagram from start to finish qualify as solutions, though many would require far more moves than the optimal one. As with the Hobbits and Orcs problem, solutions can be decomposed into smaller

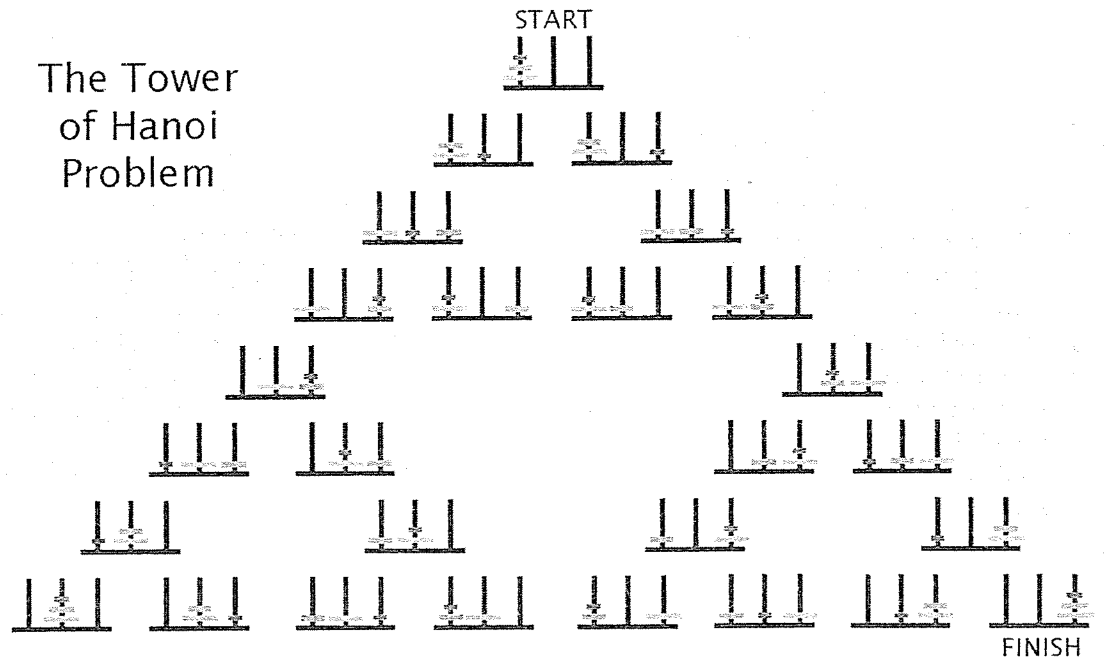

Figure 29-2 Schematic diagram of possible steps in the Tower of Hanoi problem, in which the three disks on the left peg at START must be moved one at a time with no disk ever being placed on one that is smaller, with the FINISH configuration as the target. Adjacent configurations of the pegboards show possible moves. The optimal solution is the sequence shown along the right edge from START to FINISH. (Adapted from Anderson, 1993, Figure 1)

units, as in moving just the top two disks from one peg to another. The decomposition is more obvious with a four-disk problem. With the four disks on the left peg, as in START, I'd first have to move the top three disks to the middle, as at the lower left of the figure. Then I could transfer the largest fourth disk to the right peg, after which I would again have to move the three-disk stack, this time from the middle to the right peg. Thus, these three-disk moves become subunits of the solution. With additional disks, the first steps are always those of moving the top three disks, but where should those three go in a problem that started instead with five or six or even more disks? There is a general solution, but I'll leave that to you. Solving the Tower of Hanoi problem has sometimes been regarded as an intelligence test (as for the chimpanzee Caesar in *The Rise of the Planet of the Apes*).

Artificial Intelligence: Chess

Problems in which the steps are as explicit as in the Hobbits and Orcs or the Tower of Hanoi problems lend themselves well to computer solution (e.g., Newell, Shaw, & Simon, 1958). The computer can perform large numbers of calculations in short periods of time. In problems involving well-defined alternatives it can select those branches leading to a solution more rapidly than a human can. Part of the field of *artificial intelligence*, or *AI*, is devoted to the design of computer programs for simulating such activities as problem solving. But computer programs are limited when those who write the programs can't provide an exhaustive list of alternatives at some steps in the solution of a problem or can't define the terms entering into it or can't reduce it to a manageable size even for a computer (e.g., Dreyfus, 1992; Winograd, 1980). Each of these constraints was also important in limiting the application of computer programs to language translation (cf. Chapter 25).

Computer programs designed to play chess are a case in point (Frey, 1977). Chess-playing programs have gradually evolved to play better chess. The best programs now rarely lose to humans, even chess grandmasters. For the purposes of a computer program, chess is well-defined: all possible moves of each piece in any position are exactly specified by the rules of the game. Given 16 pieces on each side and a board of 64 squares, however, possible moves and countermoves multiply so rapidly that even the most rapid programs on the largest computers are limited in the number of moves ahead that they can calculate. With increasing moves the number of alternative moves exceeds the number of atoms in the universe. If the computer doesn't find a forced checkmate of its opponent in, say, the next 50 moves, how does it evaluate the relative strengths of all those possible future positions so as to select its best next move? More accurately, how does a human write a program that allows the computer to do that?

Some of the most spectacular chess games are those in which a checkmate is forced through the sacrifice of several pieces. If the computer evaluates positions on the basis of the relative number of pieces lost and if the mating move itself is just beyond the horizon of its calculations of future moves, it will never embark upon such a forced mate. The computer plays chess carefully; it doesn't make the mistakes of the amateur human player, such as exposing an unprotected piece to capture by the opponent. Nevertheless, chess-playing programs that now defeat human grandmasters do so in part by brute force, in the sense that they include extensive libraries of many standard chess openings and explore very many alternatives a great many moves ahead. They don't play chess like human grandmasters, who not only calculate the consequences of various moves and countermoves but also see various strengths and weaknesses to be exploited in the pattern of pieces on the board. Given the computing power that has been brought to bear, it is amazing that any human can be competitive against a computer. Nevertheless, the most successful future chess-playing programs will probably be those that most closely capture the structure of what the chess grandmaster sees in a form that can be used by the computer. The writers of such programs will presumably be expert in both chess and computer programming.

Predictions that computer programs capable of defeating any human chess grandmaster are imminent go back some decades, and perhaps that time has at last come. Meanwhile, there is a curious irony in the successes of artificial intelligence. Instead of teaching things to children, the programmers of artificial intelligence have figured out how to teach things to computers. We might have hoped that what they discovered would be applicable to improvements in education, but that hasn't happened. Part of the problem may be that they have not yet defined intelligence in a way that relates it to behavior. But more likely it is that teaching chess or other skills to computers doesn't have anything to do with teaching things to children because computers and children do things in such different ways. Some computer programs are constructed to be sensitive to their consequences (e.g., Sutton & Barto, 1998), but the learning of most is still better described as governed by rules than as contingency-shaped.

Expert Performance

If chess grandmasters could describe what they did when analyzing a chess position, their descriptions could be used to design chess-playing computer programs. Playing chess by following the instructions of such a program would be verbally governed rather than contingency-shaped chess playing (cf. Chapter 23). But the higher-order behavior classes of chess masters haven't been identified; grandmasters can't describe what they do in a way that can be translated into a computer program (e.g., Chabris & Hearst, 2003). Grandmaster chess play is called intuitive, which is another way of saying that it is contingency-shaped rather than verbally governed. This kind of expert performance is not independent of verbal behavior, however, which suggests that there may be more than one variety of contingency-shaped behavior.

In discussing limits on the capacity of computers to simulate human judgments, Dreyfus and Dreyfus (1986) made some observations about expert performance. Their examples included chess and medicine among others. They noted that education in complex human skills often begins with verbal antecedents, or rules. They suggested that proficiency develops when intuitive judgments begin to mix with rules and that expertise involves performance that no longer depends on the rules. Their distinctions between the functions of rules and intuitions are much like those between verbally governed and contingency-shaped behavior, but for them verbally governed behavior comes first; intuitive or contingency-shaped behavior eventually grows out of it.

Their distinction is compatible with our earlier treatment of verbally governed behavior if we expand our taxonomy to include three classes of skilled behavior: contingency-shaped behavior that has never depended on verbal antecedents, corresponding to the kinds of contingency-shaped behavior ordinarily treated in behavior analysis, as in nonhuman behavior and human nonverbal behavior; verbally governed behavior, in which verbal antecedents override effects of nonverbal contingencies (Chapter 23); and, finally, expert performance, in which continuing contact with the environment attenuates control by verbal antecedents and allows behavior to become sensitive to subtle changes in contingencies, in what might be regarded as a second and different kind of contingency-shaped behavior that has not yet been explored in detail.

We can't write the rules of human expertise even in restricted domains such as chess. But that poses a problem, because we don't have ways to incorporate human expertise into computers except by writing programs that tell computers what to do. If we can't state the rules, we can't write the program. Unless and until new computer technologies such as parallel distributed processing (Donahoe & Palmer, 1989) can introduce contingencies of selection into computation, computer simulations of human expertise will inevitably be limited to simulations derived from verbally governed behavior.

The IBM computer, Watson, provided another example when competing with champion human players on the quiz show *Jeopardy!* (Markoff, 2011). It won easily, but its answers when it missed questions were not like the errors the human players made. For example, to a clue asking for the name of an art period (*Modern*), Watson gave the name of an artist (*Picasso*). It had a built-in advantage not only in memory (four terabytes of stored data) but also in its reaction time, which could be as short as 10 ms; humans could get in by anticipating the opportunity to answer, but without anticipations minimal human reaction times are in the 100 to 150 ms range. Thus, the touted victory for artificial intelligence was perhaps overrated. A system that can answer questions with the speed and breadth of google searches is of course fascinating, but it works in ways that have little in common with the ways humans answer questions.

The rules of chess are defined very explicitly. If the problems to be surmounted in writing expert computer programs are so formidable even in this limited domain, what about those domains in which the terms are less well-defined? Furthermore, once a computer program has been written to solve one problem, how can it be used to solve other problems? The question is relevant not just to computer programming but also to human problem solving, to which we now return. In the teaching of problem-solving skills within well-defined areas such as mathematics (e.g., Wertheimer, 1959), for example, analyses start with the solution of single problems and then move on to the transfer of solving strategies from one problem to another. Experiments on problem solving, therefore, have often dealt with the effects on problem solution of the solver's history of solving problems.

TRANSFER

Does mastery of a classical language like Greek or Latin make it easier to learn economics or history or sociology? Does the discipline of mathematics make people more logical thinkers? Does the mastery of music or art create skills useful for studying great works of literature? Research on transfer of learning or transfer of training began with simpler questions than these (Thorndike & Woodworth, 1901). A common assumption of earlier times was that traditional educational curricula were appropriate not just because each discipline might later be important but also because they "exercised the mind" or, in other words, taught general intellectual skills. Studies of how the learning of one task affected the learning of another were particularly relevant to this view. Yet each demonstration that transfer was specific to a particular learning task or was limited in some other way contradicted rather than supported the view. Meanwhile the typical college curriculum gradually evolved from a fixed sequence of required courses to electives, distribution requirements and other flexible course options.

Thorndike played a critical role in this history (Jonçich, 1968). As his research interests shifted from instrumental learning in animals to human learning in educational settings, he became more involved in curricular matters such as course requirements for potential teachers. His findings on transfer probably influenced faculty committees responsible for establishing and reviewing course requirements and curricula. The transfer of learning from one subject matter to another is difficult to analyze. Some study skills probably generalize from one course to another, but examining how the mastery of a classical language like Latin influences a student's progress in mathematics or science or philosophy requires experimental procedures that are neither practical nor acceptable in typical educational institutions (e.g., random assignment of students to course sequences). Perhaps inevitably, research on transfer concentrated on simpler tasks, such as transfer from one paired-associates list to another or motor skills from one hand to the other (e.g., McGeoch, 1942; Woodworth, 1938).

In *zero transfer*, learning Y after X is essentially equal to learning just Y. In *positive transfer*, learning Y after X is more rapid or occurs with fewer errors than learning only Y; learning X makes Y easier to learn. In *negative transfer*, learning Y after X is slower or occurs with more errors than learning only Y; learning X makes Y harder to learn. The ease or magnitude of transfer from one task to another depends on similarities and differences between the old and the new items and the order of tasks (Müller & Pilzecker, 1900). Instances of learning don't occur in isolation; other learning preceded and is likely to follow. Every instance of learning may be affected by what came before and what follows, so proaction and retroaction are relevant to all learning (cf. Addendum 27A). For example, a child's learning of the alphabet may be affected by earlier nonverbal discriminations among geometrical patterns (e.g., Gibson, 1965). Uppercase letters, unlike lowercase ones, don't include any up-down or left-right reversals (compare *P*, *B*, *D* and *p*, *b*, *d*), so transfer from uppercase to lowercase alphabets may depend on whether the child has already learned nonverbal up-down and left-right discriminations.

Functional Fixity and its Relatives

Failures of transfer are sometimes called *fixations*). The Luchins' water-jar problem provides a relevant experiment (Luchins & Luchins, 1950). Given a water supply and three jars of different capacity, the problem is to produce a specified amount of water. For example, if jars A, B and C have respective capacities of 21, 127 and 3, how can you arrive at exactly 100 units of water? In this instance, the solution is to fill B and then to pour out enough to fill A once and C twice. One group was given several successive problems that could be solved this way while a second group wasn't. Later, both groups were tested on problems such as the following: How does one arrive at a quantity of 20 if jars A, B and C have respective capacities of 23, 49 and 3? The first group solved the problem in the same old way, whereas the second group solved it more efficiently, by filling A and then pouring out enough to fill C.

In this example of negative transfer or problem-solving rigidity, the consequences of past behavior affected current behavior. The literature on problem solving includes many variations (e.g., Saugstad & Raaheim, 1960). For example, solvers instructed to mount a candle on a wall were given either a box of candles and a box of matches and a box of tacks, or separate stacks of candles, matches, tacks and boxes (Duncker, 1945). The solution was to tack a box to the wall so that it provided a flat surface on which a candle could be mounted. Those who received the candles and matches and tacks in boxes were much slower to solve the problem than those who received the boxes separately; they received the boxes as containers and continued to treat them that way. Instances of failure to solve a problem requiring an unusual use of common tools or materials are sometimes described as cases of *functional fixity*; the commonplace items occasion responses appropriate to their everyday functions instead of the novel responses appropriate to the problem solution.

Before we consider another class of examples, guess how many seven-letter words in this chapter end in *-ing*. Then guess how many end in -_n_, where the spaces can be any letter. (It might be a good idea to write your answers down.) The other class of example comes from experiments on decision theory and probability judgment. Decision theory is concerned with the finding that our estimates of the probabilities of events often differ substantially from the probabilities calculated from actual event frequencies (Kahneman, 2011; Tversky & Kahneman, 1983). For example, research on decision making shows that without regard to supporting evidence, the more often we've heard or read a sentence the greater are the odds that we'll say the sentence is true (Kahneman, 2011, p. 62).

Here's another case, based on the Monty Hall dilemma (Herbranson & Schroeder, 2010). You can choose one of three doors, behind one of which is a prize. After you've chosen a door, one of the unchosen doors is opened and has no prize behind it. You can then stick with the door you chose or switch to the other unchosen one. Which should you do? The show that is the source of this problem demonstrates that most people stick with their original choice rather than switching, even though the probability of getting the prize is twice as great for switching. The probability that you chose the door with the prize at the start was one in three. On the other two-thirds of your tries, the prize is behind a door you didn't choose, but in those cases you'll learn which of those two doors doesn't hide the prize. That means the odds of getting the prize by switching are two in three, twice the odds for getting it if you stay. But even given these probabilities the correct strategy seems counterintuitive. A version of this problem was arranged on three pigeon keys with food as the prize. Unlike humans, pigeons respond optimally by switching when one of the unchosen alternatives is removed from play, so that food is available for pecks on the other. The pigeon's behavior is contingency-governed, but once again our human verbal behavior gets in the way (cf. Chapter 23 on verbal governance).

Now let's discuss your guesses about the *-ing* and -_n_ endings. There's a good chance that they demonstrate the *conjunction fallacy*, which occurs when people judge the probability of several events occurring together as greater than the separate probabilities of each. If your first number was bigger than your second, your answer was mathematically impossible. At least as many words must end in -_n_ as end in *-ing*, because those ending in -_n_ include all of those ending in *-ing*.

Consider a different kind of judgment in the context of an imaginary diagnostic situation in a psychiatric hospital. You're a staff member who has discovered a new clinical syndrome called *narapoia*. Narapoids are the opposite of paranoids; they are under the delusion that other people are plotting to do them good. You know that narapoids cannot possibly make up more than 1% of your clinical population, and you have developed a test for narapoia that has only a 5% error rate. You give the test to someone who is being screened for admission to the hospital and the person scores as a narapoid. Does the test score provide sufficient grounds for admission?

To make this question easier to answer, assume you give the test to 1000 patients in your hospital. No more than 1% are narapoids, so they total 10 or so; with your 5% error rate most or all of them will correctly score as narapoid. But of the 990 or so others, roughly 50 will be false alarms; they are not narapoid, but because of your 5% error rate they will score so. In other words, you'll get a total of about 60 narapoid scores, but given that only 10 or so are actually narapoid the chance that anyone with that score has been correctly diagnosed is only 1 in 6. You cannot justify admission on the basis of the test alone; you need other diagnostic criteria. And perhaps that's just as well, because you'll be hard-pressed to devise an effective treatment: Whenever you try to help your narapoid patients, you'll only be confirming their delusions.

Narapoia is an imaginary syndrome (Nelson, 1951), but the probability relations it illustrates are of practical significance. For example, even with low error rates such judgments can lead to many false alarms in cases where false alarms can be very costly, as in lie detection or medical diagnoses. Imagine a criminal case in which someone facing a long jail sentence tells the truth but fails the lie-detector test, or someone gets incorrect results on a test for drugs or for AIDS or for cancer. These types of biases in probability estimation are likely whenever there is some error as we sample for events that occur with low frequencies in large populations. Most of our everyday probability judgments are made in the context of more symmetrical distributions of event probabilities. Our histories of such judgments, as in functional fixity, transfer too easily to extreme cases in which the familiar probability relations no longer hold. Teaching the mathematics of probability is one way to make such probability judgments more accurate; another is to provide experience with probability judgments over a range of situations that include these sorts of extremes. Note that the former solution involves verbally governed behavior whereas the latter involves contingency-shaped behavior.

Suppose you're shown cards with letters on one side and numbers on the other and you're told that every card with a vowel on one side has an odd number on the other. Four cards are now dealt out with A, 8, X and 7 facing up and you're asked to turn over the minimum number of cards to verify the rule *if vowel then odd number* (cf. Wason & Johnson-Laird, 1970). How many do you turn over and which should they be? Most of us will turn over the A, to see whether an odd number is on the other side. Those that do turn over a second card usually try the 7, to check for a vowel on the other side. But it won't disconfirm the rule if they find a consonant, because the rule doesn't say that consonants can't have odd numbers too. You need to turn over two cards, but the other one to try is the 8, which will disconfirm the rule if there's a vowel on the other side. Only a small proportion of people choose this card.

Why? The point is that people are more likely to seek confirmation than disconfirmation of their hypotheses. Turning over a card is an observing response, and we all have long histories of checking out our guesses about things. Presumably we check on things that are important to us, so in the past it's been nicer to have our guesses confirmed than disconfirmed. Over time we've probably learned that some kinds of guesses can only be confirmed and others can only be disconfirmed, and we guess accordingly. If that is true, it unfortunately means that sometimes we think we've confirmed something when we haven't. Perhaps then the truth was something we didn't want to know (cf. Chapter 16 on observing responses).

THE CONSTRUCTION OF SOLUTIONS

Problem solving can be affected by antecedent events and by contingencies, but what constitutes problem solving? Consider Skinner's account of the contingencies at an airport baggage claim if you're picking up a friend's suitcase:

> You have never seen the suitcase or heard it described; you have only a ticket with a number from which a match is to be found among the numbers on a collection of suitcases. To simplify the problem, let us say you find yourself alone before a large rotary display. A hundred suitcases move past you in a great ring. They are moving too fast to be inspected in order. You are committed to selecting suitcases essentially at random, checking one number at a time. How are you to find the suitcase? You may, of course, simply keep sampling. You will almost certainly check the same suitcase more than once.... A much more effective strategy is to mark each case as it is checked—say, with a piece of chalk. No bag is then inspected twice, and the number of bags remaining to be examined is reduced as rapidly as possible. Simple as it seems, this method of solving the problem has some remarkable features.... It is the use of the chalk which introduces something new. Marking each suitcase as it is checked... is constructing a discriminative stimulus. (Skinner, 1969, pp. 136-137)

Problem solving, in other words, is behavior. The discriminative features of the situation define the problem, and the reinforcer is the solution of the problem. Often we solve problems by manipulating stimuli: looking something up, converting a verbal problem into a mathematical equation, and so on. It is important to recognize how we can alter our own behavior by changing something in our environment. Suppose we multiply 23 by 14 in this form:

$$\begin{array}{r} 23 \\ \underline{\times\ 14} \\ 92 \\ \underline{230} \\ 322 \end{array}$$

In the intermediate products, 92 and 23, we create discriminative stimuli that allow us to reach the solution, 322. But what if we had to solve the problem without a pencil? Presumably the intermediate products would still enter into the solution even we had no written record of them. If we didn't say them aloud, an observer might say we'd been doing mental arithmetic (cf. Ashcraft, 1982). But the intermediate products have the same role in both cases, even though they are more public and permanent in one than the other.

SIMULATIONS

Our imagining isn't limited to the visual mode. We not only visualize; we hold imaginary conversations, take imaginary journeys and perform imaginary actions. These are all *simulations*, imitations in the absence of relevant stimulation of some parts of the behavior that might occur in some situation. Games have been used to simulate various properties of economic, political and social situations. Just as war games played on game boards or computers or in the actual field may demonstrate potential consequences of various strategies (e.g., Sabin, 2012), so also our imaginings may bring us into contact with possible consequences of our own actions. But just as military simulations may be imperfect because they fail to incorporate important variables in a combat situation, our imaginings will also often be fallible because they involve not real contingencies but only our own partial re-creations of them.

Dawkins (1976) makes the point by comparing computer simulations and human imagination. With respect to the computer simulation:

> A model of some aspect of the world is set up in the computer. This does not mean that if you unscrewed the lid you would see a little miniature dummy inside with the same shape as the object simulated. In the chess-playing computer there is no "mental picture" inside the memory banks recognizable as a chess board with knights and pawns sitting on it. The chess board and its current position would be represented by lists of electronically coded numbers.... But it does not matter how the computer actually holds its model of the world in its head, provided that it holds it in a form in which it can operate on it, manipulate it, do experiments with it, and report back to the human operators in terms which they can understand. Through the technique of simulation, model battles can be won or lost, simulated airliners fly or crash, economic policies lead to prosperity or to ruin. (Dawkins, 1976, p. 62)

In other words, the important properties of the computer program aren't whether they generate or manipulate copies of the world but rather whether they operate in ways analogous to the workings of events in the real world (cf. blueprints versus recipes in Chapter 3 on). Dawkins continues by discussing similar features of human behavior, referring to organisms as survival machines that have been built by their genes:

> when you yourself have a difficult decision to make involving unknown quantities in the future, you do go in for a form of simulation. You *imagine* what would happen if you did each of the alternatives open to you.... Survival machines which can simulate the future are one jump ahead of survival machines who can only learn on the basis of overt trial and error. The trouble with overt trial is that it takes time and energy. The trouble with overt error is that it is often fatal. Simulation is both safer and faster. (Dawkins, 1976, pp. 62–63)

Organisms that have evolved with a capacity to simulate some consequences of their own behavior have some obvious advantages over those that haven't. And once just a little bit of simulation has become possible, natural selection is likely to produce organisms that can more and more effectively simulate the environmental contingencies with which they come in contact (cf. Gallistel, 1990).

SCIENCE AS SEEING

Science is often a matter of problem solving, as when an experimenter must figure out how to build a new piece of equipment or how to keep some extraneous variable from contaminating an experiment or, sometimes more important than anything else, how to come up with funding and other support for a project. But science is also a method for developing new ways of seeing. That has become especially clear in biology: "Evo Devo can trace the modifications of structures through vast periods of evolutionary time—to *see* how fish fins were modified into limbs in terrestrial vertebrates, how successive rounds of innovation and modification created mouthparts, poison claws, swimming and feeding appendages, gills, and wings from a simple tubelike walking leg, and how many kinds of eyes have been constructed beginning with a collection of photosensitive cells" (Carroll, 2005, p. 10); "... we can now see how the endless forms of different animals actually take shape" (Carroll, 2005, p. 13).

Philosophers of science have debated the nature of knowledge and the logic of scientific methods, but science is not so much a body of knowledge as it is a way of discovering the answers to questions about how things work. A behavior analytic philosophy of science must start not with assumptions about truth and knowledge but rather with the behavior of the scientist. It must provide behavioral alternatives to the major categories and assumptions of traditional philosophies of science (cf. Skinner, 1950, 1956, 1957).

For too many students, the starting place for contact with behavior analysis is through verbal behavior: they listen to lectures or read texts. But that verbal behavior was originally established through direct contact with nonverbal behavior in the laboratory. Vocabularies treated mainly in terms of definitions cannot deal adequately with such origins. Mastering a definition cannot be counted on to produce the discriminations upon which that verbal behavior was originally based. The student who has learned to define *reinforcement* may offer a correct definition, but it doesn't follow that the student will then be able to reliably discriminate between instances of reinforcement and nonreinforcement in real-world settings (cf. Catania, 1993a; Hineline, 1990).

And what about explanations? *Explanation* is derived etymologically from roots implying the laying out or display of something (*ex-*, out, and *plain*, flat), but it has come to be defined in terms of giving reasons or of clarifying. As an example of the former usage, an auto mechanic might offer one of several explanations for the failure of my car to start: dead battery, empty fuel tank, bad starter, and so on. We accept the explanation if the car starts after we act upon it. Other explanations, such as those given as circular definitions, don't work as well for us, or at least shouldn't. If we define emotions such as happiness or anger or sadness based on how someone acts, it will not do when we see such behavior to try to explain it by saying the person was happy or angry or sad. If we define attitudes or preferences or prejudices based on how someone acts, it will not do when we see such behavior to try to explain it by saying the person had an attitude or a preference or a prejudice.

The auto mechanic's explanation for why the car didn't start is of a different sort than a scientific explanation. The scientific explanation shows us how something works. For example, to explain how a car works we might show its systems in operation, perhaps with visual displays rather than actual components: how the burning of fuel in the cylinder moves the piston, how that in turn moves the crankshaft, how crankshaft motion is transmitted through clutch and differential to the wheels, and so on. The most effective demonstrations clarify each part of a system by relating it to familiar phenomena with which it readily generalizes, as in showing the similarity between fuel ignition and other everyday kinds of combustion. In traditional philosophies of science, explanations are judged to be effective or valid when they relate what is to be explained to other familiar and well-established phenomena.

If asked how fuel is introduced into the cylinder or how its ignition is timed, we could go further and show how valves or fuel injectors work or how distributors cycle the activation of spark plugs. For some, this level of explanation might be satisfactory; others might ask for more detail. For example, few might be concerned with how exhaust valve timing is coordinated with fuel injection and ignition; of those, fewer still might care about the detailed working of timing belts.

Explanation in the sense of the car example is not very different from explanation in biology. Whenever a biologist demonstrates how strands of DNA combine and recombine, or how they can act as recipes for proteins, or how they replicate themselves, the biologist has successfully explained some part of the genetic functions of the cell. The double helical structure of DNA is no longer theory; it can be shown to us and it has become part of the definition of DNA that it can take this form. Precisely when this aspect of its structure stopped being theory it became maximally effective in explanation.

Both the car example and the DNA example illustrate that explanations vary in depth. Still, at any level explanation is showing how something works. In this sense, behavior analysts sometimes offer explanations of behavior. We sometimes show how shaping works by studying effects of consequences on subclasses of the responses being shaped, and we sometimes show how molar relations come about by showing how simpler processes combine to produce them. But we needn't be satisfied with this level of explanation and in the long run some will no doubt pursue explanations at more detailed levels, presumably including physiological ones.

Before we arrive at explanations we make guesses. Guesses are weakly determined instances of verbal behavior, but when we call them theories or hypotheses they sound as if they are more strongly determined. But it is useful to explore the implications of our guesses. When we formalize them we may call them simulations or models. Models are often mathematical, but they needn't be. For example, early human-scale physical models of DNA played a crucial role in establishing its double helical structure (Watson & Crick, 1953). What starts as a simulation or a model may eventually be shown. Synapses were once only theoretical entities, but the junctions between axons and bodies of neural cells are now studied directly, and synaptic transmission is explained in more depth whenever neuroscientists show in more detail how it works. A word that begins as a theoretical term can evolve into a name for a phenomenon.

We sometimes accept mathematical theories or models as explanatory, but they may be better treated as economical descriptions. For example, planetary motion is not explained by Kepler's laws; rather, those laws describe, to a reasonable approximation, the mathematical properties of the orbits. Newtonian mechanics is consistent with the properties of human-scale space, but other systems are more general (e.g., Einstein's relativity theory) or operate in other domains (e.g., quantum mechanics). The correspondences between mathematical systems and physical events are remarkable, but the sense in which Einstein's mathematics explains the bending of light as it passes near the sun or the implications of saying that the light bends because space is curved are matters of verbal behavior as well as physics. Just as these systems describe properties of events that occur in space, we describe properties of behavior.

Just as feelings and thoughts are best treated as accompaniments of behavior or as behavior itself rather than as causes of behavior (Skinner, 1945), so also a behavioral philosophy of science should treat theories not as causes of scientific behavior but as its products. If so, the origins of theory become of special interest. For example, what leads to good guesses about how behavior works? We might assume that good theories in a given domain come most easily to those who spend the most time exploring that domain but who also have contact with a broad range of phenomena outside it. It is inappropriate to assume that theories drive scientific behavior; they may instead be derivatives of it, with functions yet to be established. This is not an argument against theory; it is an argument for questioning traditional assumptions about its role.

Probably we should also take seriously the possibility that theories may hinder rather than enhance scientific discovery. Theories are, after all, verbal behavior, so they may generate the insensitivity to relevant contingencies that is the hallmark of verbally governed behavior. Mathematical models were one of the four distractors that Skinner (1961) identified among the alternative reinforcers that could lead psychologists to a flight from the laboratory. He did not include the flight to discussions of the philosophy of science, and that may be because a behavior analytic philosophy of science must inevitably return to behavior itself as its subject matter.

BACK TO THE INSIGHTFUL APE

Now that we've seen how problem solutions can be affected by antecedents and contingencies, we've come nearly full circle. Let's return to the problem Köhler set for his apes (Chapter 2), but with a different organism (Epstein, 1981). A pigeon has been placed in a compartment with a movable box in a far corner and, hanging out of reach, a model banana, chosen instead of a simple key in deference to Köhler. Earlier, the pigeon's pecks at the banana had been reinforced with access to food. Also, the pigeon was taught through shaping to move the box across the floor by pecking at it. Finally, the pigeon's feathers were trimmed so that it couldn't fly, and its jumping and wing-flapping dropped to a low rate in the presence of the hanging banana. Under these circumstances, the pigeon looks from the banana to the box and back again. Soon it goes to the box and moves it under the banana. Once the box is there, the pigeon climbs it and pecks the banana, operating the feeder.

This is novel behavior, a new response combination. The pigeon solved the problem of getting to the banana by moving the box and then climbing it, even though it had never done so before. This spontaneous combination of two operant classes is an example of adduction (cf. Chapter 13). Presumably it was important that the competing responses of jumping and wing-flapping had already been reduced in probability. The pigeon doesn't look like a chimpanzee, but its performance is functionally just like Sultan's in Köhler's insight experiment. The main difference between this demonstration and Köhler's is that we know a lot more about the pigeon's history than we do about Sultan's. An experimenter interested in problem solving might start simply by watching the solver's behavior. Like Köhler in his study of insight in apes, the experimenter will discover what the solver already can do. That's where we began. But now we can see much more clearly how the environments that were constructed by Köhler for his apes worked than we could then.

ADDENDUM 29A: TWO SIMULATIONS BASED ON THE RESERVE

Simulations that can show how things work are important components of science. I offer two examples here. The first is a shaping game. In the days of main frames and then Apple II and IIe computers and then pre-Windows computers running programs under MS-DOS, my colleague Eliot Shimoff began to program computer simulations for his Introductory Psychology section, a colleague who was a behavioral sociologist, Bud Matthews, did the same for his Introductory Sociology course, and I did so for the learning course. In our programs, students learned to estimate the correlations of data scatterplots; they saw how population growth grew, sometimes at a frightening pace and often accompanied by falling per capita income, as a function of different birth rates; they learned to distinguish among cumulative records produced by various reinforcement schedules and to identify the contingencies arranged by these schedules; and in simulations of both residential neighborhoods and cafeteria seating they saw how a mere mild preference for having a couple of neighbors like oneself could lead to segregated areas (Catania, Matthews, & Shimoff, 1989b; Matthews, Shimoff, & Catania, 1987b; Shimoff, Catania, & Matthews, 1987).

One shaping game was based upon the differential reinforcement of responses drawn from a probability distribution, as in Figure 9-2 (Shimoff & Catania, 1995). This and some other versions were designed on the assumption that it was better for students actually to do some shaping, even if only in a computer simulation, than just to read about it. The player saw a mouse pop up along the bottom of the screen, beginning on the left. By pressing the space bar, the player could reinforce its entry at that location. Otherwise that entry was unreinforced. Each reinforced entry raised the probability of future entries at and around that location; each unreinforced one lowered those probabilities. The player watched the probability distributions change with the progress of shaping. The game objective was to get the mouse to enter at the rightmost position before it quit either because of extinction or because of satiation.

The game is illustrated in Figure 29-3 with samples from early through late stages of two games. The mouse is shown having entered at various locations along the bottom. The graphs show the current probability distributions, where the heights of the filled bars are proportional to probabilities. The mouse is most likely to enter where the bars are highest and least likely where they are lowest. It never enters where there are no bars: at those locations, response probability is zero.

We learned some lessons from the design of this simulation. Most important, we couldn't get it to work unless reinforced responses raised probabilities and unreinforced responses also reduced them. Both reinforcement and extinction are essential, and of course that is what is implied by differential reinforcement. The system has some properties in common with Skinner's (1938) concept of the reserve, in which reinforced responses add to the reserve whereas unreinforced responses deplete it (cf. Addendum 15B). It was of course necessary for the effects of reinforced or unreinforced responses to spread to neighboring locations; the broader the spread, the more quickly shaping could progress. The simulation also showed that shaping was most efficient, in the sense of moving the distribution farthest with the fewest reinforcers, when response probabilities were kept low, as in Game A, but that made extinction a risk. Reinforcing many responses in one region, as in Game B, makes probabilities in that region high, but responding there must eventually be extinguished to keep the distributions moving in the right direction.

These features of shaping have a phylogenic parallel, in that rapid evolution is produced by relatively extreme changes in the environment (cf. Chapter 9 and Kirschner & Gerhart, 2005, p. 74). The lessons we've been taught by this simulation have been consistent with what happens in shaping the behavior of real rather than virtual organisms. To that extent, the simulation has been useful in helping us to see how shaping works.

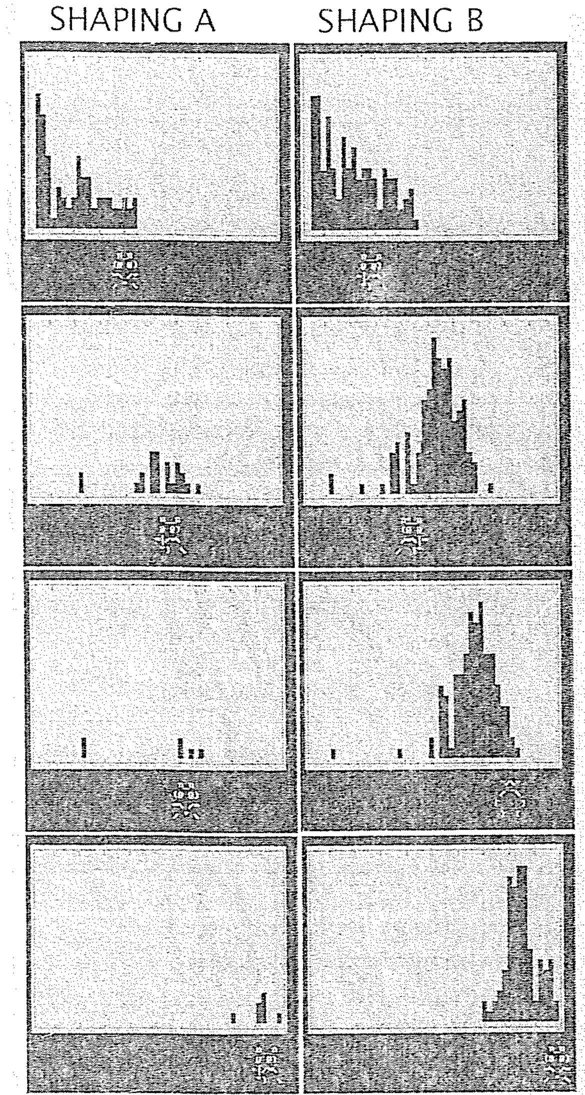

Figure 29-3 Sample probability distributions for early through late stages of two plays of a shaping-game simulation. Bars show probability distributions of the different locations of a mouse's entry at the bottom. Reinforcing a response raises probabilities at and around the entry point; not reinforcing decreases them. By differentially reinforcing only responses from the right region of any distribution, the player shifts the distributions in that direction. The player wins by getting the mouse to the far right before the game ends in extinction or satiation. Game A shows efficient shaping, but probabilities that low risked extinction, which occurs when no bars remain. Game B built up high probabilities at each location as shaping progressed, but delivering so many reinforcers slowed things down and risked satiation.

Another simulation, also sometimes called a model, extended the reserve from the response that produced the reinforcer to the earlier responses that preceded it, weighted by delay of reinforcement (Catania, 2005d; cf. Addendum 15B). In this simulation responses are produced with probabilities linearly related to the size of the reserve, so the bigger the reserve the higher the rate of responding, each response leading up to a reinforcer contributes to the reserve but weighted by the delay between it and the reinforcer, and each response reduces the reserve by a constant amount. Because the effects of delay decrease with increasing time to the reinforcer whereas the effects of response emission are constant, these functions cross at an equilibrium point. Thus, responses closer to the reinforcer than this point add more to the reserve than they take away and those further away take away more than they add.

These relations interact with different schedules in different ways. A DRL schedule keeps all but the final response from getting close to the reinforcer; this happens without regard to whether the organism engages in any timing behavior. An RI schedule keeps the interreinforcer times constant, so if response rates increase they do so uniformly on both sides of the equilibrium point. In an RR schedule, on the other hand, increasing response rates move more responses to the side of the equilibrium point closer to the reinforcer, thereby reducing delays still further and therefore increasing response rates further. By calculating the moment-to-moment size of the reserve over successive reinforcers arranged according to various schedules contingencies, the simulation can generate cumulative records. Figure 29-4 shows some samples, each of which is reasonably consistent with the records produced by real pigeons under comparable contingencies. It was particularly satisfying to see that in this simulation responding is not maintained by free reinforcers and that response rates can quickly change from the moderate rates typically maintained by RI schedules to the high rates maintained by RR schedules.

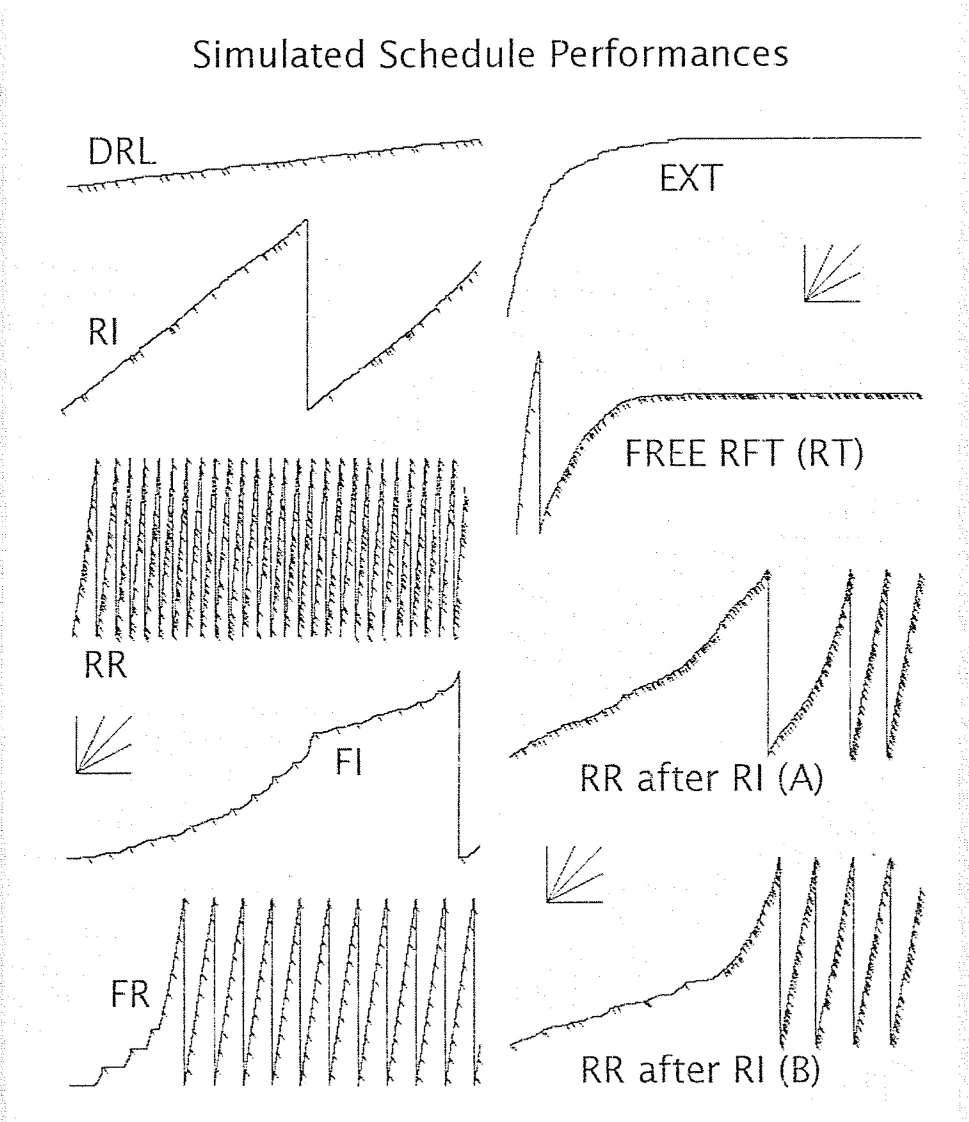

Figure 29-4 Cumulative records for various schedules produced by a computer simulation in which each reinforcer augments a reserve based on its relation not only to the most recent response but also to later ones, weighted by a delay gradient, and in which each emitted response depletes that reserve. (Adapted from figures in Catania, 2005d)

PART V
CONCLUSION

Chapter 30
Structure and Function in Learning

The Indo-European root leu-, *to loosen or divide, is the source of* loss *and the suffix* -less. *In combination with* an-, *up or again, it provides* analysis, *a taking apart. In combination with* se-, *apart, it also provides* solve *and* solution. *The Greek roots,* syn-, *together, plus* tithenai, *to put, provide the origins of* synthesis, *a putting together. We take complex contingencies apart in a behavior analysis, and we demonstrate the adequacy of our taxonomy by putting them back together again in a behavior synthesis.*

Two Psychologies of Learning: A Capsule History
Structure and Function
Learning and Evolution
Behavior Analysis and Behavior Synthesis

We've considered the varied phenomena of learning. We began with experimental operations or procedures. The first and simplest, apart from observation, was presenting stimuli. Stimuli produce responses, but they can also have other effects. For example, they can modulate how responses are distributed in time. Just as responses can follow stimuli, stimuli can follow responses. We dealt with the effects of contingencies, relations between behavior and the environment, in treating the operations of reinforcement and punishment. Their effects led us to distinguish among positive reinforcement, positive punishment, negative reinforcement (escape and avoidance) and negative punishment (timeout and omission procedures). In exploring these procedures, we noted the importance of some consequences, less dramatic than biologically significant stimuli such as food and water, that affected an organism's orientation within its environments. In this context, we examined sensory-motor learning, visual fixation and other phenomena not usually classified according to the effects of reinforcers and punishers.

We distinguished between reinforcement and punishment by whether response consequences raised or lowered response probability, but we

KEY TERMS: Structure, Function, Anatomy, Physiology; Behavior Analysis, Behavior Synthesis; Methodological Behaviorism, Radical Behaviorism; Philosophical Origins, Biological Origins.

also recognized that distinguishing between the positive and negative cases was sometimes arbitrary. As in reinforcement by a change in temperature, it was sometimes ambiguous whether a case should be treated in terms of presenting or removing stimuli (e.g., presenting heat or removing cold). Instead, we treated these cases in terms of relations between reinforced responses and the responses occasioned by their consequences. We noted that stimuli could have multiple effects, and that a stimulus that served as a reinforcer for a response might also have eliciting effects on that or other responses. The task of an experimental analysis is to separate such different effects and to observe how elicitation and consequences can combine to create a behavioral outcome.

We also turned to the signaling effects of stimuli, in stimulus control. This operation can be superimposed either on the consequential operations of reinforcement and punishment, when we speak of discrimination learning, or on stimulus-presentation operations, when we speak of respondent conditioning. Within these contexts, we explored a variety of behavioral phenomena and procedures: attention, stimulus-control gradients, reinforcement schedules, novel behavior, self-control, sensory preconditioning and conditioned suppression, to mention just a few. Along the way, we considered their various applications to significant human problems, such as parenting, education and the nurturing of the developmentally disabled.

Having examined the effects of these operations, we shifted our attention to the problem of characterizing the dimensions along which stimuli and responses vary. In discussing differentiation and discrimination, we saw that the relevant dimensions included not only relatively simple ones such as topography or intensity, but also complex ones involving the structure of stimuli and responses and their relations. The correspondences between the classes of responses with particular consequences (descriptive or nominal operants) and the classes of responses generated by those consequences (functional operants) was critical to defining behavioral classes; we distinguished between those defined by response properties, called operants, and those defined by stimulus properties, called discriminated operants.

In our examinations of natural categories, concept formation and verbal behavior, we were forced to conclude that behavioral relations were primary: We couldn't appeal to physical measures to define such classes. Responses like attending, remembering, imagining and thinking aren't easily observed directly, but we decided early that it would be inappropriate to define behavior in terms of movement. Furthermore, in our analyses of symbolic matching we discovered relations among stimuli and responses that weren't implicit in the three-term contingencies among antecedents, behavior and consequences. The relations of reflexivity, symmetry and transitivity can be demonstrated only within procedures that allow stimulus and response terms to be exchanged. They led us to define symbolic behavior in terms of equivalence classes, and those classes later entered into our treatment of verbal behavior.

Much of our account classified learning phenomena according to experimental operations or procedures. But there is no guarantee that imposing a particular procedure on a given organism's behavior will be effective. A procedure in which one organism learns may be one in which another doesn't, and an organism that learns under one procedure may not do so under another. Such findings don't invalidate our behavioral taxonomy, because the classifications are merely ways of naming phenomena and relating them to each other. That wouldn't be so if we assumed that all learning was based on some one process or some small number of processes that act across all organisms and procedures.

When we moved from learning without words to learning with words, by way of social learning, we discovered that our basic operations were relevant to verbal as well as nonverbal behavior. Verbal classes such as tacting and manding were, in effect, names for the relevant processes as they entered into verbal behavior. They also provided the context within which we more fully developed the implications of higher-order classes and recognized potentials for novel behavior in adduction and other sources of emergent behavior. The taxonomy of verbal behavior then allowed us to treat issues in a range of areas, including remembering, cognition, problem solving and metaphor, among others. Speaking metaphorically, we've traveled a long way.

We have typically preferred the concrete to the abstract. Instead of starting with assumptions about mind and personality, we have here built up our accounts of the processes relevant to such entities from much simpler parts, much as biologists have built their accounts of complex organisms from their analyses of cells and their components. We did not need to invent new processes to fill in gaps. By proceeding in that way we kept our approach parsimonious. Our frequent appeals to the details may sometimes have seemed constraining, but environments are properties of the world about which we can often do something. We need not be weighed down by prejudices about human behavior that have evolved over many centuries. Don't think of our behavior analytic accounts in terms of the role of the environment as imposing limits. Instead, think about how liberating they are. They keep us from being trapped by unconstrained theorizing and unwarranted assumptions about the sources of behavior (Skinner, 1961).

TWO PSYCHOLOGIES OF LEARNING: A CAPSULE HISTORY

The psychology of learning evolved on the basis of the different outcomes of a variety of experimental procedures. As each procedure was considered in its historical turn, it was assigned importance in proportion to its demonstrated effectiveness. At different times, the available findings led to theoretical formulations dominated by laws of association or contiguity, rules of respondent conditioning, or principles of reinforcement and punishment. Sometimes this domination was so substantial that one or another process was presumed to be the fundamental and exclusive basis of all learning. Such formulations were inevitably open to challenge, because the phenomena of learning can't be accounted for exhaustively by any single process. Thus, the history of the psychology of learning, a tale of confusions and controversies, has been more often told in terms of theorists and their systems than in terms of the phenomena of learning.

The psychology of learning has at times been regarded as a foundational component of experimental psychology, but in the psychological laboratories of the late 19th century it was overshadowed by other issues, such as those of sensation and perception. By the turn of the century, research on animal behavior had been stimulated by the work of Darwin (1859), and the fundamentals of research on human memory and verbal learning had been provided by Ebbinghaus (1885/1964). Both historical lines found homes within university laboratories, but despite their academic proximity they remained separate.

Early in the 20th century, John B. Watson (1913) addressed what came to be called his behaviorist manifesto to a psychology marked by substantial disagreements about its methods and its subject matter (Catania, 1993b). Researchers claimed to be able to study the content of consciousness through introspection, but couldn't agree on such fundamentals as the nature of the basic mental units. In this context, Watson advocated behavior, as opposed to consciousness or mind, as the only legitimate subject matter of psychology. On methodological grounds he excluded mental events such as images as proper areas of inquiry, and his version of behaviorism came to be called *methodological* behaviorism. In other words, Watson rejected the study of private events. The study of private events would require a different kind of behaviorism, sometimes called *radical* behaviorism, perhaps because in its interest in the behavior of the scientist it had turned its own science upon itself (Catania, 1993a; Skinner, 1957). Another variety of behaviorism that has broken some of the constraints of methological behaviorism is the *paradigmatic* behaviorism of Staats (Burns & Staats, 1991; Staats, 1986).

In the study of animal behavior, much early research was devoted to invertebrates (e.g., Jennings, 1906; Loeb, 1900; Lubbock, 1882), but attention gradually shifted to vertebrates. Instrumental learning had been introduced by the mazes of Small (1899-1900) and Yerkes (1907) and by the problem boxes of Thorndike (1898). Thorndike soon moved from animal studies to analyses of human learning (e.g., Thorndike, 1921), but in so doing he was exceptional. Other students of animal learning were more likely to make the extension from animal to human behavior through theoretical statements than through experiments. Pavlov (1927, 1957, p. 285), for example, manifested an interest in human language in his theory of the second signal system, but his main impact was through his canine research. Once instrumental learning and respondent conditioning had been distinguished as phenomena in the early decades of the 20th century, the stage was set for elaborations of discrimination learning, in the progression from jumping stands to rat chambers and pigeon boxes (Lashley, 1930; Skinner, 1930, 1938). The phenomena of reinforcement schedules were to follow later (Ferster & Skinner, 1957).

Preoccupation with the scientific legitimacy of psychology disposed researchers to look to other sciences for principles of scientific method, and gradually the behaviorism of Watson (1919) converged with the operationism that was newly developing in physics (Bridgman, 1927) and with the logical positivism that was being introduced as a revolutionary change in the philosophy of science (Ayer, 1946). Behaviorism became a major orientation within psychology, and the 1930s and 1940s became a period of grand behavioral theories. Applications to human behavior and to language became an issue in the competition among the systems of Guthrie (1935), Skinner (1938), Hull (1943) and Tolman (1948), among others, but attempts to integrate the processes of animal learning with the properties of language gradually lost influence as the fortunes of early behaviorist theories based on stimulus-response associations declined. Even so, methodological behaviorism has remained the foundation for virtually all of contemporary experimental psychology in such practices as the operational definition.

Meanwhile, the field of human learning and memory went its separate way (Catania, 1985). The precedence of Ebbinghaus dominated detailed studies of serial-position effects, massed versus spaced practice, meaningfulness, backward associations, interference and so on. Some controversies were long-lasting. For example, plateaus in Morse-code learning at various stages of competence, demonstrated at the turn of the century (Bryan & Harter, 1899), remained a part of psychological lore until well into the 1950s, when the phantom phenomenon was at last put to rest (Keller, 1958). Thorndike and Woodworth (1901), inspired perhaps by questions about the educational value of such classical disciplines as Greek and Latin, had begun investigations of transfer of learning. Findings from these and other studies made the practical relevance of studies of human learning seem obvious, and the effects on verbal learning of such variables as sleep and distraction and motivation became the basis for advice on study habits. The theoretical underpinnings of these areas were the same as those of animal learning, but although laws of effect, contiguity, association and generalization gave a superficial appearance of unity, the two experimental lines originated as separate entities and remained so.

By the 1940s, the pattern was firmly entrenched and was most clearly illustrated by two textbooks of the time. *Conditioning and learning* by Hilgard and Marquis (1940) was devoted primarily to animal research; *The psychology of human learning* by McGeoch (1942) was concerned mainly with human learning and remembering. Each appeared later in revised editions (Kimble, 1961; McGeoch & Irion, 1952). Beyond the common appeal to theoretical principles such as association and some lip service about likely mutual relevance, there was little evidence in either the original or the revised editions that the two research traditions had substantially influenced each other.

It can reasonably be claimed that these two psychologies of learning are still separate in contemporary psychology. Having examined their histories we might ask the point of treating these disparate subject matters together. Perhaps animal learning and conditioning and human learning and remembering have so little relevance to each other that they should go their separate ways. But here we've argued otherwise. Such a course would be a serious mistake for several reasons: Learning phenomena studied with animals also occur in human behavior; the nature of complex human learning is clarified by analyses in terms of more elementary processes; and perhaps most important, human behavior especially is characterized by the interplay between verbal behavior and the nonverbal behavior on which it depends, and we have yet to fully understand the origins of either.

We must identify not just the properties of our behavior unique to humans but also those we hold in common with other organisms. Given the myriad properties of human verbal behavior, both elementary phenomena of animal learning and conditioning and the complexities of human learning and remembering must enter into any effective account. Our treatments of such phenomena as equivalence classes, awareness, instructional control and problem solving couldn't have proceeded without prior analyses of antecedent stimuli. To the extent that such important varieties of human behavior as autoclitics, self-control, and metamemory are based on both verbal behavior and discriminations of our own behavior, any analysis of them that fails to build upon the more elementary processes of stimulus control and contingencies is bound to be deficient. We must therefore conclude that the union of these two psychologies of learning is long overdue; we need a single psychology of learning encompassing all of the varieties of animal and human learning.

STRUCTURE AND FUNCTION

Early in our account we distinguished between structural and functional analyses (cf. Catania, 1973b; Titchener, 1898). A structural analysis holds relations among stimuli and responses constant while varying critical properties of one or the other. We noted the properties of formal verbal units (as in phonemes and letters), hierarchical organizations in text structure (as in relations among letters, words, phrases and sentences) and correspondences between stimulus and response structure (as in relations between text and speech). A functional analysis holds the stimuli and responses of interest constant while changing their relations. We study interactions between behavior and environment in terms of contingency relations among discriminative stimuli, responses and consequences, and when three-term contingencies are nested within other contingencies we extend our analyses to four terms or whatever larger number of terms the environment may call for.

Structural and functional problems are often interrelated. For example, the design of an illustrated reader for an elementary school class should consider both its structural features, in the details of text and pictures, and its functional features, in how the pictures are related to the sentences they illustrate. A problem exists if the pictures set the occasion for verbal descriptions that correspond so closely to the accompanying text that a teacher can't tell whether a child is responding to the picture, reading the text, or doing both in some combination. It can be addressed by changing the relation between the pictures and the child's responses, as in making the pictures reinforcing consequences of reading rather than antecedents of it. A systematic analysis might show that, as consequences, pictures maintain the picking up of books and the turning of pages, but as antecedents they compete with the text in setting the occasion for verbal responses. This analysis might lead to a rationale for sequencing the relations among words and pictures to maximize both stimulus control by the words and reinforcing effects of the pictures. At the same time, a structural analysis of text and pictures could provide the basis for organizing the textual material, deciding on appropriate levels of complexity, and perhaps even determining how quickly pictures could be faded out (e.g., Harzem, Lee, & Miles, 1976).

For far too long, educational systems have focused on what the teacher does instead of what the student does. What the student does is what the student learns. The student is not a vessel into which knowledge can be poured. It matters little whether a teacher changes from lectures to presentation software or from large class sections to small if no contingencies are placed on the student's behavior. Too many proposed educational innovations involve little more than changing the curriculum. But changes in what a student is supposed to learn or the order in which the student is supposed to learn it are largely irrelevant to changing how the student is taught. The student must behave with respect to what is being learned.

Educators must attend to the behavior of the student rather than to the behavior of the teacher. There is no evidence for significant effects of individual learning styles (Pashler *et al.*, 2008). Students should be able to learn in many different ways: by rote, by understanding, by discovery, and so on. The more time a student spends with the material to be learned, the more the student will learn, so anything a teacher can do to get the student to put additional time in will probably be a good thing, provided the contingencies do not become aversive (school absenteeism is a measure of whether the school environment provides reinforcers).

But changing curricula will not accomplish this if teachers present materials to students but fail to arrange contingencies for the successful learning of their students. Effective methods for teaching based on behavior analytic principles have been around for decades (e.g., Binder, 1996; Greer & McDonough, 1999; Greer & Ross, 2004; Johnson & Layng, 1992; Lindsley, 1992; Potts, Eshleman, & Cooper, 1993; Skinner, 1968). It is shameful that in contemporary western culture they are so little recognized, much less implemented. Teachers must attend to the structure of what they teach, but they must also know the functions of their interactions with their students. Above all, they must attend to their students' behavior. In the best of educational worlds the shaping should be reciprocal.

Structure and function are complementary, not mutually exclusive. There is no reason why structural concerns shouldn't enter into functional experiments, or vice versa. But in the evolution of psychology the distinction between structural and functional approaches became correlated with and eventually confused with another distinction, between the two languages of cognitive psychology and behaviorism. The cognitivist prefers to summarize the organization of the organism's behavior in terms of structures the organism knows, for which the language of mind is convenient. The behaviorist prefers to anchor accounts of action to the detailed functional relations among observable events, for which the language of stimuli and responses is appropriate. Yet relations between stimuli and responses must be implicit in the cognitive vocabulary, as it deals with relations between the structures of the environment and knowledge, just as structure is implicit in the behavioral vocabulary, as it deals with the properties that define operants and other classes.

Biology also distinguished between structure and function, in anatomy and physiology, but an equivalent schism didn't emerge within biology because the languages of anatomy and physiology didn't diverge. For example, the debate over mechanism versus vitalism, perhaps analogous to that between behaviorism and mentalism, remained reasonably independent of the distinction between anatomy and physiology (e.g., Catania, 1978; Hein, 1972). In psychology, the respective correlations of cognitive and behavioral languages with structural and functional problems made it difficult to recognize that the problems were different and therefore that these research areas might be complementary rather than mutually exclusive.

There is always a certain ambiguity in the distinction between structure and function. But when the biologists of an earlier time debated the issues, their concerns weren't so much with whether structural or anatomical problems could be distinguished from functional or physiological ones as they were with whether one or the other problem should be given priority (e.g., Russell, 1916). The emergence of structure from undifferentiated beginnings has long been a fundamental problem in science. For example, debates over the development of the embryo pitted the unfolding of preexisting structure (preformationism) against the functional differentiation of unstructured systems (epigenesis), and gave rise to the widely cited but usually misleading generalization that ontogeny recapitulates phylogeny (Gould, 1977).

The arguments were based on assumptions about how, in the evolution of a species, the functions of an organ might determine its structure or how, in the development of an organism, the structure of an organ might determine its functions. With the advent of Darwin's (1859) account of evolutionary process in terms of selection, both structure and function in biology came to be seen as derivatives of selection, each reciprocally constrained by the other. For example, the properties of animal locomotion, whether on land or at sea or in air, are jointly determined by common functions (e.g., capturing prey, escaping predators) and by structural limitations that arise from differences in vertebrate and invertebrate body plans (e.g., muscle configuration, number of limbs). The problem was resolved not with the domination of one or the other position, but rather with the recognition that structure and function are mutually determined by selection.

Biology has evolved from a theoretical discipline to a descriptive one. The double helical structure of DNA was once a theory of sorts, but once Watson and Crick (1953) had constructed their model and showed how it worked the detailed confirming investigations followed soon after. Biologists and their colleagues in related disciplines have created a taxonomy of fundamental biochemical and other processes. That taxonomy provides the parts, and to demonstrate how those parts fit together and interact with each other is to show what a biological system does and how it does it (cf. Carroll, 2005, p. 10, 13, 295). The contemporary microbiologist does not trade in theories of cell metabolism but instead shows how cell metabolism works. "Throughout scientific history, new ways of seeing have played a critical role in the discovery and the acceptance of new ideas" (Carroll, 2006, p. 217) and "The payoff for learning about the DNA code is being able to see, and therefore so much better understand, the process of evolution at its most fundamental level" (Carroll, 2006, p. 73). Seeing how behavior works is perhaps sometimes trickier, because behavior does not sit still for us when we try to look at it. It is appropriate for us to expand our taxonomy, now that it is more respectable than once it was to speak of taxonomies for categories other than plants and animals (e.g., Ross & Spurrett, 2004, p. 604).

Controversies over the primacy of structure or function similar to these in the history of biology also exist in the history of the psychology of learning, although the issues are usually expressed differently. For example, consider latent learning: A hungry rat explores a maze with an empty goalbox; when later given food there, it demonstrates that it has learned the maze by negotiating it as rapidly and with as few blind-alley entries as a rat that has always found food there. The rat's running, itself not learning, occurs because of its relation to food in the goalbox; presumably it reflects what the rat learned. It was once important to distinguish between learning and performance, but the distinction was merely a basis for asserting the primacy of structure over function. What is structure here? It seems to be the particular sequence of turns. To say that the structure was learned, however, isn't to say that it caused the learning. A theorist who wanted to reinstate the primacy of function over structure might then argue that the structure was learned because of the contingencies: Certain turns at certain choicepoints led to certain new maze locations, and the rat learned these functional relations even without food in the maze. Even in negotiating the environment, the consequences of behavior matter. Then comes the argument that if contingencies define what is learned, they too can't be a cause of learning. And so it goes.

The theorist arguing for primacy of structure was probably a cognitivist, and the theorist arguing for primacy of function was probably a behaviorist, and their different languages weren't likely to help matters. Yet these structures and functions of behavior were both outcomes of learning. The problems can be resolved only by recognizing that behavioral structure and behavioral function are mutually determined by the relations between behavior and the environment. Both the structure and the function of behavior are to be understood in terms of their origins, and neither has primacy over the other. The relative provinces and provenances of structure and function remain controversial, but the controversies are perhaps amenable to the same sorts of solutions that worked for biology.

Both structure and function have limits, and as learning theories evolved, they were necessarily refined and qualified, thereby restricting the range of phenomena to which they could be applied. Recognizing the boundary conditions for learning was implicit in these restrictions. Limits on learning pose no problems when they can be easily traced to an organism's sensory or motor capacities. For example, we aren't surprised if certain stimuli are more likely to produce responses in some species than in others. We know that the visual and auditory systems of pigeons and bats make pigeons capable of visual discriminations impossible for bats and make bats capable of auditory discriminations impossible for pigeons. We also aren't troubled by different capacities for responding. Pigeons and bats fly differently at least in part because the anatomical structures of their wings are different. Although the examples are obvious, they aren't trivial. They illustrate how much we take the different sensory and motor capacities of different species for granted.

Like sensory and motor capacities, the effectiveness of different stimuli as reinforcers or punishers varies across species. Once we recognized the relativity of reinforcement, it was no longer good enough just to identify which reinforcers or which punishers might be effective for a given species. These consequences can only be defined in relation to the responses that produce them. Within a species, a reinforcer effective for one response may not be effective for others. As demonstrated by phenomena such as food aversion and acknowledged by the concept of preparedness, we can't specify limits on learning in terms of stimuli alone or responses alone. Instead, we must express the limits in terms of the types of relations that can be created in a given species. Organisms may be predisposed to learn different relations among stimuli and responses in different situations. These predispositions are limitations on the structure of behavior.

An organism's behavior in its current environment is determined by phylogeny as well as ontogeny. The relative contributions of nature and nurture have been a long-standing issue in psychology, and though the emphasis has often shifted in one direction or the other it remains clear that neither operates to the exclusion of the other. Some aspects of behavior are highly determined by evolutionary factors (e.g., the human vocal apparatus) and others by experience (e.g., the particular human language one speaks). Nature and nurture are extremes on a continuum, and we must therefore recognize that learning as well as evolutionary history can impose constraints on behavior.

Like nature and nurture, behavior too is best represented not by all-or-none categories but rather by dimensions along which processes can be located. Reinforcement and punishment are extremes on a continuum of contingencies ranging from those that increase responding through those that don't affect responding to those that decrease responding. Differentiation and discrimination are extremes on a continuum of the relative contributions of response properties and stimulus properties to the criteria for differential reinforcement. Contingency-shaped and verbally governed behavior are extremes on a continuum of the relative contributions of contingencies and verbal antecedents to the behavior of a listener. Behavioral and cognitive processes are extremes on a continuum that represents the relative accessibility or inaccessibility of what organisms do. And so on.

One crucial part of the distinction between behavioral and cognitive psychologies is the insistence by the former that its subject matter is behavior. Our interpretation of cognitive processes such as remembering and imagining, for example, has been in terms of what organisms do. When some type of cognitive process, such as the processing of information, is expressed in terms that aren't explicitly related to behavior, its status becomes similar to that of the tacting of private events (cf. Watkins, 1981). A consistent vocabulary can be developed for a cognitive process only if at some point it makes contact with the environment, just as the consistent tacting of a private event can be developed only if some correlate of that event is publicly available to the verbal community. That may be one reason why the status of representations as copies or as transformations of stimuli has been the basis for so much controversy.

Another source of controversy has been the relation between behavior and physiology and the nervous system in particular. Asserting that behavior is a subject matter in its own right doesn't deny its intimate dependence on physiology. Consider again the evolutionary analogy. Natural selection overcame challenges from orthogenesis and other alternatives and emerged as the primary account of evolution long before molecular biology worked out mechanisms of genetic transmission. Early geneticists had no biochemical evidence about genes and based their conclusions only on the data of reproduction. Molecular biologists needed the findings of genetics and mutation and natural selection to know what to look for in cells. Evolutionary biologists needed the findings of embryology and biochemistry and paleobiology to know what to look for in the interactions between genes and developing embryos and their environments. They'd have looked differently, and undoubtedly with less success, if they'd started with orthogenetic or Lamarckian assumptions.

Similarly, those studying the neurophysiology of learning need to know what happens in learning and behavior to know what they should look for in the nervous system. The neurophysiologist who thinks of learning mainly in terms of stimulus-response associations will look for different things than the one who thinks of learning in terms of ontogenic selection. Skinner compared the relation between behavior analysis and neuroscience to that between genetics and biochemistry: "It is the function of the science of behavior at the present time to give neurologists their assignments, as it was the function of genetics prior to the discovery of DNA to give modern geneticists their assignment with respect to the gene" (Skinner, 1988, p. 60; see also p. 461).

We've made progress in understanding the neurophysiology of some relatively simple systems (e.g., Carew, 1992; Kandel & Schwartz, 1982), but neuroscience has given only scant attention to mechanisms for ontogenic selection (Stein, Xue, & Belluzzi, 1993). Neural mechanisms presumably underlie the environmental contingencies that create behavioral structure (cf. Hebb, 1949). But where Hebb saw cell assemblies, we might see the processes that produce boundaries and compartments as the crucial dimensions of development, where operant classes are among the modules through which behavior becomes organized (cf. K. C. Catania & Kaas, 1996).

The accumulating evidence for growth and reorganization within the nervous system is consistent with a selectionist account (e.g., Donahoe, Burgos, & Palmer, 1993; Donahoe & Palmer, 1994; Edelman, 1987; Kaas, 1991; Recanzone *et al.*, 1992; Shull, 1995; Yates, 1986). Parallels with biology have often provided hints for useful directions in the analysis of behavior. That should be no surprise, because after all is said and done behavior remains a central component of biological systems. Despite origins in philosophy and psychology, we are a part of biology and must always cultivate our intellectual ties with the biological sciences. If we do so, we can anticipate continuing insights from ongoing developments in evo devo and related biological disciplines (e.g., Bolker, 2000; Carroll, 2005; Chédotal & Richards, 2010; Irvine & Rauskolb, 2001; Kiecker & Lumsden, 2005, 2009; West-Eberhart, 2003); attending to them will serve us well.

LEARNING AND EVOLUTION

In discussing types of selection, we've argued that properties of learning parallel those of evolution, because the selection or survival of patterns of behavior in an organism's lifetime has much in common with the selection or survival of individuals in the evolution of a species (e.g., Catania, 1978; Skinner, 1966, 1975). Our discussion of social learning made a similar case for the selection of cultural practices (cf. Catania, 2001b; Dawkins, 1976; Harris, 1977; Petroski, 1992; Skinner, 1981). Our own verbal behavior makes this variety of selection of special interest to us. Each type of selection involves some kind of variation that provides the source materials upon which it operates, and each involves some mechanism for selecting what survives. Whatever else happens at every level, behavior gets selected and the environment does the selecting.

The parallels between Darwinian natural selection and operant selection also extend to the problems of acceptance each has faced (Catania, 1987). Like Galileo's displacement of the earth from the center of the universe to an orbit around the sun and like Freud's challenge to the central status of human consciousness, these accounts overturned traditional ways of thinking about the place of our species in nature. We should note, by the way, that Freud's message at least took consciousness for granted while reducing its scope, in the sense that his unconscious processes, in the interactions of ego and superego and id, were supplements to consciousness (Freud, 1917). In a behavioral account, however, consciousness itself is derivative, because it demands discrimination of one's own behavior, which has to exist already if it is to be discriminated. In other words, the unconscious has to be there first or there'll be nothing to be conscious of.

In any case, phylogenic and ontogenic selection have faced similar substantive challenges. For example, artificial selection was familiar in Darwin's time; what was questioned was whether selection could operate naturally. The operant parallel is provided by shaping, which is also an artificial selection procedure, as when an experimenter shapes a pigeon's figure-8 turns or as when a behavior therapist shapes the vocalizations of a nonverbal child on the autism spectrum. The effectiveness of shaping is self-evident; what is questioned is whether it operates naturally to produce some of the varied patterns of behavior that we see in everyday life.

It isn't good enough to argue that in humans the effects of shaping are likely to be often masked by verbally governed behavior. It would be best to document cases in which the changes in contingencies are identified early and tracked. Typically, however, we only have outcomes, after the natural contingencies have done their work. For example, we can assume that ontogenic selection was involved in shaping the skill with which grizzly bears catch salmon in the rivers of the Pacific Northwest, but we mainly see the differences between the inefficient performances of the young novices and the well-coordinated actions of the experienced adults; we don't see the shaping itself, because it continues over too long a time.

Furthermore, shaping can be hard to see if one doesn't know what to look for. Someone who has actually done shaping is more apt to notice it when it happens naturally than someone who has only read about it. Thus, the parents who always wait a while before attending to a crying child may not notice that they've gradually shaped louder and more annoying cries. The attention reinforces the crying, and annoying cries are, by definition, the ones most likely to get attention. If you watch what a parent does when a child throws tantrums, it is usually easy to guess where the tantrums came from.

Time is another factor in the acceptance of both types of selection. For Darwinian natural selection, the question was once whether the earth had existed long enough for such selection to have taken place; upward revisions of the age of the earth resolved the problem. The comparable problem is easier to deal with for operant selection. Even with rapid breeders like the fruit fly, genetic experiments take days. Shaping, however, can be demonstrated within minutes. If reinforcers can do so much to behavior when contingencies are deliberately arranged over relatively short periods of time, isn't it reasonable to assume that they will also affect behavior when natural contingencies operate over substantial periods in an organism's lifetime? Many contingencies may take hold of behavior over the course of a year in the life of a young child. Compared to how long most artificial examples of shaping last, a year is an extremely long time. Some contingencies may be subtle, especially when we recognize the very broad range of events that can serve as reinforcers. Some may produce behavior that is desirable; others may do the opposite. Given what artificial contingencies can do in a short time, natural contingencies should be able to do a lot in a long time. It is certainly more appropriate to be alert for the effects of such contingencies than to assume that they don't exist.

In both natural and artificial environments, it is difficult to determine the boundaries of behavioral classes. Here again is a parallel between ontogenic and phylogenic selection. In each case, we must deal not with particular instances but rather with populations or classes of events. We speak of populations of organisms as species, and we speak of behavioral populations as response classes, such as operants, discriminated operants and respondents (and in social learning we speak of populations of socially maintained response classes as cultural practices and as classes of verbal behavior).

It might be argued that our notions of response classes are much vaguer than the classes Darwin spoke of as species in his account of evolution. Yet even though the word *species* is in Darwin's famous title, *On the Origin of Species*, Darwin knew the term could not be defined unambiguously. In his book, he commented often on the problem of definition:

> I look at the term species, as one arbitrarily given for the sake of convenience to a set of individuals closely resembling each other [p. 52]; ...the amount of difference necessary to give two forms the rank of species is quite indefinite [p. 59]; [and]...we shall have to treat species in the same manner as those naturalists treat genera, who admit that genera are merely artificial combinations made for convenience. This may not be a cheering prospect; but we shall at least be freed from the vain search for the undiscovered and undiscoverable essence of the term species. (Darwin, 1859, p. 485)

In Darwin's account of evolution, the relations among populations of organisms could not be adequately expressed in terms of similar topographies (for example, males and females within some species differ more from each other in form than some pairs of organisms within completely unrelated species; the social insects in particular provide striking instances).

For Darwin, the important basis for distinguishing among populations was descent. We define relations among populations by looking at where they came from. Darwin's achievement, in fact, was more description than explanation. His account of evolution didn't depend on any theory specifying the mechanics of evolution. His work predated genetics, and he even argued at times that acquired traits might be the source of variation on which evolution acted. He had described the properties of evolution (Gould, 1975, p. 824). Darwin himself was skeptical about explanation: "It is so easy...to think that we give an explanation when we only restate a fact" (Darwin, 1859, pp. 481-482).

What does this have to do with learning? The analogies between behavior and biology suggest that some solutions appropriate to biology will be appropriate to behavior analysis. We've seen the importance of descriptions of what happens in learning and gave such descriptions priority here over theories or mechanisms or models. For example, we found that *reinforcement*, a term that once served explanatory functions, now functions just as a name for a behavioral phenomenon. We no longer ask whether it explains behavior. Instead, we regard questions about the generality of reinforcement as about the range of circumstances to which the name can be applied. Similarly, the study of memory doesn't explain what's remembered; instead, what's remembered is the basis for determining behavioral structure, as in the analysis of syntactic or semantic classes. Theories and models come and go, but the descriptions of behavior remain; it is therefore fitting that we've sometimes considered the results of experiments conducted more than a century ago.

Darwin's treatment rejected the concept of immutable species in favor of classes defined by their descent or evolution. The psychology of learning has sometimes moved in the opposite direction. It sought to *explain* learning by inventing sources for responses (neural traces, associations, cognitive structures). But that was getting it backwards. We should use the development of behavior to define behavioral classes. We should define behavioral classes in terms of their descent: where they come from or how they are learned. To some extent, we already do that, though seldom explicitly. We distinguish innate behavior from behavior acquired through experience. We speak of behavior produced by stimuli as elicited and we speak of behavior occurring independently of stimuli as emitted. We say that responses elicited by stimuli are respondent behavior and responses occasioned by stimuli that signal consequences are operant behavior. We say that responses engendered by verbal antecedents are verbally governed and responses engendered by consequences are contingency-shaped. Such distinctions constitute our behavioral taxonomy.

In the analysis of behavior, we deal with populations of responses. These populations are sometimes not well-defined. The problem isn't different from Darwin's. Darwin clearly recognized the arbitrary nature of the concept of species, but unambiguous definitions of species were no more critical to his account than unambiguous definitions of stimulus or response classes are to a behavioral account (like natural categories, these too are probabilistic classes). When we distinguish among words by the circumstances in which they are uttered (e.g., *fire* as a mand, a tact, an echoic or a textual response), we are simply distinguishing among classes of verbal responses on the basis of their origins. Ambiguous cases will necessarily occur because, just as organisms have many ancestors, responses have many origins. But if that weren't so there would be no need for analysis.

The relations among behavior and its consequences in operant contingencies may seem simple, but they have subtle properties some of which reveal themselves only in certain contexts. For example, when side effects are not taken into account, contingencies can appear to be ineffective. Side effects of operant contingencies may affect their acceptance, because they allow the effects of contingencies to be masked in various ways. It is therefore prudent to consider how the properties of operant contingencies may mislead us as we deploy them and evaluate their effects. In the interests of preventing misconceptions and misunderstandings, it is probably even more important to remind ourselves of them whenever we present what we know about operant contingencies to those outside of behavior analysis (cf. Schneider, 2012). To those who argue that these contingencies should not be studied because they can be misused, the appropriate rejoinder is that detailed familiarity with their properties may be the best defense against their misuse. Alone or in combination, the factors considered here may sometimes give the appearance that operant contingencies don't work. On examination, we might instead conclude that they work more ubiquitously and more profoundly than we had originally imagined.

Selection operates on behavior analysis itself. Will these ways of dealing with behavior survive in our culture? Behavior analysts must be attentive not only to the maintenance of their practices in the laboratory and in applied settings but also to the nurturing of continuing generations of students. Contingencies of cultural selection operate on our science as on other human activities. Just as we distinguish verbal behavior as the behavior of speakers and listeners from language as the practices of the verbal community, we must distinguish the social and cultural behavior of individuals from the practices of the societies and cultures. Shaping the behavior of a parent or a child is a far cry from changing practices at the level of cultures or nations.

We might do well to see whether our analyses of social contingencies can effectively be extended to human behavior on a larger scale. Can the analysis of verbal behavior contribute to our understanding of the rapid cultural changes that led to such moral revolutions as the demise of dueling and the overthrow of slavery and still ongoing changes in the role of women (Appiah, 2010)? Can the analysis of conditional discriminations contribute to our understanding of the social contingencies that help to maintain institutions that foster tolerance (Meyer & Brysac, 2012)? Can the analysis of the side effects of reinforcement and of aversive control contribute to our understanding of how social and political precursors determine not only whether conflicts are peacefully or violently resolved, but also how one or the other mode of resolution may favor the kinds of political and cultural systems that follow (Chenoweth & Stephan, 2011)?

We have here discussed cultural selection, the selection of behavior as it is passed on from one individual to another. It operates along with phylogenic and ontogenic selection (Skinner, 1981), but selection at any one of these levels need not be consistent with selection at the other two. For example, it may not matter how valuable one way of doing things is relative to some other way; the one that is easier to pass on may spread quickly and come to dominate relative to the other even if the other would be more beneficial in the long term. A case in point is the application of techniques of reinforcement relative to those of punishment (cf. Chapter 19); sadly, the advantages of reinforcement do not make it more likely than punishment to spread through cultures (Catania, 2000).

Even as reinforcement begins to be more widely appreciated in our culture, we must not be complacent about teaching what we know about it. Despite the advantages of reinforcement, it is easier to teach the use of punishers than to teach the use of reinforcers, and reinforcement can be misunderstood or obscured by other processes in many different ways. Some people are very good at shaping even without explicit instruction, but mostly the effective use of reinforcers has to be carefully taught.

BEHAVIOR ANALYSIS AND BEHAVIOR SYNTHESIS

A behavior analysis begins with complex behavior and breaks it down into its components. These are the elements of our behavioral taxonomy. They can be combined in various ways, when it may be appropriate to speak of behavior synthesis. For example, we can synthesize some kinds of sequential performances through chaining procedures. In a more complex instance, we may combine discriminative stimuli, reinforcement schedules and delays of reinforcement so that the resulting contingencies are analogous to those when we speak of self-control. In the verbal domain, we may integrate echoic behavior, tacting and listener behavior into a higher-order class called naming. If our syntheses are successful, we may use them to clarify the properties of behavior; if they are unsuccessful, we may use them to identify components of performance that were missed or taken for granted in preceding analyses (e.g., as when we study variables affecting the likelihood of commitment responses in self-control procedures, or as when we discover, in designing an animal analogue of some human performance, that verbal behavior must have had a role we hadn't allowed for). Because many important human problems involve creating new behavior (e.g., teaching developmentally delayed children), the applications of our methods are often matters of behavior synthesis (cf. Catania & Brigham, 1978; Cooper, Heron, & Heward, 2007; Fisher, Piazza, & Roane, 2011; Lattal & Perone, 1998; Madden, 2012; Mayer, Sulzer-Azaroff, & Wallace, 2012).

The term *learning* has receded into the background in all of this; perhaps it has outlived its usefulness. New behavior has many sources. We can modify behavior hierarchies, shape new responses, construct higher-order classes, generate discriminations, form equivalence classes, solve problems and create novel behavior through adduction. Our understanding of these phenomena depends at least in part on whether we've developed a vocabulary consistent with and worthy of them. We've recognized ambiguities in our current behavioral vocabulary; we can assume it will evolve as research progresses. Yet in emphasizing behavioral procedures and outcomes, it at least adheres closely to what is done and what is observed as we expand our science and apply it. The success of behavior analysis will be measured by its survival in the behavior of those who practice it and by the effectiveness of the behavior syntheses that follow from it.

Glossary

As behavior analysis has evolved, its terminology has been progressively refined. This glossary defines some of that terminology. A set of definitions must be treated as a preliminary guide to the basic classifications and concepts in the relevant literature rather than as an inflexible set of rules. This glossary has been prepared in that spirit. A fuller and more technical glossary restricted to the experimental analysis of behavior is available in Catania (1991) The evolution of some of the terminology can be examined by comparing current entries with those in an earlier version now out of print (Catania, 1968). An expansive glossary has been developed by Verplanck (1957, 2000). Most chapters in Goldiamond & Thompson (2002) include glossaries, but they typically deal in behavior analytic terms with vocabularies outside behavior analysis rather than with those within it.

The present glossary attempts to acknowledge alternative definitions and to point out difficulties or potential ambiguities in existing usages. Nevertheless, sooner or later the reader must expect to encounter particular usages in the literature that disagree with the usages defined here. Definitions are merely words that can substitute for other words, and the substitution is sometimes just an approximation. Because the framing or mastering of a definition is primarily verbal, it cannot be counted on to produce the discriminations on which the development and evolution of that verbal behavior was based. For example, the student who has learned to define *reinforcement* may be able to offer a correct definition, but it does not follow that the student will then be able to discriminate reliably between actual instances of reinforcement and nonreinforcement in laboratory or real-world settings.

Glossaries are hardly ever exhaustive. This one is no exception. Over time old terms are modified or dropped and new ones are added. This glossary covers some major terminology of behavior analysis and the psychology of learning as it appears in this volume and in closely related literature. With some exceptions, it doesn't cover aspects of vocabulary consistent with everyday usage (e.g., technical terms defined adequately in standard dictionaries), terms likely to be encountered in specialty areas or other disciplines (e.g., drug classification in psychopharmacology), or specialized technical terms that appear only in passing in the text and that can be located via the index. For a review of other glossaries in behavior analysis, psychology and related disciplines, see Catania (1989); for a useful general dictionary of psychology, see Reber, Allen, & Reber (2009); for a general English dictionary that includes Indo-European etymologies, see *The American Heritage Dictionary* (2011) and compare Skinner (1989b).

Time is usually expressed in seconds (abbreviated s in technical contexts); the abbreviation for minutes is *min*, and for milliseconds (thousandths of a second) it is *ms*. Arbitrary quantities of time or number are indicated by t and n, which are constants unless otherwise stated. In cross-references among definitions, identified by FULL CAPITALIZATION, *cf.* (as opposed to *see*) usually refers to useful contrasts and critical distinctions among related terms rather than to synonymous usages. Most cross-references are given at the end of entries, but some are contained within entries. Citations within glossary entries are noted in the References section by *gloss* followed by the initial of the entry; those without an initial were cited in this introduction. Some matters of usage pertinent to the glossary as a whole are discussed under OPERATION.

A

Abstraction: discrimination based on a single stimulus property, independent of other properties; thus, generalization among all stimuli with that property (e.g., all red stimuli as opposed to specific red objects). Cf. CONCEPT.

Acquisition: the addition of new behavior to an organism's repertory; learning. The behavior may be a discriminated operant, a topographically complex operant, a conditional reflex relation or the performance controlled by a schedule, or new behavior produced by a wide range of other contingencies. Cf. LEARNING, REPERTORY.

Activation: when reinforcement is interpreted as made up of two components, the arousal or energizing part of the action of reinforcers, as opposed to the coupling or linking of reinforced responding to that activation (Killeen, 1994).

Adaptation: a reduction, usually during the prolonged presentation of a stimulus, in the behavior produced by that stimulus (e.g., adaptation to an experimental chamber). Cf. HABITUATION, POTENTIATION.

Adduction: the production of novel behavior when new combinations of stimulus properties that separately control different classes or properties of behavior engender new combinations of those classes or properties, as when a child appropriately combines a color name and an animal name on seeing a horse of a different color for the first time; the novel coming together of different repertories. The combinations can be either sequential (*serial coordination*) or simultaneous (*parallel coordination*). Adduction is most evident in verbal behavior.

Adjunctive behavior: responding that reliably accompanies some other response produced or occasioned by a stimulus, especially with stimuli presented according to temporally defined schedules. The stimulus rather than the responding it engenders is usually emphasized. For example, in rats, food deliveries reliably produce eating followed by drinking; the adjunctive drinking is said to be *induced* by the schedule of food delivery and not by the eating.

Adjusting (adj) schedule: a schedule varying as a function of some property of performance, as when the ratio of an adjusting FR schedule increases or decreases depending on the duration of postreinforcement pauses, or as when the parameters of an adjusting avoidance schedule change as a function of how often aversive stimuli occur.

Aggression: a side effect of presenting aversive stimuli or removing positive reinforcers. These events may generate responses that injure other organisms (e.g., biting) and/or make the opportunities for such responses effective as reinforcers.

Alternative (alt) schedule: a schedule in which a response is reinforced when either of two or more requirements is satisfied, as when, in alternative FR 10 FI 60-s, the 10th response or the first response after 60 s is reinforced, whichever comes first, and then both schedule requirements start over.

Analysis of behavior: see BEHAVIOR ANALYSIS.

Antecedent: a stimulus or event that precedes some other event or a contingency; a discriminative stimulus in a three-term contingency is one kind of antecedent.

Anxiety: see EMOTIONAL BEHAVIOR, PREAVERSIVE STIMULUS.

Appetitive stimulus: usually, a positive reinforcer, especially one the effectiveness of which is modifiable by deprivation.

Applied behavior analysis: see BEHAVIOR ANALYSIS.

Arbitrary matching: see MATCHING-TO-SAMPLE.

Arousal: a state of increased behaving or readiness for behaving, metaphorically extended from arousal in the colloquial sense of awakening.

Artificial reinforcer: see CONTRIVED REINFORCER, EXTRINSIC REINFORCER.

Artificial selection: in the Darwinian account of evolution, the variety of selection practiced by humans, in selective breeding in horticulture, animal husbandry, etc. The distinction between artificial and NATURAL SELECTION is also relevant in ontogenic selection.

Association: see CONTIGUITY.

Attention (attending): looking at, orienting toward or otherwise responding to the functional features of a current discriminative stimulus. Such responding may sometimes be made explicit as observing behavior, as when by pecking one key a pigeon can turn on the stimuli correlated with an on-going reinforcement schedule on a second key.

Alternatively, discriminative responding based on some stimulus or stimulus property. An organism is said to *attend to* a stimulus or property when variations in that stimulus or property change behavior. For example, a pigeon discriminating blue light from its absence is said to attend to color rather than brightness if variations in wavelength but not intensity change its performance. Cf. DISCRIMINATION, FUNCTIONAL STIMULUS.

Attribution theory: in social psychology, accounting for the behavior of others by attributing the causes of their behavior to properties of the individual rather than to properties of the individual's environment, as when someone is said to have done something because of having a particular disposition or trait rather than because of particular environmental contingencies. Individuals tend to interpret their own behavior in terms of environmental causes but the behavior of others in terms of organismic traits. Thus, behavior analytic accounts sometimes seem counterintuitive (Hineline, 1990).

Audience: the discriminative stimuli that set the occasions on which verbal behavior has consequences. Different audiences set the occasion for different verbal classes. Audience stimuli are typically social (as when a speaker is influenced by cues provided by an attentive listener) but they are not exclusively so (as when someone interacts verbally with a computer terminal).

Augmenting stimulus: see INCENTIVE.

Autobiographical memory: episodic memory. See REMEMBERING.

Autoclitic: a unit of verbal behavior that depends on other verbal behavior for its occurrence and that modifies the effects of that other verbal behavior on the listener. Two classes of autoclitics must be distinguished.

Descriptive autoclitics involve discriminations of one's own behavior, as when the word *not* depends on a mismatch between what one is inclined to say and the appropriateness of saying it; including *not* in the statement cancels some of its effects on the listener. A descriptive autoclitic is verbal responding under the stimulus control of one's own current or pending verbal behavior, including the situation occasioning it.

Relational autoclitics involve verbal units coordinated with other units in such a way that they cannot stand alone, as when plurals depend on quantitative features of events or grammatical tenses depend on temporal features. Novel verbal behavior is sometimes the product of novel combinations of such units: see ADDUCTION.

Autology: the scientific study of the self. "It is because our behavior is important to others that it eventually becomes important to us" (Skinner, 1957, p. 314). Cf. PRIVATE EVENTS.

Automaintenance: the maintenance of autoshaped responding by continuing the AUTOSHAPING procedure. In *negative automaintenance*, reinforcers are omitted on trials with and presented on trials without responses.

Automatic reinforcer: a reinforcer related to a response in such a way that it is usually produced automatically by the response, as in the relation between masturbation and orgasm. Cf. CONTRIVED REINFORCER, INTRINSIC REINFORCER, PRIMARY REINFORCER.

Autoshaping: a respondent procedure that generates skeletal responses. In the most familiar example, a pigeon's key pecks are engendered by presentations of a fixed duration of keylight followed by food, which is not presented at other times. In some procedures, pecks on the key, once they occur, produce food immediately rather than at keylight offset. Cf. AUTOMAINTENANCE.

Aversive control: see specific cases: ESCAPE, AVOIDANCE, PUNISHMENT, PREAVERSIVE STIMULUS.

Aversive stimulus: a stimulus effective as a *negative reinforcer* or as a *punisher*, or that suppresses positively reinforced operant behavior during another stimulus that precedes it (cf. PREAVERSIVE STIMULUS). A stimulus with any one of these effects is likely also to have the others but is not guaranteed to do so. Cf. NOXIOUS STIMULUS, PUNISHMENT, REINFORCEMENT.

Avoidance: the prevention of an aversive stimulus by a response. In *deletion* procedures, the response cancels presentations of the aversive stimulus; in *postponement* procedures, the response only delays them. In *discriminated, discrete-trials* or *signaled avoidance*, an exteroceptive stimulus, sometimes called a *warning stimulus*, precedes the aversive stimulus; a response during this stimulus prevents the aversive stimulus on that trial. If no response occurs and the aversive stimulus is presented, escape from it typically depends on the same response that is effective for avoidance. In *continuous, free-operant* or *Sidman avoidance*, no exteroceptive stimulus is arranged and, typically, there is no provision for escape. Each response postpones the aversive stimulus (usually, brief shock) for a fixed period called the *response-shock (RS) interval*; in the absence of responses, shocks are delivered regularly according to a *shock-shock (SS) interval*. Cf. NEGATIVE REINFORCEMENT.

B

Backward conditioning: respondent conditioning in which the CS follows rather than precedes the US. This procedure can work with aversive stimuli but is otherwise usually ineffective.

Bait-shyness: see TASTE AVERSION.

Bar: lever. See OPERANDUM.

Baseline: a stable and usually recoverable performance on which effects of experimental variables are superimposed. For example, a drug effect may be expressed as the change produced in baseline response rate by a dosage of the drug. The term is also used occasionally to refer to the horizontal starting position (zero responses) of a cumulative recorder pen. Sometimes what the baseline should be is ambiguous; for example, if pecks occur at higher rates when they produce shock than when they don't because shocks are correlated with reinforcers, should the baseline be a no-shock condition with reinforcers or a shock condition without them? Cf. STEADY STATE.

Behavior: anything an organism does. The definition is too inclusive as it stands but cannot easily be restricted further. Behavior is not to be equated with movement. For example, shifts of attention needn't involve eye movement but still qualify as behavior, as when one attends to different performers in listening to a recording of a musical event. The word is often used as a substitute for *responses* (*a behavior, several behaviors*), but this text adheres to colloquial usage, in which *behavior* is a collective term (*kinds of behavior*). See specific cases: COVERT, EMOTIONAL, SPECIES-SPECIFIC, OPERANT, OVERT and RESPONDENT BEHAVIOR.

Behavioral contrast: see CONTRAST.

Behavior analysis: breaking complex behavior down into its functional parts. A successful analysis should allow the behavior to be synthesized by putting the parts back together. Those who conduct behavior analyses are sometimes called *behavior analysts*, those who do so in applied settings are sometimes called *applied behavior analysts*.

Bias: a systematic error in measurement. For example, if a device cannot record all responses when they follow each other rapidly, data recorded with the device will be biased toward low response rates. For other usages, see MATCHING LAW, PREFERENCE, SIGNAL DETECTION ANALYSIS.

Biofeedback: FEEDBACK based on physiological measures (e.g., blood pressure, heart rate, muscle tension).

Blackout: a timeout arranged by turning off all lights in the chamber.

Blocking: an attenuation of respondent conditioning with one stimulus because of prior conditioning with another. For example, if tone and bell together precede food but bell is already a CS, tone may remain ineffective as a CS even though it and bell have the same contingent relation to food. Cf. OVERSHADOWING.

Boundaries: See COMPARTMENTS.

Break: an abrupt transition from responding to no responding (cf. RATIO STRAIN).

Bribes: statements of contingencies that substitute for a parent's differential reinforcement of a child's behavior, so that compliance with requests comes to depend upon the statements as discriminative stimuli. The child who is frequently bribed learns to discriminate between conditions in which bribes are in effect and those in which they are not. The parent who often bribes will probably find that the child complies only when bribes are offered.

Burst: high-rate responding bounded by lower-rate responding.

Butterfly Effect: See CHAOS THEORY.

C

Categorical perception: discrimination involving a discrete change in responding with continuous changes along a continuum, as when gradual changes in the acoustic formants that distinguish the phonemes *b* and *p* produce responses that shift from reporting *b* to reporting *p*, without producing any reports of intermediate forms such as blends of *b* and *p*.

Categories: see specific cases: ABSTRACTION, CONCEPT, EQUIVALENCE CLASS, NATURAL CONCEPT, POLYMORPHOUS STIMULUS CLASS, PROBABILISTIC STIMULUS CLASS, PROTOTYPE.

CER: conditioned emotional response. See PREAVERSIVE STIMULUS.

Celeration: see RATE OF RESPONDING.

Chaining: the creation of maintenance of a sequence of discriminated operants such that responses during one stimulus are followed by other stimuli that reinforce those responses and set the occasion for the next ones (see CHAINED SCHEDULE, CONDITIONAL REINFORCER). Not all sequences are maintained through chaining; those that are should be distinguished from those that are not (cf. CHUNKING). Parts of a chain are called *components, links* or *members*. Procedures for creating chains often start at the end closest to the reinforcer and then work back (*backward chaining*); starting from the other end (*forward chaining*) is more difficult, because early responses may extinguish while later ones are being shaped. Chains with topographically similar responses are *homogeneous* (e.g., pecks maintained by a chained schedule); those with topographically dissimilar responses are *heterogeneous* (e.g., alley running and then lever pressing).

Chained (chain) schedule: a compound schedule in which reinforcers are produced by successive completions of two or more component schedules, each operating during a different stimulus. Equivalent arrangements with the same stimulus during each component are *tandem* schedules.

Chamber: a space designed to minimize interference from stimuli irrelevant to experimental conditions (e.g., laboratory noises) and including devices for presenting stimuli and for recording behavior (see OPERANDUM). Typical chambers may contain mechanisms for delivering reinforcers (e.g., food dispensers), discriminative stimulus sources (e.g., speakers for auditory stimuli, lamps or projectors for visual stimuli), aversive stimulus sources (e.g., see SHOCK), a houselight for general illumination, feedback devices that produce stimuli such as clicks after each response, and auditory sources that mask outside noise (sometimes, a fan that provides masking noise along with ventilation).

Changeover: the switch from one response to another, as when a pigeon in a two-key chamber moves from left-key pecking to right-key pecking. Sometimes the changeover is made explicit, as when a peck on one key changes the schedule available on a second key. Cf. CONCURRENT OPERANTS.

Changeover delay (COD): a feature sometimes used with concurrent schedules to prevent sequences in which a reinforcer produced by one response closely follows the other response. As usually arranged, the COD provides that no response can be reinforced within t s after a changeover.

Changeover ratio (COR): a changeover contingency that provides that no response can be reinforced until at least n responses after the last changeover. Cf. CHANGEOVER DELAY.

Chaos theory: a branch of mathematics dealing with nonlinear systems, which are drastically affected even by very small changes in initial values. For example, in what is called the Butterfly Effect the path of a storm system may be influenced by the flaps of a butterfly's wings weeks before. Like the weather, behavior is a nonlinear system. For measurement at any level of precision, the mathematics of chaos implies that we can predict kinds of things that will happen but we cannot predict specific details. For example, we can predict that a pigeon will peck a key but not precisely when. This makes interpretation far more important; in many behavioral applications it is all that is feasible.

Choice: the emission of one of two or more alternative and, usually, incompatible responses. The term is sometimes applied to other behavior that precedes the response, especially in human examples, as when someone deliberates before selecting an alternative, but even in this instance deliberating precedes choosing. Choice is not some other response that precedes the selection;

it is the selection itself. Cf. CONCURRENT OPERANTS, PREFERENCE.

Chunking: the creation of temporally extended units in which sequences become so well established that each response no longer depends on being occasioned by prior responses or reinforced by subsequent ones (cf. CHAINING), as in learning to play rapid musical passages (cf. INTRAVERBAL); in verbal learning, the arbitrary creation of larger verbal units, as when a mnemonic system is used to convert a number sequence to a single word.

Classical conditioning: see RESPONDENT CONDITIONING.

Closed economy: in operant contexts, the availability of appetitive stimuli only within the session as reinforcers, with none provided independently of behavior on a supplementary basis outside of the session. Cf. OPEN ECONOMY.

Clustering: In free recall, the reorganization of items by the learner so that related ones are recalled together rather than in the order in which they appeared on the list.

cm: centimeters (abbreviation). One cm is about 0.4 inches.

COD: see CHANGEOVER DELAY.

Coding, coding response: an inferred variety of mediating behavior, as when humans remember visually presented letters on the basis of their sound rather than their geometric properties, perhaps as a result of saying or subvocally rehearsing them. Errors based on acoustic rather than visual similarity support the inference. Tacting is one kind of coding. Cf. DECODING, ENCODING.

Cognition, cognitive processes: knowing and the ways in which it takes place. Processes said to be cognitive are often varieties of BEHAVIOR that are not manifested as movements and so must be measured indirectly, as in doing mental arithmetic, shifting attention, imagining). Cf. COVERT BEHAVIOR.

Cognitive map: a spatial SCHEMA or representation. In learning a coordinated set of spatial relations, an organism is sometimes said to develop a cognitive map. The term is most likely to be invoked when an organism orients toward locations it cannot see or otherwise directly respond to.

Collateral behavior: responding that, like MEDIATING BEHAVIOR, appears in a consistent sequential relation to reinforced behavior while not itself instrumental in producing reinforcers. It does not carry the implication that this responding makes the reinforcer more likely.

Comparison stimulus: see MATCHING-TO-SAMPLE.

Compartments: separate regions in physiological structures, as in segmentation during embryonic development. These regions may isolate functions, sometimes called *modules*, so that different classes of processes can change independently. In evolution, for example, boundaries limit the effects of mutation to compartments, so that a mutation does not produce potentially lethal changes in other compartments. Modularity permits a given body part to be modified and specialized in evolution independently of other body parts. Modularity may play a similar role in behavior, as when different operant classes are created and become more finely differentiated; it is likely to be the basis for CATEGORICAL PERCEPTION.

Component: one of the schedules, or the stimulus associated with it, in a compound schedule. The term is usually restricted to cases in which the schedules making up the compound operate successively rather than simultaneously.

Compound schedule: a schedule that combines two or more component schedules. The components may operate successively, in regular or irregular alternation (MULTIPLE and MIXED) or as a sequence (CHAINED and TANDEM), or simultaneously (CONCURRENT and CONJOINT); they may also interact (ALTERNATIVE, CONJUNCTIVE, INTERLOCKING). See these and other specific cases: ADJUSTING, CONCURRENT-CHAIN, HIGHER-ORDER, PERCENTILE-REINFORCEMENT, PROGRESSIVE.

Concept: a class of stimuli such that an organism generalizes among all the stimuli within the class but discriminates them from those in other classes. Concepts are to analyses of discriminative stimuli as operants are to analyses of response classes (Keller & Schoenfeld, 1950). Cf. ABSTRACTION, DISCRIMINATION, GENERALIZATION, STIMULUS.

Concurrent-chain schedules: concurrent schedules in which the reinforcers are themselves schedules that operate separately and in the presence of different stimuli, as when equal and independent concurrent VI VI schedules operate for a pigeon's pecks on two white keys and, according to the VI schedules, left pecks produce an FI schedule on a blue key and right pecks produce an FR schedule on a yellow key. The concurrent VI VI schedules are *initial links* and the separate schedules they produce are *terminal links*. The initial links can be thought of as doors that admit the pigeon to separate rooms containing the terminal links. PREFERENCE for terminal links is given by relative response rates in initial links.

Concurrent operants: two or more classes of alternative responses. Concurrent operants may be compatible, as when a rat simultaneously presses one lever with its left paw and another with its right, or incompatible, as when the pigeon, having only one beak, pecks only one of two keys at a time, as long as the organism can emit either or change over from one to the other at any time; responding and not responding are occasionally treated as concurrent operants. Discriminated operants can be concurrent if the organism has an opportunity to produce the stimuli that occasion them at any time. For example, in a *changeover-key procedure*, a pigeon changes the stimuli and their associated schedules on one key (the *main* key) by pecking a second key (the *changeover* key). In this case, two CONCURRENT SCHEDULES operate on one key, and the CHANGEOVER is an explicit response on a second key. See also PREFERENCE.

Concurrent (conc) schedules: two or more schedules operating simultaneously and independently, each for a different response, as when separate VI schedules are arranged for a pigeon's pecks on each of two keys. Cf. CONJOINT SCHEDULES.

Conditional: a preferred alternative to *conditioned*.

Conditional discrimination: a discrimination in which reinforcement of responding during a stimulus depends on (is conditional on) other stimuli. For example, matching-to-sample involves a conditional discrimination in the sense that the comparison response that will be reinforced depends on the sample. Conditional discrimination procedures involve four-term contingencies: They arrange stimuli during which different three-term contingencies operate.

Conditional probability: the PROBABILITY of one event given another event. For example, if A and B occur equally often but A is followed by A 75% of the time and by B 25% of the time, the simple probability of A is .5 but its conditional probability given a prior A is .75. See INTERRESPONSE TIME for another example.

Conditioned: see CONDITIONAL.

Conditioned aversive stimulus: a stimulus that acquired its aversive properties by reliably accompanying another aversive stimulus. For example, in discriminated avoidance, the warning stimulus may become a conditioned aversive stimulus. Cf. PREAVERSIVE STIMULUS.

Conditioned emotional response (CER): see PREAVERSIVE STIMULUS.

Conditioned reflex or conditional reflex: a reflex produced by a contingent relation between stimuli (see CONTINGENCY). One stimulus, originally neutral, sets the occasion for a second stimulus, the *unconditioned stimulus* (*US*). A conditioned reflex is created when the neutral stimulus becomes a *conditioned stimulus* (*CS*), eliciting a response by virtue of its contingent relation to the US. This response, a *conditioned response* (*CR*), is often related to but is not necessarily the same as the *unconditioned response* (*UR*) elicited by the US. Responses elicited by the CS before conditioning (e.g., orienting responses) tend to disappear as conditioning progresses (cf. HABITUATION).

Although the procedure is usually characterized in terms of CS-US contingencies, it may depend on the CS-UR relations produced by those contingencies. The most typical respondent procedure, in which a CS is followed within no more than 5 s by the US, is called *simultaneous conditioning*; brief delays were incorporated into most so-called simultaneous procedures because the CR cannot be measured independently of the UR if the CS and US are simultaneous, and strict simultaneity is less effective in conditioning than a brief delay between CS and US. Effective USs in respondent conditioning are often effective positive or negative reinforcers in operant selection, and an older usage of *reinforcement* referring to US presentations still survives in some parts of the learning literature. See BACKWARD, DELAY, HIGHER-ORDER, TEMPORAL and TRACE CONDITIONING; cf. UNCONDITIONED REFLEX, RESPONDENT.

Conditional reinforcer or **conditioned reinforcer:** a stimulus that functions as a reinforcer because of its contingent relation to another reinforcer. Such stimuli have also been called *secondary* reinforcers, but this designation is best reserved for cases in which the modifier specifies how many stimuli separate the conditioned from the primary reinforcer (e.g., secondary related directly to primary, tertiary to secondary, etc.). Convenience often dictates the assigned order. For example, feeder operations are usually called primary reinforcers even though the auditory and/or visual stimuli that accompany them are actually conditional reinforcers that precede eating.

Conditioned response or **conditional response (CR)** and **conditioned stimulus** or **conditional stimulus (CS):** see CONDITIONED REFLEX.

Conditioned suppression: see PREAVERSIVE STIMULUS.

Conditioning: see RESPONDENT CONDITIONING. The term appears occasionally in conjunction with *operant*, but *operant conditioning* is becoming a rarer usage.

Conflict: a situation in which a single response produces both reinforcers and punishers (*approach-avoidance* conflict), or two or more incompatible responses produce different reinforcers (*approach-approach* conflict), or two or more incompatible responses each avoid only one of two or more aversive stimuli (*avoidance-avoidance* conflict). Approach-avoidance conflicts often produce vacillating behavior, as in love-hate relationships.

Conjoint (conjt) schedules: two or more component schedules, usually involving different reinforcers, operating for a single response. For example, in conjt FR avoid, lever presses simultaneously are reinforced according to an FR schedule and postpone shock according to an avoidance schedule. Cf. CONCURRENT SCHEDULES.

Conjugate reinforcement: reinforcement in which some property of a reinforcer varies systematically with some response property, as when the sharpness of focus of a visual reinforcer increases with the momentary rate of responding or as when the position of a car in a lane varies continuously with movement of the steering wheel.

Conjunctive (conj) schedule: a schedule that reinforces a response when each of two (or more) schedule requirements is satisfied. For example, in conj FI 60-s FR 10, a response is reinforced only after at least 60 s have elapsed and at least nine other responses have been emitted since the last reinforcer.

Consequence: an event produced by some other event; especially, in operant contexts, an event produced by a response, such as stimulus presentation or removal, a change in contingencies or any other environmental change. The term is particularly useful for referring to events with an unknown status as reinforcers or punishers; presenting such events contingent on responding has sometimes been called *consequation*.

Consolidation of memory: a theoretical process based on the assumption that remembering is relatively impermanent immediately after learning and takes some time to become fixed or consolidated.

Constituent grammars: phrase-structure grammars. See GRAMMARS.

Constructive memory: see RECONSTRUCTION.

Consummatory response: the behavior occasioned by a reinforcer. The term originated with reinforcers that were consumed (food, water) but has been extended to other kinds. For example, if opportunities to run in a wheel are reinforcers, wheel running is the consummatory response.

Context: the constant features of an situation, such as the chamber in which an operant session occurs. Experimental contexts acquire behavioral function because they are embedded in the still larger contexts that include the experimental session.

Contiguity: the juxtaposition of two or more events when they occur simultaneously or very closely together, as in the incidental succession of a response and a reinforcer in free reinforcement or of a CS and a US in a respondent procedure. Cf. CONTINGENCY.

Contingency: in the operant case, the conditions under which a response produces a consequence. For example, in an FI, the reinforcer is *contingent on* a response of a given force, topography, etc., as well as on the passage of time. An organism is said to *come into contact with a contingency* when its behavior produces some consequences of the contingency.

Studies of reinforcement schedules analyze contingencies and their effects, as when contingencies of reinforcement for various IRTs in VI and VR schedules are compared. In this most general usage, contingencies describe any relation, whether completely specified in a procedure or the incidental and perhaps fortuitous consequence of them. In a more specific sense, contingencies are the conditional probabilities relating some events (e.g., responses) to others (e.g., stimuli). When responses produce reinforcers, the contingent relation is defined by two conditional probabilities: probability of the reinforcer (1) given a response and (2) given no response. Without both probabilities specified, the contingent relations cannot be distinguished from the incidental temporal *contiguities* of responses and reinforcers that occur independently over time. Response-reinforcer relations involve two terms, but when correlated with discriminative stimuli they produce a *three-term contingency*. Conditional discriminations add a fourth term, and so on for other contingency relations of various orders of complexity.

When applied to respondent cases, *contingency* refers to the conditions under which some stimuli are followed by others. By analogy to the operant case, *stimulus-stimulus contingencies* expressed as conditional probabilities specify conditions more completely than descriptions in terms of temporal *contiguities*, and distinguish cases in which two stimuli always occur together from those in which they are frequently paired but also occur independently. Stimuli correlated with stimulus-stimulus contingencies, sometimes called *occasion setters*, may enter into three-term or higher-order relations.

Contingency-governed behavior or **contingency-shaped behavior:** operant behavior. The terminology is ordinarily used to contrast responding that is not occasioned by verbal behavior with *verbally governed behavior*, behavior controlled by verbal antecedents (e.g., instructions).

Contingency space: any coordinate system within which contingencies expressed as conditional probabilities are plotted.

Contingent associations: associations that involve all possible combinations, so that their learning is contingent on the entire configuration of each item, as when a verbal discrimination task uses two lists of three-letter words arranged so that every possible letter combination appears in both lists. For example, with list-1 items FIT FAN PIN PAT and list-2 items FIN FAT PIT PAN, every possible letter and pair of letters appears as often in list 1 as in list 2 (e.g., F--, -I-, --T; PA-, P-N, -AN). Only whole items can consistently occasion correct responses; no part will work. Lists constructed this way can be exceedingly difficult to learn.

Contingent stimulus: a stimulus the presentation of which depends on a response-stimulus contingency.

Continuous avoidance: see AVOIDANCE.

Continuous reinforcement (CRF): reinforcement of every response qualifying as a member of a descriptive operant class.

Continuous repertory or **continuous repertoire:** behavior that tracks continuous changes in some environmental property, as when a driver steers a car to keep it in its lane. Cf. CONJUGATE REINFORCEMENT.

Continuum: see STIMULUS CONTINUUM.

Contrast or **behavioral contrast:** a change in the rate of one response that occurs when either the rate of a second response or the reinforcement rate produced by that response changes in the opposite direction, where the reinforcement rate maintaining the first response remains constant. For example, the rate of reinforced responding in one multiple-schedule component typically increases if reinforcement is reduced or discontinued in the other. The term usually applies to effects during successive stimuli, as in multiple schedules, but similar phenomena occur in concurrent schedules.

The appropriate baseline is responding maintained by reinforcement in one component alone rather than responding maintained by reinforcement in both components. Extinction in one component does not excite responding in the other; instead, reinforcement in each component reduces responding in the other. See INHIBITION.

Contrived reinforcer: an artificial reinforcer. Cf. CONDITIONAL REINFORCER, EXTRINSIC REINFORCER.

Control: the systematic modification or maintenance of behavior by changes in relevant conditions. The manipulation of conditions distinguishes control from *prediction* and *interpretation*. If control is not possible because relevant conditions are not manipulable, adequate information about relevant variables may allow prediction (as in the history of astronomy before space flight). Interpretation usually occurs after the fact. Given an outcome, a plausible account of the relevant variables can be offered, but it may be difficult to determine its adequacy. Nevertheless, such analyses are often expected or demanded of students of behavior, as when a psychologist is asked to explain in a court of law why a defendant acted in some way. In its commonest behavior analytic usage, the term appears in conjunction with some variable that has a demonstrable effect on behavior (e.g., *schedule control, stimulus control*). Cf. CHAOS.

Coordinate: the value of a point plotted on a graph. In a two-dimensional system, the one plotted along the x-axis is called the *abscissa* and the one plotted along the y-axis is called the *ordinate*; the terms do not refer to the axes themselves.

COR: see CHANGEOVER RATIO.

Correction procedure: the repetition or continuation of conditions and/or stimuli after certain responses or their absence, especially after errors in simultaneous discrimination trials. For example, stimuli may be repeated in a new trial if an error occurred in the last one or, with trials of limited duration, if no response occurred; or a multiple VI extinction schedule may be arranged so that each response during the extinction component delays the onset of the VI component. The term can refer to any procedure that arranges continued or repeated opportunities for responses in alternative classes until a given response does or does not occur. Its colloquial origins imply procedures that eventually force an organism to emit a correct response (cf. ERROR), but the usage does not exclude procedures in which the alternative response classes cannot easily be categorized as corrects and errors. Correction procedures sometimes make it likely that errors will be followed by reinforced correct responses that contribute to their maintenance; cf. DELAY OF REINFORCEMENT.

Correlated stimuli or **correlated reinforcers:** see CONJUGATE REINFORCEMENT.

Correlation: see STATISTICS. The term is often applied to stimulus control procedures, as when a schedule operating during the presence but not the absence of a stimulus is said to be *correlated with* the stimulus. *Molar* analyses may include examining correlations between overall response rates and overall reinforcement rates.

Correspondence (between saying and doing): see VERBALLY GOVERNED BEHAVIOR.

Cost: see RESPONSE COST.

Counterconditioning: in clinical applications, overcoming through new conditioning the effects of a CS that, by virtue of its relation to an aversive US, has become aversive. Respondent contingencies involving a new US are arranged so as to produce new CRs incompatible with those of the original conditioning.

Coupling: when reinforcement is interpreted as made up of two components, the linking of reinforced responding to the behavior activated by reinforcers, as opposed to the arousal or energizing effects of those reinforcers.

Covert behavior: behavior that isn't observed or observable and so is only inferred. Alternatively, behavior inside an organism, but of such a sort or on such a small scale that it is not recordable or is recordable only with special equipment. Examples include thinking or counting to oneself, perhaps inferred from a verbal report, and muscle contractions too small to produce obvious movement.

CR: conditioned response. See CONDITIONED REFLEX.

CRF: see CONTINUOUS REINFORCEMENT.

Critical feature: a feature, perhaps one of several, on which discrimination among stimuli depends. For example, letters of the alphabet include straight versus curved, open versus closed, etc., as critical features. Cf. FUNCTIONAL STIMULUS.

Critical period: the time during which a stimulus can become imprinted (cf. IMPRINTED STIMULUS). More generally, any time period to which the operation of some behavioral process is limited. For example, arguments have been advanced for a critical period for language acquisition in humans.

CS: conditioned stimulus. See CONDITIONED REFLEX.

Cue-dependent learning: see STATE-DEPENDENT LEARNING.

Cultural selection: the selection of behavior as it is passed on from one organism to another. Examples include imitated behavior and verbal behavior, sometimes also called *memetic selection*. Cf. OPERANT SELECTION.

Cumulative record: a record that shows total responses plotted as a function of time, usually made by a marker or pen that moves a fixed distance with each response across a paper advancing at a constant speed. Thus, the faster the responding, the steeper the slope. Moment-to-moment changes in slope show the details of changing response rates over time. Cumulative recorders typically include additional features; for example, the pen can be reset to its starting position (sometimes called the *baseline*) after a full excursion across the paper or after some event; it can be briefly displaced downward or to one side, producing a pip, to indicate a reinforcer or other brief event; it can be held in its downward position, producing a displaced line, to indicate stimuli or other extended conditions; and an event pen at the bottom of the record can be used to indicate other events. Cf. RATE OF RESPONDING.

Cusps: tipping points. When behavior that has been changing gradually reaches a point where it makes contact with new contingencies, it may abruptly begin to change very rapidly.

CVC trigram: a verbal item made up of a consonant, a vowel and a consonant, in that order (often but not necessarily a nonsense syllable).

D

Data; datum: recorded information, usually in numerical form. *Data* is plural; *datum* is singular.

Declarative memory: see REMEMBERING.

Decoding: the learner's behavior with respect to an item to be remembered at the time it is recalled. Cf. ENCODING.

Deep structure: the common structural features of sentences related to each other by grammatical transformations. Cf. SURFACE STRUCTURE.

Defensive conditioning: respondent conditioning with an aversive or noxious stimulus as US.

Deictic verbal behavior, deixis: verbal behavior in which terms have functions that vary relative to the speaker, as with personal pronouns, *here* versus *there*, and *this* versus *that*. Deixis depends on discriminations of the speaker's own behavior and so shares some properties with autoclitic behavior.

Delay conditioning: respondent conditioning in which the CS is presented for some fixed extended time before the US (in most usages, no less than 5 s).

Delay of reinforcement: the time from a response to a reinforcer. Reinforcers lose their effectiveness with increasing delays, but procedural complications stand in the way of measuring delay-of-reinforcement effects. When delay procedures interpose a stimulus between a response and its delayed reinforcer (*signaled delay of reinforcement*), the stimulus probably functions as an immediate conditional reinforcer. When procedures interpose no stimulus, either additional responses reset the delay, thereby limiting response rate because a reinforcer is delivered only after a pause equal to the delay, or they do not reset the delay, thereby allowing the actual delay to be reduced to the now shorter time between those additional responses and the reinforcer.

Delayed reinforcement may operate whenever a succession of responses is followed by a reinforcer; the effect of the reinforcer depends not only on its relation to the last response, which produced it, but also on the others that preceded it. Thus, reinforcement schedules allow each reinforcer to reinforce many responses and not just one response. This may be a problem when the earlier responses include errors, as in matching-to-sample that includes a CORRECTION PROCEDURE.

Delayed matching-to-sample: see MATCHING-TO-SAMPLE, DELAYED RESPONSE.

Delayed response: a response that occurs some time after a discriminative stimulus is removed, as when the sample in a matching-to-sample procedure is turned off several seconds before the comparison stimuli are presented. See MEDIATING BEHAVIOR for an example. Cf. REMEMBERING.

Deletion: see AVOIDANCE.

Density: a synonym for rate, (as in *reinforcement density* or *shock density*. Cf. RATE OF REINFORCEMENT.

Dependency: roughly, a contingency completely specified by the experimenter or with a conditional probability very close to 1.0. Cf. CONTINGENCY.

Deprivation: a reduction in the availability of a reinforcer. With food reinforcers, percentage of free feeding weight and the time since free feeding have been used as criteria for deprivation levels. Deprivation may be a condition for making any positive reinforcer effective, as when locking a running wheel deprives the rat of an opportunity to run. Cf. ESTABLISHING OPERATION.

Descriptive autoclitic: see AUTOCLITIC.

Development of language: see LANGUAGE DEVELOPMENT.

Dictation-taking: A formal verbal class in which a spoken verbal stimulus occasions a corresponding written response. The correspondence is defined by the one-to-one relation of verbal units (e.g., letters or words). Cf. ECHOIC BEHAVIOR, TEXTUAL BEHAVIOR, TRANSCRIPTION.

Differential conditioning: usually, producing a discrimination in respondent conditioning. See RESPONDENT DISCRIMINATION.

Differential reinforcement: reinforcement of some responses but not others, depending on intensive, temporal, topographical or other properties of the responses (including the stimuli during which they are emitted; cf. DISCRIMINATED OPERANT); differential reinforcement defines operant classes. When the proportion of responses within the limits of the operant class increases as a result of differential reinforcement, responding is said to be *differentiated*.

Differential-reinforcement schedules: schedules of differential reinforcement, especially when reinforcers depend on the temporal spacing of responses. Contingencies can be based on interresponse times (IRTs), response rates, or periods of no responding. They are usually arranged for free-operant responding but also can be arranged within discrete trials.

In *differential-reinforcement-of-long-interresponse-times*, *differential-reinforcement-of-low-rate* (*DRL*) or $\mathit{IRT}>t$ schedules, a response is reinforced only if at least t s has passed since the last response. An alternative and less common method is based on rates rather than IRTs; it reinforces a response only if fewer than n responses were emitted during the last t s.

In *differential-reinforcement-of-high-rate* (*DRH*) or $\mathit{IRT}<t$ schedules, a response is reinforced if at least n responses were emitted during the last t s. The alternative based on IRTs rather than rates is less common because reinforcing single short IRTs tends to produce short IRTs separated by pauses rather than sustained high response rates.

A *paced-response* or *differential-reinforcement-of-pacing* (*DRP*) schedule arranges upper and lower limits on reinforced IRTs (e.g., an IRT between 2 and 4 s) or on reinforced response rates (e.g., between 10 and 15 responses during the last 5 s). When the schedule is based on IRTs rather than rates, it is sometimes called *DRL with limited hold*.

Another type delivers a reinforcer after a response is followed by t s of no responding (cf. DELAY OF REINFORCEMENT). Still another reinforces not responding by delivering a reinforcer after t s without a response; this schedule is called *differential reinforcement of zero behavior* or *other behavior* (*DRO*). See also INTERRESPONSE TIME, RATE OF RESPONDING.

Differentiation: see DIFFERENTIAL REINFORCEMENT.

Discrete trials: see TRIAL.

Discriminated operant: an operant defined in terms of the stimuli during which it occurs as well as its environmental effect. This operant depends on the relations among three events (the *three-term contingency*): a *stimulus* in the presence of which a *response* may have *consequences*. In one sense, the stimulus sets the occasion on which the response may be reinforced; in another, it defines a property of the operant and so sets the occasion for the response. The joint dependence of the response on both stimulus and reinforcer distinguishes the relation from that of a reflex. See also OPERANT.

Discrimination: any difference in responding in the presence of different stimuli; in a more restricted usage, a difference resulting from differential consequences of responding in the presence of or during different stimuli. See also DISCRIMINATED OPERANT, GENERALIZATION, RESPONDENT DISCRIMINATION, SIMULTANEOUS DISCRIMINATION, STIMULUS, SUCCESSIVE DISCRIMINATION, VERBAL DISCRIMINATION. Usually the organism is said to discriminate among relevant stimuli. In some cases, however, it is useful to speak of responses as discriminating. For example, if response rate varies with color whereas response location varies with form, rate discriminates color while location discriminates form.

When discriminated responding is produced by differential contingencies in the presence of different stimuli, it is appropriate to say that the stimuli are discriminated but not that the contingencies are discriminated.

Discriminative stimulus: any stimulus with a discriminative function; according to an older usage, a stimulus correlated with reinforcement when another is correlated with extinction. The latter usage has become less common because it is not applicable to stimuli correlated with different schedules (e.g., multiple FI FR); it was the source of S^D ("S-dee": discriminative stimulus) and S^Δ ("S-delta": absence of discriminative stimulus) as abbreviations for stimuli correlated with reinforcement and with extinction. The abbreviations have lost some ground to S+ (positive stimulus) and S- (negative stimulus). Strictly, S^0 ("S-zero") is more appropriate for the absence of reinforcement but S- is typographically more convenient.

Displacement activity: an ethological term referring to a response occurring not in the presence of the RELEASER that usually produces it but rather in the presence of one that usually produces some other response. Displacement activity and VACUUM ACTIVITY depend on deprivation of opportunities to complete a FIXED ACTION PATTERN, but displacement activity is likely to occur at lower levels of deprivation than vacuum activity.

Distributed practice: spacing periods of activity on a task. Cf. MASSED PRACTICE.

Distribution: a classification of events by location along a continuum. For example, an IRT distribution classifies IRTs into several temporal categories (e.g., less than 1 s, 1 but less than 2 s, 2 but less than 3 s, and 3 or more s). *Frequency distributions* show the number of events per category; *relative frequency distributions* show the events per category as a proportion of the total. Each category is called a *class interval*, and class intervals are ordinarily of equal size. Distributions often include a category for all events falling beyond some point on the continuum (e.g., in the above example, 3 or more s) so that a category exists for any event no matter how extreme. Cf. INTERRESPONSE TIME, STATISTICS.

DMTS: delayed matching-to-sample. See DELAYED RESPONSE, MATCHING-TO-SAMPLE.

DRH: differential reinforcement of high rate. See DIFFERENTIAL-REINFORCEMENT SCHEDULES.

Drive operation: see MOTIVATING OPERATION.

DRL: differential reinforcement of long interresponse times or low rate. See DIFFERENTIAL-REINFORCEMENT SCHEDULES.

DRO: differential reinforcement of zero behavior or other behavior. See DIFFERENTIAL-REINFORCEMENT SCHEDULES.

DRP: differential reinforcement of paced responding or pacing. See DIFFERENTIAL-REINFORCEMENT SCHEDULES.

Duplic verbal relations: verbal responding characterized by one-to-one correspondences between verbal constituents of the stimulus and those of the response, as in echoic, textual, dictation-taking and transcription behavior; verbal behavior in which verbal responding replicates verbal stimuli.

Duration of response: the time from the beginning to the end of a response (sometimes called *holding time*). Analyses of this property of responding depend heavily on procedural details. For example, if a reinforcer is delivered when a lever is pressed, response duration is short because the reinforcer occasions quick release, but if it is delivered on lever release, then each member of the chain (press, hold, release) may be differently affected by the contingencies.

Duration of stimulus: see TEMPORAL DISCRIMINATION.

E

Echoic behavior: A formal verbal class in which a spoken verbal stimulus occasions a corresponding spoken verbal response. The correspondence is defined by the one-to-one relation of verbal units (e.g., phonemes or words) and not by acoustic similarity. Cf. DICTATION-TAKING, TEXTUAL BEHAVIOR, TRANSCRIPTION.

Echoic memory: see REMEMBERING.

Eidetic memory: sometimes called photographic memory, a rare type of remembering, usually in children, in which visual stimuli are described in detail as if seen long after they were presented.

Effect, Law of: see LAW OF EFFECT.

Elaborative rehearsal: see REHEARSAL.

Elicitation: the reliable production of a response by a stimulus in unconditioned or conditioned reflexes. Cf. RESPONDENT BEHAVIOR.

Emergence, Emergent relation: a new behavioral relation, including the combination of behavior classes in adduction or extended equivalence relations, coming about as a byproduct of other relations rather than through differential reinforcement. For example, if arbitrary matching has been trained only for AB and BC, where the first letter of each corresponds to the sample and the second to the matching comparison, and transitivity testing demonstrates matching with the new pair AC, this new matching relation is said to be emergent. Cf. ADDUCTION, EQUIVALENCE CLASS.

Emission: the occurrence of behavior in the absence of an eliciting stimulus. Responding need not imply something responded to. A response that occurs without an eliciting stimulus is said to be *emitted*. The term applies to responding occasioned by discriminative stimuli as well as to undiscriminated responding. An emitted response need not be treated as operant until it can be shown to be modifiable by its consequences. Cf. OPERANT BEHAVIOR, SPOND.

Emotional behavior: correlated changes in a range of response classes. For example, if a preaversive stimulus simultaneously alters heart rate, respiration, blood pressure, defecation and operant behavior maintained by reinforcement, it is said to produce emotional behavior. This and related terms evolved from an imprecise colloquial vocabulary, so types of emotional behavior cannot be defined unambiguously in terms of response classes. The operations that produce them allow them to be more consistently defined: for example, *fear, anxiety* or, with another organism present, *anger*, produced by primary or conditioned aversive stimuli; *relief*, produced by the termination of aversive stimuli; *joy* or *hope*, produced by primary or conditional reinforcers; and *sorrow*, produced by the termination of reinforcers. Nevertheless, they still do not function as technical terms. Cf. AGGRESSION, FRUSTRATION, PREAVERSIVE STIMULUS.

Encoding: the learner's behavior with respect to the item to be remembered at the time it is presented. Cf. DECODING, REHEARSAL.

Environment-based theory. Cf. ATTRIBUTION THEORY.

Episodic memory: Autobiographical memory. See REMEMBERING.

Equivalence class: a stimulus class, usually produced through conditional discrimination in matching-to-sample, that includes all possible EMERGENT RELATIONS among its members. The following assumes at least two classes, A_1, B_1, C_1 and A_2, B_2, C_2, but for convenience the examples show just one class with the subscripts omitted. The properties of an equivalence class are derived from the logical relations of reflexivity, symmetry and transitivity. *Reflexivity* refers to the matching of a sample to itself, sometimes called identity matching: AA, BB, CC. In these examples, each letter pair represents a respective sample and its matching comparison. *Symmetry* refers to the reversibility of a relation: if AB, then BA. *Transitivity* refers to the transfer of the relation to new combinations through shared membership: if AB and BC, then AC. If these properties are characteristics of matching-to-sample, then training AB and BC may produce AC, BA, CA and CB as emergent relations; reflexivity provides the three other possible relations, AA, BB and CC. Given AB and BC, for example, the combination of symmetry and transitivity implies the CA relation. The emergence of all possible stimulus relations after only some are trained through contingencies is one criterion for calling the three stimuli members of an equivalence class. Each class can be extended by training new stimulus relations. For example, if CD is learned, then AD, DA, BD, DB and DC may be created as emergent relations. Stimuli that are members of an equivalence class are likely also to be *functionally equivalent*. It remains to be seen whether the logical properties of these classes are fully consistent with their behavioral ones. Cf. EQUIVALENCE RELATION.

Equivalence relation: a term with various usages, including functional equivalence (the relation between stimuli that have become members of a *functional class*) as well as the mathematical relations that define an *equivalence class* (especially the CA relation). The terminology of equivalence relations has often been interchanged with that of equivalence classes, but functionally equivalent stimuli needn't be members of an equivalence class. Cf. EQUIVALENCE CLASS.

Error: in simultaneous discriminations, a response to a stimulus not correlated with reinforcement; in successive discriminations, a response during a stimulus correlated with extinction. Because of its colloquial origins, the term often assumes evaluative as well as descriptive functions. Cf. CORRECTION PROCEDURE.

Escape: the termination of an aversive stimulus by a response. A reduction in the magnitude of an aversive stimulus by a response is sometimes called partial or fractional escape. Cf. REINFORCEMENT.

Establishing operation: any operation that changes the status of a stimulus as a reinforcer or punisher, and especially stimulus presentations that change the reinforcing or punishing status of other stimuli, as when an already available screwdriver becomes a reinforcer in the presence of a screw that needs tightening. SEE MOTIVATING OPERATION

Estes-Skinner procedure: see PREAVERSIVE STIMULUS.

Ethology: an area of biology concerned with the analysis of behavior patterns that evolve in natural habitats, either in species or in individual organisms, with particular emphasis on those patterns that do not depend on or are not known to depend on prior operant selection or respondent conditioning. Cf. FIXED ACTION PATTERN, SPECIES-SPECIFIC BEHAVIOR, RELEASER.

Event recorder: a device that records events as displacements of markers or pens along timelines generated at a constant rate.

Evocation: the production of a response, usually by an motivational or establishing operation, as when food deprivation is said to *evoke* behavior that has led to food in the past. Sometimes responding is said to be evoked if it is unclear whether it was emitted or elicited.

Evolution: changes in populations over time. Evolution can operate at phylogenic, ontogenic and cultural levels. Evolution happens; NATURAL SELECTION is a theory of how it works. Cf. CULTURAL SELECTION, ONTOGENIC SELECTION.

Excitation: roughly, the production of behavior or the variables that produce it, used especially in contrast with *inhibition*.

Exercise, Law of: see LAW OF EXERCISE.

Expectancy: a colloquial term referring to behavior that precedes predictable events. Expectancy depends on a history with respect to the events; it cannot depend on events that have not yet occurred.

Expert performance: behavior that was originally verbally governed but, after extended exposure to natural contingencies, has become independent of verbal antecedents. It is sometimes called *intuitive* and is directly sensitive to contingencies

EXT: extinction. See EXTINCTION.

Exteroceptive stimulus: any stimulus presented at or outside of the organism's skin. Cf. INTEROCEPTIVE STIMULUS.

Extinction: in operant behavior, discontinuing the reinforcement of responding (or the reduction in responding that it produces). In negative reinforcement (escape and avoidance), extinction has often referred to the discontinuation of aversive stimuli, although the term applies more appropriately to discontinuing the consequences of responding, so that aversive stimuli occur but responses no longer prevent them. The discontinuation of punishment (see RECOVERY) is rarely referred to as a variety of extinction. In respondent conditioning, extinction is presenting the CS without, or no longer in a contingent relation to, the US (or the diminution in conditioned responding that follows this operation).

Extinction gradient: following differential extinction, a gradient obtained along a continuum that contains the extinction stimulus. In one type, responding is first reinforced during several stimuli along the continuum and is then extinguished during only one of them. In another, reinforcement is correlated with one stimulus and extinction with another, but only the second is represented on the continuum along which the gradient is determined (e.g., the first stimulus is a form and the second is a color, and the gradient is determined for color). Cf. INHIBITORY GRADIENT.

Extrinsic reinforcer: a reinforcer that has an arbitrary relation to the responses that produce it, as when a musician plays for money rather than because the playing produces music. The term has also been applied to stimuli presumed to function as reinforcers because their function has been instructed, as when children are told that it is important to earn good grades; despite their label, such stimuli are often ineffective reinforcers. Cf. INTRINSIC REINFORCER, CONDITIONAL REINFORCER.

F

Facilitation: an occasional synonym for *potentiation*.

Fading: a procedure for transferring control of responding from one stimulus or set of stimuli to another by gradually removing one while the other is gradually introduced. Stimuli may be faded in or out. For example, once a pigeon discriminates key colors, the discrimination may be transferred to line orientation by maintaining differential reinforcement while gradually fading out color intensity and fading in line intensity. Cf. SHAPING.

Feature-positive stimulus: in a successive discrimination between reinforcement and extinction stimuli, a stimulus property present only during reinforcement components, as when, in discrete trials with pigeons, a star appears somewhere within a grid of squares during reinforcement trials but the grid appears without the star during extinction trials. Stimulus control is more easily produced when such stimuli are correlated with reinforcement (*feature positive*) than when they are correlated with extinction (*feature negative*). Cf. SIGN-TRACKING.

Feedback: roughly, a stimulus or stimulus property correlated with or produced by the organism's own behavior. The stimulus may in turn change the behavior, which again changes the stimulus, and so on, as when rumble strips on either side of a traffic lane produce vibrations whenever the driver moves too far from the lane's midline. Feedback may be discrete, as in this example, or continuous, as when the driver's speedometer reading varies with changing pressure on the gas pedal. Mathematical relations between the behavior and the stimuli are called *feedback functions*.

FI: fixed interval. See INTERVAL SCHEDULE.

Fixed action pattern: an ethological term for a sequence of responses, usually but not necessarily produced by a RELEASER, the consistent patterning of which cannot be attributed to an operant chain. One reason for distinguishing fixed action patterns from reflex relations is that the pattern is typically released as an all-or-none event; unlike the relation between stimulus and response in a reflex relation, it does not vary in magnitude with changes in the magnitude of the releaser.

When the stimuli that elicit or set the occasion for a fixed action pattern are absent, their presentation, and thereby an opportunity to engage in the fixed action pattern, may serve as a reinforcer. Cf. DISPLACEMENT ACTIVITY, VACUUM ACTIVITY.

Fixed consecutive number (FCN): a two-operandum procedure with trials initiated by responses on one operandum and terminated by a changeover to the other. The changeover is reinforced if at least *n* responses preceded it, as in reinforcing a pigeon's left peck only if at least 10 right pecks preceded the changeover to the left.

Fixed-interval schedule: see INTERVAL SCHEDULE.

Fixed-ratio schedule: see RATIO SCHEDULE.

Fixed-time schedule: see TIME SCHEDULE.

Fixity, functional: see FUNCTIONAL FIXITY.

Flashbulb memory: the detailed remembering of the context of a surprising and significant event in one's life,, such as where one was on hearing the news of the Challenger disaster.

Fluency: accurate performance that occurs at high rates and/or with short latencies, and that is well retained after long periods without practice. Skills that become fluent, such as mastery in mathematics, may be more likely than skills not so well mastered to combine with other behavior in novel and productive ways. Cf. ADDUCTION.

Foraging: searching for food. Foraging in natural habitats has been treated as a chain that includes search, prey identification, prey capture and handling and/or consumption of prey, with concurrent or concurrent-chain performances treated as analogous to parts of this chain. For example, the foraging patterns according to which organisms switch from partially depleted patches of food to fresh ones can be characterized in terms of strategies examined within concurrent performances, such as *matching*, *momentary maximizing* and *optimizing*.

Forced choice: see FREE CHOICE.

Formal verbal classes: see DICTATION-TAKING, ECHOIC BEHAVIOR, TEXTUAL BEHAVIOR, TRANSCRIPTION.

FR: fixed ratio. See RATIO SCHEDULE.

Fractional escape: see ESCAPE.

Free choice: the availability of two or more concurrent operants even if one is consistently chosen over the other. With only one operant available, choice is said to be *forced*, as when one of the two arms in a T-maze is blocked.

Free feeding weight: the stable weight maintained by a mature organism with unlimited access to food and water. A percentage of this weight (e.g., 80%) may serve as a criterion for a level of deprivation; 80% of free feeding exceeds the weight at which many wild animals maintain themselves.

Free operant: see OPERANT.

Free recall: a verbal learning procedure in which the learner recalls the items of a list, usually after a single presentation, without regard to the original order of items.

Free reinforcer: a reinforcer delivered without reference to the organism's behavior. This arrangement has been variously called *noncontingent reinforcement* or *NCR* (sometimes referred to as a *non sequitur*), and *response-independent reinforcement*. Although NCR is a frequent usage, free reinforcer has some precedent and is often conveniently brief and unambiguous. Cf. SUPERSTITION, TIME SCHEDULE

Frequency of reinforcement: total reinforcers over a fixed time; occasionally, total reinforcers over a session of variable duration, over a fixed number of responses or, in a trial procedure, over a fixed number of trials). Cf. FREQUENCY OF RESPONDING, RATE OF REINFORCEMENT.

Frequency of responding: total responses over a fixed time, over a session of variable duration or, in a trial procedure, over a fixed number of trials. Cf. FREQUENCY OF REINFORCEMENT, RATE OF RESPONDING.

Frustration: any operation that reduces an organism's opportunities for highly probable responses, or the consequences of such operations, especially including emotional behavior, aggression or escape from correlated stimuli. The term is most commonly applied to extinction after food reinforcement, which eliminates the opportunity to eat; in this usage it is a label for some side effects of extinction.

FT: fixed time. See TIME SCHEDULE.

Functional analysis: an analysis in terms of behavioral functions (effects of responses); alternatively, an analysis in terms of functional relations. For example, the production of pupillary constriction by light might be discussed as a pupillary reflex, but a functional analysis deals with it as a transition from one point to another on a continuous mathematical function relating pupillary diameter to light intensity.

Functional class: a class in which members have common behavioral functions, either produced by similar histories or acquired through emergent relations. If two stimuli are members of a functional class, then the behavior occasioned by one will also be occasioned by the other; such stimuli are *functionally equivalent*. Cf. EQUIVALENCE CLASS; see also EQUIVALENCE RELATION, OPERANT, STIMULUS.

Functional communication training (FCT): the replacement of disruptive problem behavior with new behavior that allows the child to obtain reinforcers from caregivers in alternative, nondisruptive ways. It must not be confused with *facilitated communication*, a procedure used with autism that has been shown to be an unproductive and sometimes destructive variation on the Clever Hans effect.

Functional fixity: problem-solving behavior in which the usual function of a tool, object or material makes it less likely that the solver will use it effectively in a novel way.

Functional relation: a mathematical function appealed to by a FUNCTIONAL ANALYSIS.

Functional stimulus: the properties of a stimulus that control behavior, as opposed to the properties of the *nominal stimulus*. For example, for a pigeon attending to the color but not the form of a green circle, the functional stimulus is just green even though the nominal stimulus is a green circle.

Fuzzy set: see PROBABILISTIC STIMULUS CLASS.

G

g: gram (abbreviation).

Gene: a hereditary unit, inferred during early studies of genetics, and later with arbitrary boundaries as a sequence on a DNA molecule. It is now more precisely defined as a portion of DNA that codes for an individual protein (Carroll, 2006, p. 74).

Generalization: the spread of the effects of reinforcement, or other operations such as extinction or punishment, during one stimulus to other stimuli differing from the original along one or more dimensions. To the extent that responding is similar during two different stimuli, the organism is said to *generalize* between them; the stimuli are said to be *generalized*. If responding is identical during different stimuli, generalization between them is said to be complete; this outcome may also be described as the absence of DISCRIMINATION. Cf. ATTENTION, INDUCTION, STIMULUS.

Generalization gradient: a gradient obtained after reinforcement correlated with a single stimulus (occasionally two or more stimuli, in studies of the summation of gradients), when no discrimination has been trained between this and other stimuli along the continuum of the gradient.

Generalized imitation: see HIGHER-ORDER CLASS OF BEHAVIOR.

Generalized reinforcer: a conditional reinforcer based on several primary reinforcers. It is more likely to remain effective across different motivational or establishing operations than a conditional reinforcer based on just one primary reinforcer. Money is often offered as an example of a generalized reinforcer of human behavior.

Genotype: the genetic makeup of an organism, as opposed to its phenotype, its makeup as produced by the course of its development. The development of two or more phenotypes from a single genotype, as in the progression from tadpole to frog or from caterpillar to butterfly, is called *polyphenism.*

Goal gradient: systematic changes in responding with changes in an organism's spatial or temporal separation from a reinforcer, such as changes in running speed as a rat approaches the goalbox of a maze.

Go-no go discrimination: usually, a successive discrimination using trials with reinforcement in the presence of one stimulus (*go*) and extinction in presence of the other (*no go*).

Gradient: a measure of responding during different stimuli as a function of their location along a continuum (cf. STIMULUS). Gradients are usually determined by presenting stimuli successively but in irregular order in extinction. The slope or steepness of a gradient is determined by how much responding changes from one point on the continuum to another; the larger the change, the steeper the gradient. The case in which responding does not change is usually called a *flat* gradient, although it may also be called the absence of a gradient. See specific cases: EXTINCTION GRADIENT, GENERALIZATION GRADIENT, INHIBITORY GRADIENT, POSTDISCRIMINATION GRADIENT.

Grain: a reinforcer effective with food-deprived pigeons; also a characteristic of cumulative records (see RATE OF RESPONDING).

Grammars: descriptions of the structural or syntactic properties of verbal behavior. *Phrase-structure* grammars describe structures in terms of relations among sentence *constituents* (e.g., noun and verb phrases); *transformational* grammars describe them in terms of relations among different sentences (e.g., active and passive voice). The two are complementary.

H

Habit: recurrent behavior or behavior that is highly resistant to change. The term is mainly colloquial; it often occurs in conjunction with evaluations of the behavior (good habits versus bad habits) but without specifying the contingencies that maintain it.

Habituation: a reduction, over repeated presentations, in the respondent behavior elicited by a stimulus. Cf. ADAPTATION, POTENTIATION.

Hearsal: the first response to verbal stimuli in verbal learning tasks, as opposed to subsequent responses during continued REHEARSAL.

Helplessness, learned: see LEARNED HELPLESSNESS.

Hierarchical organization: the nesting of some classes of behavior within others. Cf. HIGHER-ORDER CLASS OF BEHAVIOR.

Hierarchy, of behavior or responses: a ranking of response classes based on their relative probabilities. A more probable class is said to be higher in the response hierarchy.

Higher-order class of behavior: an operant class that includes within it other classes that can themselves function as operants, as when generalized imitation includes all the component imitations that could be separately reinforced. Higher-order classes may be a source of novel behavior, as in the imitation of modeled behavior that the imitator hadn't seen before. Contingencies operate differently on the higher-order class than on the classes that are its components. For example, if all instances of imitation are reinforced except those within one component class, such as jumping whenever the model jumps, that class may change with the higher-order class rather than with the contingencies arranged for it, as when imitations of jumping do not extinguish even though they are no longer reinforced. Control by the contingencies arranged for the higher-order class defines membership in the class; the component classes are sometimes said to be *insensitive* to the lower-order contingencies arranged for them. A higher-order class may be called a *generalized* class, in that contingencies arranged for some of its components generalize to all the others, including novel instances (cf. EMERGENCE). Generalized matching and verbally governed behavior are examples of higher-order classes.

Higher-order conditioning: respondent conditioning in which the stimulus that functions as the US in producing one conditioned reflex is itself the CS of another.

Higher-order schedule: a schedule that reinforces a higher-operant consisting of completion of a schedule requirement. For example, with FR 10 reinforced according to an FI 50-s schedule, every 10th response that occurs at least 50 s after the last reinforcer is reinforced; the FR is the first-order schedule and the FI the second-order schedule. Such schedules often include a stimulus presented on each completion of the first-order schedule (in the example, flashing a light after every 10th response). The notation for such schedules includes the first-order schedule and the stimulus it produces in parentheses: FI 50-s (FR 10: stimulus). A *percentage-reinforcement* schedule is a higher-order schedule in which the second-order schedule is an RR schedule.

Hill-climbing: changes in behavior that come about because the organism moves gradually from one performance to another that produces a somewhat higher reinforcement rate. The metaphorical paths made possible by different contingencies determine whether the organism approaches *optimizing* or gets trapped in a valley or a dead end. Cf. MAXIMIZING, SHAPING.

History: conditions that an organism has been exposed to and its performances under them; often an abbreviation for *experimental history*, simply because experimental organisms are rarely observed continuously throughout their lives. History is particularly important when its effects are irreversible or only slowly reversible.

Houselight: see CHAMBER.

I

Iconic memory: see REMEMBERING.

Identity matching: see MATCHING-TO-SAMPLE.

Imitation: behavior that duplicates some properties of the behavior of a model. Imitation needn't involve a matching of stimulus features. For example, when one child imitates the raised hand of another, the felt position of the child's own limb has different stimulus dimensions than the seen position of the other child's. Cf. HIGHER-ORDER CLASS OF BEHAVIOR.

Immediate memory: usually, short-term memory. See REMEMBERING.

Implicit learning: in human learning, contingency-shaped learning, as in learning to speak grammatically even though one cannot state the grammatical rules.

Implicit memory: remembering demonstrated by the effect of an item on other behavior rather than by its recall, as when a priming stimulus enhances a learner's later recognition of a semantically related word even though the learner cannot report what the priming stimulus was.

Imprinted stimulus: a stimulus that, by virtue of the conditions of its presentation, has become effective as a reinforcer. Imprinting is noted primarily in some bird species (e.g., ducks) and usually occurs within a few days of hatching. Conditions that affect imprinting include movement of the stimulus and time spent in its presence. Imprinting does not occur as easily with older ducklings because they avoid novel stimuli and therefore do not spend enough time in their presence.

Impulsiveness or **impulsivity:** see SELF-CONTROL.

Incentive: discriminative effects of reinforcing stimuli, as when the smell of food makes responses reinforced by food more likely; occasionally, a stimulus that changes the reinforcing or punishing status of other stimuli. A verbal response that has such effects is sometimes called an *augmenting* stimulus. Cf. ESTABLISHING OPERATION.

Incidental chaining or **incidental reinforcement:** see SUPERSTITION.

Incidental learning: human learning that occurs in the absence of instructions or consequences, usually contrasted with *intentional learning.*

Induction: the spread of the effects of reinforcement to responses outside the limits of an operant class (sometimes called *response generalization*). This phenomenon is essential to shaping because without it new responses more closely approximating some final form might never be emitted. For example, reinforcement of a 10-N key peck may be followed by the first instance of a 15-N peck (cf. SHAPING). With discriminated operants, induction occasionally refers to the spread of the effects of reinforcement to stimuli other than those defining the operant class (as when, after extinction during green and red, reinstating reinforcement during green produces both responding during green and a transient increase in responding during red; cf. GENERALIZATION).

Information: strictly, the reduction in uncertainty provided by a stimulus, as quantified in *bits*, the number of binary decisions needed to specify the stimulus. One bit specifies 2 alternatives, two bits 4, three bits 8, and so on in increasing powers of 2. The term more often appears in its nontechnical sense, as when applied to *information processing*, the value of nontechnical usages is diminished by the finding that reinforcers rather than information maintain observing behavior.

Informative stimulus: a predictive stimulus; a discriminative stimulus though not necessarily a conditional reinforcer. For example, a stimulus correlated with differential punishment that is superimposed on ongoing reinforced behavior is informative, but its onset does not ordinarily maintain observing responses.

Inhibition: a process inferred from a response decrement. The term, extended to behavior by analogy to physiological usage, is appropriate only when it has been shown that the decrement is produced by an increment in something else. For example, if reinforcing one response reduces the rate of another, reinforcement of the first response may be said to inhibit the second one. When the reinforcers maintaining each of two response act upon each other, their interaction may be characterized as reciprocal inhibition; these interactions parallel those seen in sensory systems.

The term has sometimes been extended to accounts of the extinction process, in part because extinction may be accompanied by increments in other responses, such as behavior characterized as emotional. Such accounts have typically been unsupported by demonstrations that the increments produce the decrement rather than simply accompany it. They may be misleading when they do not distinguish between conditions that reduce responding and those failing to maintain it.

Inhibitory gradient: an extinction gradient in which responding increases along some continuum as the distance from a stimulus previously correlated with extinction increases. This is sometimes taken to indicate that the extinction stimulus reduces response rate rather than just failing to maintain it; it can also be interpreted as generalization from the reinforcement stimulus along the entire continuum, which the effects of extinction concentrated around the stimulus correlated with nonreinforcement. Cf. EXTINCTION GRADIENT, INHIBITION.

Inhomogeneous data: data derived from more than one type of performance and that misrepresent the performances from which they were derived when summarized statistically, as when an avoidance schedule produces both moderate response rates and bursts of high-rate responding after shock, so that the average rate does not represent either contribution to overall rate.

Initial links: see CONCURRENT-CHAIN SCHEDULES.

Innate behavior: see SPECIES-SPECIFIC BEHAVIOR.

Insensitivity to contingencies. see HIGHER-ORDER CLASS OF BEHAVIOR.

Insight: the sudden solution of a problem, especially in contrast to gradual "trial-and-error" learning. This now nontechnical term is mainly of historical interest.

Instinctive behavior: see SPECIES-SPECIFIC BEHAVIOR.

Instructional stimulus, instruction: in verbal settings, a verbal antecedent of either verbal or nonverbal behavior; in nonverbal settings, a conditional discriminative stimulus. Occasionally a simple discriminative stimulus is also said to have instructional functions. Cf. VERBALLY GOVERNED BEHAVIOR.

Instrumental behavior: see OPERANT BEHAVIOR.

Intentional learning: the opposite of *incidental learning*.

Interdependent schedules: schedules in which the operation of one depends on some property of the other. For example, in one version of interdependent concurrent VI VI schedules, each VI arranges setups only during runs of responding on the other.

Interdimensional: between or across dimensions.

Interference: see PROACTION, RETROACTION; cf. INHIBITION.

Interim behavior: varying responding that occurs, usually early or midway rather than late within interstimulus intervals, in superstition procedures or temporal conditioning. Cf. SUPERSTITION, TERMINAL BEHAVIOR.

Interlocking schedule: a schedule in which time, number and/or IRT requirements vary together according to some function, as with interlocking intervals and ratios scheduled so that the number of responses that will produce a reinforcer decreases linearly as time passes since the last one.

Intermittent reinforcement: reinforcement of some but not all responses. See specific schedules.

Interoceptive stimulus: a stimulus inside the organism. The stimulus may be presented from outside, as when an experimenter passes electric current through a brain area, or it may be produced by the organism itself, as when responses produce proprioceptive stimulation on the basis of which the organism discriminates among different movements. With self-produced stimulation, the stimuli and their discriminative functions are usually inferred rather than demonstrated.

Interpretation: see CONTROL.

Interresponse time (IRT): the time between two responses or, more strictly, if response duration is appreciable, from response onset to onset, or from response offset to offset. The time from a reinforcer to the next response is a latency and not an IRT, even if the reinforcer is response-produced. Reinforcing the response that ends an IRT is said to reinforce that IRT. An *IRT distribution* summarizes the temporal spacing of the responses making up a response rate; it does not show sequential patterning. In assessing probabilities of different IRT classes, calculating proportions of IRTs falling into a class may be misleading because short IRTs reduce the organism's opportunity to produce longer ones. For this reason, conditional probabilities, *IRTs per opportunity* (*IRTs/Op*), are sometimes calculated: the probability of IRTs in a class interval, given enough time elapsed since the last response to permit an IRT to end in that class interval. For example, if 80 of 100 IRTs were less than 1 s, 10 were 1 but less than 2 s, and the remaining 10 were 2 or more s, then the organism had only 20 opportunities to complete the 10 IRTs of 1 but less than 2 s and the conditional probability for this IRT class was .5 (10/20). Cf. CONDITIONAL PROBABILITY, RATE OF RESPONDING, DIFFERENTIAL-REINFORCEMENT SCHEDULES.

Intertrial interval (ITI): see TRIAL.

Interval schedule: a schedule in which some minimum time must elapse before a response is reinforced; early responses have no effect. The time is measured from some event, usually a stimulus onset or the last reinforcer; an alternate method times each interval from the end of the last one, without regard to the time from the end of the last interval to the reinforced response. In *fixed-interval* (*FI*) schedules, the time is constant from one interval to the next, and performance is characterized by a pause after the reinforcer followed by a gradual or an abrupt transition to a moderate response rate. In *variable-interval* (*VI*) schedules; the time varies from one reinforcer to the next; compared with FI schedules, the response rate is relatively constant between reinforcers. Interval schedules are usually identified by average interval. For example, FI 50-s arranges one reinforcer per 50 s.

Historically, VI schedules used intervals selected in irregular order from a set of intervals, often described by a mathematical progression (e.g., arithmetic or geometric). Current practice favors schedules with a constant reinforcement probability over time within the interval, with probability measured by *reinforcers per opportunity* or *Rf/Op*, the probability that a response will be reinforced at a given time in an interval given that the organism has reached that time (cf. INTERRESPONSE TIME). Such conditions are met by a type of VI schedule called *random interval* (*RI*), which arranges a setup, or makes the next response eligible to produce a reinforcer, with a fixed probability every *t* s. In RI schedules, the average interval equals *t* divided by the probability; for example, arranging a setup once per second with a probability of .02 produces RI 50-s. In one version, the schedule stops operating after a setup until the scheduled reinforcer is produced, so that low response rates make the obtained reinforcement rate lower than what had been scheduled; in another, the schedule continues and successive setups accumulate, so that obtained and scheduled reinforcement rates remain about equal even with low response rates (cf. SETUP).

Intradimensional: within a dimension.

Intraverbal: a verbal response occasioned by other verbal stimuli, where the relations between stimuli and responses are arbitrary ones established by the verbal community. Intraverbal responding can work in one of three ways: by strengthening other members within a thematic class, as when the words *apple*, *orange* and *fruit* raise the likelihood of saying *banana*; by providing the first link of a chain of one or more verbal responses, as when *7 times 9* sets the occasion for the response *63*; or by initiating an integrated verbal sequence or chunk, as in reciting numbers or the alphabet or a well-learned verse. Either the speaker or someone else may provide the verbal stimulus. In most cases intraverbals are not autoclitic because they do not require discriminations of one's own behavior.

Intrinsic reinforcer: a reinforcer that is naturally related to the responses that produce it, as when a musician plays not for money but because the playing produces music. Cf. EXTRINSIC REINFORCER.

IRT: see INTERRESPONSE TIME.

ITI: intertrial interval. See TRIAL.

J

Joint control: a change in stimulus control that depends on the correspondence of responses occasioned at the same time by different stimuli, as when a word is heard or spoken at the same time it is seen. Joint control is "a change in stimulus control that occurs when a response topography evoked by one stimulus ... and preserved by rehearsal, is emitted under the additional (and thus joint) control of a second stimulus" (Lowenkron, 1998, p. 332). For example, while reciting a lottery ticket number, someone reads the list of winning numbers; the person says "I won" only if the recited number and the number read from the list correspond. This relation must be discriminated, so joint control may be regarded as a variety of AUTOCLITIC verbal behavior.

Joint attention: what has been shaped when a child looks not at the caregiver but what the caregiver is pointing to or looking at.

Jumping stand: an apparatus used to study discrimination, especially with rats. The rat is forced to jump from a platform to one of two doors on which stimuli are displayed. One door is unlocked and by jumping to it the rat gains access to a reinforcer behind it; if the rat jumps to the other door, which is locked, it falls into a net below the doors.

K

Key: see OPERANDUM.

Kinesis (plural: kineses): undirected movement that depends on stimulus magnitude, as when the random movements of insect larvae increase with light and stop when they reach the dark.

Kinesthetic stimulus: see PROPRIOCEPTIVE STIMULUS.

Knowledge of results (KOR): A kind of feedback, usually verbal, given during human performance in verbal learning or other tasks.

L

Language: the practices shared by the members of a verbal community, including consistencies of vocabulary and grammar. Cf. LINGUISTICS.

Language development: the emergence of language in the individual. Much of the controversy about language development revolves around assumptions about the respective phylogenic and ontogenic contributions to it.

Language relativity: the dependence of the behavior within a verbal community, both verbal and nonverbal, on the verbal discriminations incorporated within its language.

Language universals: the structural features common to all human languages, especially if they have phylogenic sources. If there are such features, what they are is controversial; furthermore, some may be spandrels.

Lashley jumping stand: see JUMPING STAND.

Latency: the time from an event, usually a stimulus onset, to a response.

Latent learning: see LEARNING.

Law of effect: Thorndike's classic statement of the principle of reinforcement and, in one version, punishment. Reinforcers and punishers were called satisfying and annoying states of affairs that an organism tended respectively to maintain or renew and to put an end to or avoid. The *strong* law included both cases; the *weak* law deleted the effects of annoyers or punishers.

Law of exercise: the statement, in early accounts of learning, that repetition of a response contributes to its strength. The law has also been expressed in terms of effects of use and disuse. A more contemporary version states that the elicitation of a response may increase the likelihood of its emission.

Learned helplessness: a retardation in the acquisition of escape or avoidance responding produced by a history in which responding during aversive stimuli has had no consequences.

Learning: roughly, acquisition, or the process by which behavior is added to an organism's repertory; a relatively permanent change in behavior. The term has been used in so many different ways in both technical and colloquial vocabularies that it is of limited usefulness. Decisions about whether learning has occurred and about what has been learned sometimes depend on what the experimenter looks at. Latent learning offers an example. A rat explores a maze and the results of its exploration are assessed later when food reinforcers are available for the first time at the end of its run through the maze. Latent learning is said to have occurred if the rat negotiates the maze more rapidly and/or accurately than if it hadn't explored. The difficulty is that exploring the maze involves other contingencies (e.g., which turns lead where); these contingencies act on behavior but their effects are harder to get at than those involving food reinforcers. Cf. ACQUISITION, PERFORMANCE.

Learning set or **learning-to-learn:** a case of transfer in which, on the basis of similar relations among stimuli in a sequence of discrimination problems, accuracy in later ones improves more rapidly over trials than in the earlier problems, (perhaps to a point at which correct responses occur on the first trials of a new problem. Cf. HIGHER-ORDER CLASS OF BEHAVIOR, TRANSFER.

Level of processing: the abstractness or richness of coding. For example, encoding the number 2001 as a year or as the title of a film involves deeper levels of processing than encoding it as an arbitrary sequence of four digits.

Lever: see OPERANDUM.

LH: see LIMITED HOLD.

Limited hold (LH): termination of reinforcer availability if the response to be reinforced doesn't occur soon enough. For example, in FI 100-s with a 10-s limited hold, the first response between 100 and 110 s after the start of the interval is reinforced; if there is no response during that time, the interval ends without a reinforcer.

Linguistics: the study of language, usually divided into the topics of *syntax* or grammatical structure, *semantics* or meaning, *phonology*, or the structure of speech sounds, and *pragmatics* or the functions of language. *Psycholinguistics* is a branch of psychology concerned with demonstrating the psychological reality of linguistic categories and concepts. Cf. PSYCHOLOGICAL REALITY.

Link: a response in a chain or a component in a chained schedule.

Linking: see COUPLING.

Local rate: see RATE OF RESPONDING.

Long-term memory (LTM): see REMEMBERING.

LTM: long-term memory. See REMEMBERING.

M

Maintenance: continuation of the conditions that generated a performance. As a subject matter, maintained performance differs from but is compatible with acquisition. For example, research concerned with effects of schedule parameters on performance does not really begin until acquisition is completed. Cf. STEADY STATE.

Maintenance rehearsal: see REHEARSAL.

Mand: reinforcement as it enters into verbal behavior; a verbal response that specifies its reinforcer. In human verbal behavior, manding is usually a higher-order class, in the sense that a newly acquired tact can be incorporated into a novel mand, as when a child asks for a toy upon learning its name.

Manipulandum: see OPERANDUM.

Massed practice: uninterrupted activity on a task, as in cramming for an exam. Massed practice is usually less effective than distributed practice.

Matching: in performances involving concurrent operants, distributing responses so that the relative response rate of each roughly matches its relative reinforcement rate. See MATCHING LAW; cf. MAXIMIZING, MELIORATION, OPTIMIZING.

Matching law: a quantitative formulation stating that the relative rates of different responses tend to equal the relative reinforcement rates they produce. The *generalized matching law* summarizes this relation in an equation in which relative response rate equals a constant multiplied by the relative reinforcement rate raised to a power. The constant takes into account units of measurement and includes *bias*. For example, one response might call for a larger constant than another that is more effortful; the performance is described as *undermatching* when the exponent (the power to which the function is raised) is less than 1 and *overmatching* when it is greater than 1.

Matching-to-sample: a simultaneous conditional discrimination procedure, or the performance maintained by it. As it is typically arranged for pigeons, a *sample* stimulus is presented on the middle key of three keys. A peck on the sample turns on *comparison* stimuli on two side keys. A peck on the matching side key is reinforced, perhaps according to a schedule; a peck on the other side key may produce timeout or invoke a *correction procedure*. When the criterion for a match is physical correspondence, as when a pigeon must peck a green comparison given a green sample and a red comparison given a red one, the procedure is sometimes called *identity matching*, though accurate matches may be based on features other than the identity relation, such as stimulus configurations. When matches are based on arbitrary relations, as when a pigeon must peck a circle given a green sample and a triangle given a red one, the procedure is called *arbitrary matching*. *Symbolic matching*, an alternative terminology, has the disadvantage of suggesting that the samples and comparisons have other functions besides those involved in the matching-to-sample procedure.

These arrangements allow numerous variations: *many-to-one matching*, where one set of sample stimuli occasions just one comparison response while a second set occasions a different one (e.g., samples consisting of all letters of the alphabet, with one comparison correct when the letter is italic and the other when the letter is boldface); *one-to-many matching*, where one set of comparison responses is occasioned by one sample stimulus while a second set is occasioned by another (e.g., samples consisting of the words *italic* or **boldface**, where the correct comparison is the one with a matching format); number of available comparisons; number of sample and comparison stimuli; and so on. Cf. CONDITIONAL DISCRIMINATION, ODDITY PROCEDURE.

Maximizing: given two or more responses, emitting the one with the higher probability of reinforcement. If reinforcement probabilities change from moment to moment and responding follows the one currently highest, the maximizing is said to be *momentary maximizing*. Measuring matching requires a population of responses whereas a single response can demonstrate maximizing. Cf. HILL-CLIMBING, MATCHING, OPTIMIZING.

Maze: an apparatus through which an organism locomotes, usually from a startbox to a goalbox containing some reinforcer such as food, and often including alternative paths that divide at *choicepoints* and some of which end in *blind alleys* or *culs-de-sac*. Mazes come in a variety of configurations, including T-mazes or Y-mazes with a single choicepoint, mazes with a single sequence of choices between blind alleys and a continuing path, and radial mazes with paths arranged like the spokes of a wheel.

Meaning: in verbal behavior, a response to verbal stimuli; or the defining properties of classes, usually including verbal components, in which the members can serve as either stimuli or responses.

Mediating behavior: behavior that occurs in a consistent relation to reinforced behavior and that, though reinforcers are not explicitly arranged for it, is maintained because it makes reinforcers more likely. For example, a stereotyped pattern of drinking is said to mediate spaced responding if, when the pattern is maintained, the next response is more likely to be late enough after the last response that it will be reinforced; or, two different postures held after one or another sample stimulus are said to mediate delayed matching-to-sample if a correct match is more likely when the organism has held the posture since the sample was presented, as when it leans to the right after a red but not a green sample and later is more likely to respond to red rather than green if still leaning to the right. Cf. COLLATERAL BEHAVIOR, SUPERSTITION.

Mediation: the contribution of intervening behavior to the relation between other events, as when coding mediates between the presentation of an item and its recall. Consider a mediational account of a pigeon experiment in which (i) response-independent food is delivered during green but not red; then (ii) pecking is shaped and maintained with food reinforcers with the key white; finally (iii) green and red alternate during extinction and for a while the pigeon pecks more during green, formerly correlated with food, than during red, even though pecking had never been reinforced during either (Morse & Skinner, 1958). In (i) food presumably generated incidental responding occasioned by both green and food; in (ii), peck-produced food generates the same behavior, followed by pecks that in turn produce food. With green and red reinstated in (iii), more pecks occur during green because green but not red occasions responding that had often been followed by reinforced key pecks in (ii). Cf. MEDIATING BEHAVIOR.

Melioration: allocating time to two or more response classes so all local reinforcement rates are equal. For example, assume a pigeon whose pecks in a changeover-key procedure are maintained by concurrent VI 20-s VI 60-s schedules. In an hour, the former provides about 180 reinforcers and the latter about 60, but if the pigeon allocates 45 min to the VI 20-s schedule and 15 min to the VI 60-s schedule both local reinforcement rates will equal about 4 per min (180 in 45 min and 60 in 15 min). Cf. MATCHING, MAXIMIZING, OPTIMIZING.

Memetic selection: see CULTURAL SELECTION.

Memory: see REMEMBERING.

Memory search: see SEARCH.

Memory span: the number of items that can be remembered after a single presentation, given that they are not coded or rehearsed.

Metacognition: differentiation and discrimination of one's own cognitive processes, as in shifting attention among tasks or distinguishing between seeing something and just imagining it.

Metamemory: differentiation and discrimination of one's own remembering, as when keeping track of a constantly changing list of items in *running* or *working* memory, or judging whether some material just studied will be remembered).

Metaphor: the extension of concrete terms to complex and/or abstract events or relations for which relevant verbal responses are otherwise unavailable, as when pain is described not by how it feels but rather by the properties of objects that can produce it (e.g., pains as sharp or stinging or pounding or dull).

Metastability: see STABILITY.

Metathetic stimuli: see STIMULUS CONTINUUM.

Microanalysis: see MOLAR AND MOLECULAR ANALYSES.

min: minute (abbreviation).

Misbehavior: a nontechnical term sometimes used to refer to the intrusion of behavior with phylogenic origins into ongoing operant behavior. For example, raccoons ordinarily rub and wash food before eating it; if food is used to reinforce their deposit into a container of objects they've picked up, they may begin to rub the objects together instead of releasing them into the container. The procedure raises the probability of rubbing, so the effect illustrates the relativity of reinforcers. Intrusions in the opposite direction, as when food-reinforced behavior intrudes into a fixed action pattern, are rarely referred to as misbehavior.

Mixed (mix) schedule: a compound schedule with two or more component schedules operating in alternation, all during the same stimulus. Occasionally, a VI or VR schedule with a limited number of schedule values is called a mixed schedule. For example, a VR schedule that randomly alternates between FR 10 and FR 20 may be called mixed FR 10 FR 20. Cf. MULTIPLE SCHEDULE.

Mnemonics: techniques for enhancing remembering.

Modeling: providing behavior to be imitated. Cf. IMITATION.

Modules: see COMPARTMENTS.

Molar and molecular analyses: analyses distinguished by the level of detail in the data they consider. Molar analyses consider overall measures such as average response rates over sessions, whereas molecular analyses break such measures down into components such as the distribution of the IRTs that make up a response rate. Many levels of analysis are possible, so molar or molecular are sometimes defined relative to each other. Both nevertheless rely on data sampled over some time and should be distinguished from *microanalysis*, which proceeds at the level of individual stimuli and responses.

Momentary maximizing: see MAXIMIZING.

Motivating operation: any operation that changes the status of a stimulus as a reinforcer or punisher: deprivation, satiation, procedures that establish formerly neutral stimuli as conditional reinforcers or as conditioned aversive stimuli, and stimulus presentations that change the reinforcing or punishing status of other stimuli, as when an already available screwdriver becomes a reinforcer in the presence of a screw that needs tightening. These terms and *establishing operation* have become essentially interchangeable, and though *motivating operation* is somewhat favored in this text it has not been used exclusively.

Motor programs: coordinations that do not depend on response feedback. For example, movements of lungs, vocal chords, tongue and lips must be initiated at different times in producing phonemes; thus, their coordination must be organized before the sound begins. Such coordinations cannot be based on chaining.

Movement: cf. BEHAVIOR.

ms: a millisecond or thousandth of a second (abbreviation).

Multiple causation of behavior: the determination of behavior by two or more variables acting at the same time. Behavior is always determined by multiple variables; some may be more important than others. The aim of *behavior analysis* is to examine one at a time the multiple factors controlling behavior.

Multiple (mult) schedule: a compound schedule in which two or more component schedules alternate, each during a different stimulus, and possibly irregularly. Alternation of the components is usually arranged after reinforcers or after fixed or variable time periods. The equivalent arrangement with the same stimulus during each component is a *mixed* schedule.

N

***n*:** usually, number.

N: Newtons, a unit of force.

Naming: a higher-order class that involves arbitrary stimulus classes (things or events with particular names) and corresponding arbitrary verbal topographies (the words that serve as their names) in a bidirectional relationship. Prerequisites for naming include at least three components: (i) listener behavior, in looking for things and pointing based on what has been said (cf. JOINT ATTENTION); (ii) echoic behavior, in repeating names when they are spoken; and (iii) tacting, in saying the names given the objects. Naming is generated from the ordinary interactions between children and their caregivers. Once available as a higher-order class, naming allows expansions of vocabulary in which the introduction of new words in particular functional relations, as in tacting, involves those words in a range of other emergent functions.

Natural concept: a discriminative class produced through presentations of complex natural stimuli, as when a pigeon discriminates between pictures with and pictures without trees in them. Cf. PROBABILISTIC STIMULUS CLASS.

Natural reinforcer: sometimes used in place of *primary reinforcer* or *intrinsic reinforcer*. The relativity of reinforcers limits the usefulness of the term: cf. REINFORCEMENT.

Natural selection: at the phylogenic level, the Darwinian account of EVOLUTION as the selection of members of a population over generations. Different features survive in a population as a result of interactions between the available range of genetic variations in the population and the properties of evolutionary environments. Natural selection can also occur at ontogenic or cultural levels; cf. ARTIFICIAL SELECTION.

NCR: Noncontingent reinforcer. See FREE REINFORCER.

Negative automaintenance: see AUTOMAINTENANCE.

Negative reinforcement: see REINFORCEMENT.

Negative stimulus: see DISCRIMINATIVE STIMULUS.

Negative transfer: see TRANSFER.

Neophobia: avoidance of novel stimuli, especially of new foods.

Nominal operant or **nominal stimulus:** see FUNCTIONAL STIMULUS, OPERANT.

Noncontingent reinforcer: See FREE REINFORCER.

Nonsense syllable: an arbitrary sequence of letters, usually a *CVC trigram*, that is not a word. Nevertheless, nonsense syllables vary in meaningfulness; for example, the resemblance of LUQ to LUCK makes it more meaningful than QUL.

Novel behavior: see ADDUCTION, HIGHER-ORDER CLASS OF BEHAVIOR, SHAPING.

Noxious stimulus: often used as a synonym for aversive stimulus, but more strictly defined as a stimulus that affects pain receptors or produces tissue damage. In this strict sense, the term is useful for referring to an extensive class of stimuli without specifying behavioral consequences.

O

Observational learning: learning based on observing the responding of another organism, and/or its consequences. Observational learning does not imply imitation. For example, organisms may come to avoid aversive stimuli on seeing what happens when other organisms produce them.

Observing response: a response that produces or clarifies a discriminative stimulus and that may be maintained by the effectiveness of the stimulus as a conditional reinforcer. Some observing responses are only inferred, as when a pigeon's head movements are assumed to bring a visual stimulus into view or better focus, but conditions may be arranged to control them. For example, in matching-to-sample the pigeon may be more likely to observe the sample if a peck on the sample key is required; in a more explicit arrangement, pigeon's pecks on one key may produce stimuli correlated with components of a multiple schedule on a second key.

Occasion: an opportunity for a response or some other event, or the circumstances under which a contingency operates, as when discriminative stimuli *set the occasion* on which responses have some consequence. When a stimulus is said to *occasion* a response, the term serves as a verb and distinguishes responses emitted in the presence of discriminative stimuli from those elicited by stimuli in reflex relations.

Occasion setter: see CONTINGENCY.

Oddity procedure: a conditional discrimination procedure in which one of three or more stimuli differs from the others in some property (e.g., color) and responses to the odd one are reinforced. Versions of matching-to-sample in which responses to the comparison that does not match the sample are reinforced (mismatching) also qualify as oddity procedures. Cf. MATCHING-TO-SAMPLE.

Omission training: a version of differential reinforcement of zero behavior (see DIFFERENTIAL-REINFORCEMENT SCHEDULES): a reinforcer is delivered only if no response has occurred in a trial or within a given time. It is formally analogous to avoidance, with reinforcers substituted for aversive stimuli.

Ontogenic selection: the selection of populations of responses within the lifetime of the individual organism. See DIFFERENTIAL REINFORCEMENT, OPERANT SELECTION, SHAPING. Cf. CULTURAL SELECTION, NATURAL SELECTION, ARTIFICIAL SELECTION.

Ontogeny: the development or life history of an individual organism. Cf. PHYLOGENY.

Open economy: in operant contexts, the availability of appetitive stimuli not only as reinforcers within the session but also, independently of behavior, on a supplementary basis outside the session, as when food is provided after a session of food-reinforced responding to maintain the organism at a standard percentage of free feeding weight. Cf. CLOSED ECONOMY.

Operandum: any device operable by an organism that defines an operant class in terms of an environmental effect (*descriptive* or *nominal operant*; see OPERANT). Many operanda (plural) consist of switches, as in rats' lever presses and pigeons' key pecks, or as when a rat operates a switch by stepping off a platform. In the broadest sense, an operandum is any apparatus by means of which behavior is recorded. The term replaces an earlier one, *manipulandum*, which suggested a device that is handled. For other examples, see JUMPING STAND, WHEEL RUNNING.

Operant: a class of responses. Responses are assigned to classes because no two can be exactly the same. Special cases include the *free operant*, in which the completion of one response leaves the organism in a position to emit the next, and the discrete or constrained operant (see TRIAL). Classes defined descriptively (*descriptive* or *nominal operant*) are usually distinguished from those defined functionally (*functional operant*).

In the descriptive usage, usually for the purpose of recording responses, the class is defined in terms of its environmental effect, as when a lever press is defined by operation of a switch (see OPERANDUM). To count as a member of an operant, a response must have a certain force, topography and so on; another defining property may be the stimuli in the presence of which it occurs (see DISCRIMINATED OPERANT). The effect that defines an operant in this usage may be different from scheduled consequences of responses. For example, in a schedule, every response in the class doesn't necessarily produce a reinforcer.

In the functional usage, an operant is a class modifiable by the consequences of responses in it. It is defined by the relation between consequences and subsequent responding. According to this definition, a response class is not an operant until its modifiability has been demonstrated. In most cases, operants defined descriptively and those defined functionally include roughly the same responses. If they don't, it may be appropriate either to change recording methods or to search for the variables that are limiting the modifiability of the class. See also OPERANT BEHAVIOR.

Operant behavior: behavior that can be modified by its consequences. It is not defined by its topography and it is not limited to movement; for example, shifts in attention can be modified by their consequences, but they do not consist of movements. Cf. BEHAVIOR.

Operant behavior has also been called *instrumental* and often corresponds closely to behavior colloquially called purposive. Because of its relation to consequences, it is said to be *emitted* rather than elicited. Few responses, however, are either exclusively emitted or exclusively elicited. Many emitted responses (e.g., a pigeon's pecks) can be made more probable by certain stimuli (e.g., spots on the pigeon's key); many elicited responses can occur in the absence of typical eliciting stimuli (e.g., spontaneous salivation). Operant and respondent classes are best regarded as extremes on a continuum along which the probability varies that a stimulus will produce a response. See also OPERANT.

Operant level: the baseline level of an operant; the rate at which responses occur before they have been reinforced.

Operant selection: the selection of behavior during the lifetime of an individual organism; the modification of operant behavior by its consequences (see DIFFERENTIAL REINFORCEMENT, SHAPING). This type of selection was once called operant or instrumental conditioning.

Operation: any experimental procedure or condition (e.g., presenting a stimulus, reinforcing a response, arranging a schedule); the term is roughly interchangeable with *procedure*. The behavioral vocabulary often fails to provide separate terms for operations and their behavioral outcomes, processes. For example, *a response was reinforced* may mean that the response produced a reinforcer or that it increased in rate because it produced a reinforcer; the correct reading is usually given by context. This dual usage is common to several fundamental terms (e.g., conditioning, extinction, punishment). In this glossary, the process or outcome definitions of such terms are usually will usually follow the procedure usage. Ambiguity can be avoided by restricting such terms to operations or procedures and describing outcomes directly in terms of changes in responding: for example, *A response was reinforced and as a result its rate increased.*

Optimizing: responding that produces the maximum possible reinforcers over some extended time rather than from moment to moment, especially in concurrent schedules. Contingencies can be designed under which optimizing requires a performance different from *matching, melioration* or *momentary maximizing.*

Organism-based theory. See ATTRIBUTION THEORY.

Orienting response: in operant behavior, a response that puts an organism in a position to emit other responses or allows it to attend to a discriminative stimulus (cf. OBSERVING RESPONSE). In respondent behavior, a response elicited by initial presentations of a stimulus. For example, the first few times a bell sounds or its sound is paired with food, a dog may prick up its ears and/or turn its head toward the bell. Cf. CONDITIONED REFLEX.

Outcome: see PROCESS.

Overall rate: see RATE OF RESPONDING.

Overmatching: see MATCHING LAW.

Overshadowing: an attenuation of respondent conditioning with one stimulus because of the presence of another stimulus. For example, if soft tone and loud bell together precede food, tone may remain ineffective as a CS even though it and bell have the same contingent relation to food. Cf. BLOCKING.

Overt behavior: behavior that is observed or observable, or that affects the organism's environment. Cf. COVERT BEHAVIOR.

P

Paced response; pacing: see DIFFERENTIAL-REINFORCEMENT SCHEDULES.

Paired associates learning: a verbal learning procedure in which each of several stimuli (usually verbal) sets the occasion for a different verbal response. The stimulus items are presented repeatedly in varied order until the learner meets some learning criterion.

Pairing: see CONTIGUITY and cf. CONTINGENCY.

Paradigm: a symbolic representation of relations. For example, a three-term contingency in which a response (R) produces a reinforcer (Rf) in the presence of a discriminative stimulus (S^D) might be written as: S^D:R→Rf. *Paradigm* is often used incorrectly as a synonym for *procedure.*

Parameter: a variable that is held constant while some other variable changes. When different values of a parameter are examined, the parameter distinguishes different functions within a family of functions. For example, a graph of avoidance behavior can show avoidance response rate as a function of either RS interval with SS interval as a parameter or SS interval with RS interval as a parameter.

Parity: the approximate matching of one's own behavior to some properties of the behavior in one's community, especially with regard to verbal behavior, as when a child's babbling produces sounds similar to those made routinely by the child's caregivers (Palmer, 1996). Parity may be a reinforcing consequence, given that caregivers are often sources of reinforcers; deviation from parity may sometimes serve as a punisher.

Passive avoidance: a misnomer for punishment. To avoid passively is to not respond when responding has been punished. For example, a rat is on a platform above an electrified grid. Its failure to step down onto the grid is called passive avoidance, in the sense that in doing so it is not shocked. But defining contingencies in terms of the absence of responses may be misleading. It is more appropriate to say that stepping down is punished by shock.

Pause: a period of no responding, not necessarily bounded by responses. Cf. INTERRESPONSE TIME, LATENCY, RATE OF RESPONDING.

Pavlovian conditioning: see RESPONDENT CONDITIONING.

Peak procedure: omitting some proportion of the reinforcers arranged by an FI schedule and thereby allowing responding to continue for some time after the usual end of the interval. Response rate typically passes through a maximum (the peak) and then decreases over time; the increasing and then decreasing rates are sometimes treated as the two sides of a temporal generalization gradient.

Peak shift: see POSTDISCRIMINATION GRADIENT.

PECS: Picture Exchange Communication System

Percentage reinforcement: the omission of a fixed proportion of scheduled reinforcers. For example, in an FR 100 schedule with 50% reinforcement, only half of the completed ratios end with a reinforcer. A stimulus (e.g., a brief tone) is sometimes substituted for the omitted reinforcer; without such a stimulus, the above schedule is the same as a VR 200 schedule in which the constituent ratios are all multiples of 100 responses.

Percentile-reinforcement schedule: a schedule in which the eligibility of a response to produce a reinforcer depends on its location within a distribution. For example, a schedule for long IRTs might reinforce any IRT in the top 25% of an IRT distribution taken over the last 100 responses. The schedule must specify both the percentile criterion for reinforcement and the reference response distribution. Its criteria for differential reinforcement are relative rather than absolute, so it operates consistently over a range of changes in performance and makes automated shaping possible.

Performance: behavior, usually over extended time periods. A subject matter in itself, performance has often been treated instead as an index of something else (e.g., learning, motivational states).

Phenomenon (plural phenomena): an event; something that happens.

Phenotype: see GENOTYPE.

Phonology: see LINGUISTICS.

Phrase structure grammar: see GRAMMARS.

Phylogenic constraints: limitations on learning or differential capacities for learning that depend on phylogenic selection, including properties of the behavior classes that can be produced and limits on the contingencies that can modify behavior. For example, it may be impossible to shape alternating as opposed to synchronized wing flapping in newly hatched birds. The terminology is rarely invoked when the limitations involve obvious anatomical features. See TASTE AVERSION for an example; cf. PREPAREDNESS.

Phylogenic selection: see NATURAL SELECTION; cf. ONTOGENIC SELECTION.

Phylogeny: the development or evolutionary history of a species. Cf. ONTOGENY.

Place learning versus response learning: the historical issue of whether organisms learned stimuli or movements, as, for example, whether a rat in a maze learns a sequence oriented to stimuli outside the maze or just a particular sequence of turns; the outcome can go one way or the other depending on stimuli outside the maze.

Pliance: instruction-following based on social contingencies rather than the correspondence between verbal behavior and environmental events. Cf. TRACKING, VERBALLY GOVERNED BEHAVIOR.

Polydipsia: the schedule-induced enhancement of water intake. See ADJUNCTIVE BEHAVIOR.

Polymorphous stimulus class: a probabilistic stimulus class in which each member includes exactly *n* of *m* distinguishing features (as when a stimulus is a member of a class by virtue of containing exactly 2 of 3 critical features). In such cases, any feature may also appear in stimuli outside the class.

Polyphenism: see GENOTYPE.

Positive reinforcement: see REINFORCEMENT.

Positive stimulus: see DISCRIMINATIVE STIMULUS.

Positive transfer: see TRANSFER.

Postdiscrimination gradient: a gradient obtained after a discrimination between one stimulus correlated with reinforcement and another correlated with extinction (occasionally, between two stimuli correlated with different reinforcement schedules), usually with both stimuli represented on the continuum along which the gradient is determined. It often includes a *peak shift*, a displacement of the point of maximum responding to one side of the reinforcement stimulus in a direction away from the extinction stimulus.

Postponement: see AVOIDANCE.

Postreinforcement pause: the period of no responding following a reinforcer, especially in an FR or an FI. In an FR, the pause is sometimes measured as the time to some response other than the first (e.g., the fifth response in FR 100), because pauses may separate the first few responses before the roughly constant rate of the FR run begins. It might more appropriately be called a preresponse pause; for example, in multiple FR FR, the duration of the pause is influenced by the stimulus of the current rather than the prior component.

Potentiation: an increase, over repeated presentations, in the respondent behavior elicited by a stimulus (especially, an aversive stimulus). Cf. HABITUATION.

Pragmatics: see LINGUISTICS.

Preaversive stimulus: a stimulus that reliably precedes an aversive stimulus and thus may become a conditioned aversive stimulus. Such stimuli may reduce the responding maintained by positive reinforcers, an effect variously called *anxiety, conditioned emotional response* (*CER*) or *conditioned suppression*. In some contexts, the stimulus increases responding, as when presented during avoidance responding or during positively reinforced responding after a history of avoidance; this has been called *conditioned acceleration* or *conditioned facilitation*.

Prediction: see CONTROL.

Predictive stimulus: a discriminative stimulus. A stimulus predicts an event if the probability of the event given the stimulus differs from that given no stimulus. Cf. INFORMATIVE STIMULUS.

Preference: the probability of one of two or more alternative responses, derived from the relative frequencies of the responses over an extended sequence of choices. The term does not apply when the different probabilities of each response are engendered by the different schedules according to which each is reinforced; For example, higher VR than DRL rates do not imply that VR is preferred. Preferences quantify the relative effectiveness of different consequences as reinforcers (cf. REINFORCEMENT); when each response produces a different consequence, the organism is said to prefer the consequence produced by the response that is most probable. If different probabilities of two or more responses cannot be accounted for, as when they occur despite identical consequences and schedules for each response, the preference may be called a *bias*. Cf. CHOICE, CONCURRENT-CHAIN SCHEDULES, CONCURRENT OPERANTS.

Premack principle: the relativity of reinforcers and punishers. See REINFORCEMENT.

Preparedness: a capacity, presumably of phylogenic origin, to learn some response-stimulus or stimulus-stimulus contingencies more readily than others. For example, organisms may learn relations between tastes and gastrointestinal consequences more easily than those between lights or sounds and such consequences. See TASTE AVERSION.

Primacy: see SERIAL-POSITION EFFECT.

Primary memory: an earlier term for short-term memory. See REMEMBERING.

Primary reinforcer: a reinforcer the effectiveness of which does not depend on its contingent relation to another reinforcer. Cf. CONDITIONAL REINFORCER.

Priming: presenting a stimulus that affects behavior after the stimulus is removed, as when the brief presentation of one word lowers the recognition threshold of a semantically related word presented later.

Private events: in verbal behavior, events accessible only to the speaker (usually, events inside the skin). Private events have the same physical status as public events, but it is more difficult for the verbal community to shape tacts of private events.

Proaction: effects of learning at one time on other learning later. When the later learning is impaired, the effect is a variety of negative TRANSFER called *proactive interference* or *proactive inhibition*. Cf. RETROACTION.

Probabilistic stimulus class: a class in which each member contains some subset of features but none is common to all members. The number of features in the subset may vary from one class member to another (cf. POLYMORPHOUS STIMULUS CLASS). Such classes, sometimes called *fuzzy sets*, do not have well-defined boundaries, though class members may have family resemblances. Examples include *natural concepts* and classes defined by reference to a *prototype*.

Probability: a proportion or relative frequency, either scheduled or derived from data. The probability of an event is given by how often it occurs divided by how often it is possible. For example, if a response occurs on 40 of 50 occasions during which it is sampled, its probability is .8 (40/50). Response probability can be based on response frequencies relative to other responses, in the presence of a stimulus, or within successive time periods. Cf. CONDITIONAL PROBABILITY.

Probe: a condition or stimulus introduced into a performance to clarify the variables controlling it. For example, the interruption of FR responding by an occasional brief stimulus correlated with reinforcement of another response can be used to probe how strongly the FR responses are chained together, or a question can probe variables determining current verbal behavior. For verbal instances of probes, see PROMPT.

Problem solving: constructing discriminative stimuli, either overtly or covertly, in situations involving novel contingencies; these stimuli may set the occasion for effective behavior,as when a verbal problem is converted into a familiar mathematical formula or a listing of options clarifies complex contingencies. Cf. SIMULATION.

Procedural memory: see REMEMBERING.

Procedure: An experimental arrangement or operation. See OPERATION and cf. PARADIGM.

Process: the outcome of an operation; the changes in behavior it produces. Cf. OPERATION.

Processing: whatever goes on within an organism between the presentation of a stimulus and subsequent responding. Cf. COGNITIVE PROCESSES.

Processing, level of: see LEVEL OF PROCESSING.

Productivity: the generation of novel behavior through the recombination and reorganization of existing classes of behavior. Cf. ADDUCTION.

Programming: arranging experimental conditions such as reinforcement schedules. In some usages, *programming* is restricted to arranging systematic changes in conditions, as in shaping or the transfer of stimulus control through fading, and so is distinguished from *scheduling*, which arranges maintained conditions.

Progressive schedule: a schedule in which requirements change progressively with each reinforcer, as when, in a progressive ratio schedule, the ratio increases by 5 responses after each reinforcer. The schedule sometimes allows the requirement to be reset to its initial value (e.g., by a second response and according to some schedule).

Prompt: a verbal stimulus that raises the likelihood of a verbal response by the listener, where the speaker can specify the listener's verbal response. For example, a director may prompt or cue an actor with the beginning of a line that the actor has forgotten. Prompts are distinguished from *probes* in that with probes the speaker cannot say what the listener's verbal response should be. For example, a criminal investigator questioning a suspect is probing, not prompting.

Proprioceptive stimulus: an interoceptive stimulus produced by the effects of movements and postures on receptors in muscles, tendons or joints.

Prospective memory: see REMEMBERING.

Prothetic stimuli: see STIMULUS CONTINUUM.

Prototype: a typical member of a probabilistic class, described by a weighted average of all features of all members of the class. For example, feathers are weighted more heavily than webbed feet among birds because more birds have feathers than have webbed feet; thus, a robin is a more prototypical bird than a duck because it shares more features with other birds than does a duck. Cf. PROBABILISTIC STIMULUS CLASS.

Pseudoconditioning: the elicitation of responding by one stimulus as a result of its presentation in the same context as another, even though neither had been presented in a contingent relation to the other. See SENSITIZATION for an example.

Pseudo-reinforcer: a reinforcer in name only, as when a child is told that it is important to earn gold stars but the effectiveness of the stars as reinforcers is not assessed.

Pseudotrial: a time period corresponding to that of a trial but within which no trial stimuli occur. Pseudotrials are used to assess response probability in the absence of the trial stimuli over time periods that correspond to those of trials.

Psycholinguistics: see LINGUISTICS.

Psychological reality: the role various classes of events play in behavior, especially in psycholinguistics, as when grammatical transformations are demonstrated as something that speakers do with sentences.

Psychophysics: an area of psychology that evolved out of the philosophical concern with the relation between mind and body. Psychophysics relates behavioral properties of stimuli to properties defined in physical terms. Studies of detection or discrimination examine *absolute thresholds* or minimum effective stimulus intensities, and *difference thresholds* or minimum effective differences between stimuli along some continuum (see also SIGNAL DETECTION ANALYSIS). Studies of scaling relate the effects of changes in the properties of one stimulus to those of changes in the properties of another: for example, if responding depends on auditory or visual stimulus intensity, determining how much one must increase to equal the effect of doubling the other.

Punisher: see PUNISHMENT.

Punishment: the response-produced presentation of positive punishers or the termination of negative punishers, or the response decrement or suppression that results. The terminology closely parallels that of reinforcement. *Punishers* are stimuli, *punishment* is an operation or outcome, and responses rather than organisms are said to be punished. A stimulus is a *positive punisher* if its presentation reduces the likelihood of responses that produce it, or a *negative punisher* if its removal reduces the likelihood of responses that terminate it. Like reinforcers, punishers are relative and may be defined independently of their behavioral consequences. For example, the probabilities of two responses can be assessed by forcing the organism to choose between engaging in one or the other, and if the more probable response then forces the organism to engage in the less probable one, the forced responding will punish the more probable response. These definitions parallel the definitions of reinforcers; punishers are equivalent except for the difference in sign. Cf. AVERSIVE STIMULUS, REINFORCEMENT.

R

R, r: usually, response.

Random control: a procedure for presenting two stimuli randomly in time, as a baseline against which to compare effects of stimulus-stimulus contingencies. The random presentations are usually arranged in the context of a sequence of pseudotrials and therefore typically include incidental stimulus-stimulus contiguities as well as presentations of each stimulus alone.

Random-interval schedule: see INTERVAL SCHEDULE.

Randomness: variability generated by a process that produces events that are completely independent of one another, in the sense that none can be predicted from any of the others. It is a property of a distribution of events, or the process that generates the distribution; no single event can be random.

Random-ratio schedule: see RATIO SCHEDULE.

Rate dependency: changes in the effects of a variable that depend on baseline response rate, especially in reference to drug effects, as when a drug dosage increases low response rates but decreases high ones.

Rate of reinforcement: reinforcers per unit time; often used in preference to reinforcement frequency because frequency sometimes refers not to reinforcers per unit time but to reinforcers per response, per session or per trial.

Rate of responding: responses per unit time. Several types have been distinguished: *overall* or *average rate*, determined over a substantial time such as an experimental session; *local, momentary* or *moment-to-moment rate*, determined over a short time, particularly when rate is relatively constant throughout that time; *running rate*, roughly equivalent to local rate, but usually with the provision that it is determined over a time bounded by pauses; and *terminal rate*, determined over a short time just before a reinforcer, especially in an FI. Criteria for distinguishing rates, such as how a running rate is determined, may be informal (e.g., visual inspection) or explicitly defined, as when a period of no responding qualifies as a pause only if it is more than 5 s long.

Other terms distinguish changes in rate: *acceleration* or *positive acceleration* is increasing rate, appearing as concave upward curvature in a cumulative record; *deceleration* or *negative acceleration* is decreasing rate, appearing as concave downward curvature; rate change in general, without specifying direction, is called *celeration*. Cyclic changes are repeated increases and decreases, each completed over a roughly constant time. Compensation is a low rate immediately after an unusually high one, or a high one immediately after an unusually low one. The acceleration typically produced by an FI schedule is often called a *scallop*, especially in reference to what it looks like in a cumulative record. When extended to other schedules, the term is mostly restricted to accelerations bounded by some event, such as a reinforcer. The curvature in an FI has been measured in terms of *quarterlife* (time to complete one quarter of the responses within an interval) and *index of curvature* (a statistic based on the number of responses in successive fractions of an interval). Moment-to-moment changes in rate are often described as *grain*, such as the fine grain of a roughly constant rate or the coarse or steplike grain of a rate that fluctuates, again especially in reference to what they look like in a cumulative record. Cf. CUMULATIVE RECORD, INTERRESPONSE TIME.

Ratio schedule: a schedule in which the last of a specified number of responses is reinforced. In a *fixed-ratio (FR) schedule*, the number is constant from one reinforcer to the next; performance is characterized by pauses after the reinforcer followed by a relatively high and constant response rate. In a *variable-ratio (VR) schedule*, the number of responses varies from reinforcer to reinforcer; relative to FR schedules, the postreinforcement pause is reduced or eliminated. A VR schedule is usually identified in terms of the average responses per reinforcer. In the variety of VR schedule called *random ratio* (RR), the ratio specifies the probability with which responses are reinforced. For example, in RR 20 that probability is .05 (1/20) and is independent of the number of responses emitted since the last reinforcer. In early VR schedules, successive ratios were selected in irregular order from a set of ratios described by a mathematical progression, analogous to those used in VI schedules (see INTERVAL SCHEDULE); these arrangements are now uncommon.

Ratio strain: the appearance of pauses in VR responding, or in FR responding at times other than after a reinforcer (cf. POSTREINFORCEMENT PAUSE); a result of large ratio size and/or low reinforcement frequency.

Reaction time: usually equivalent to latency.

Reality, psychological: see PSYCHOLOGICAL REALITY.

Recall, verbal: see FREE RECALL.

Recency: see SERIAL-POSITION EFFECT.

Recognition, verbal: see VERBAL DISCRIMINATION.

Reconstruction: memory interpreted as a way of reconstructing rather than replicating what is remembered. Remembering interpreted in this way is more like following a recipe than like reading a blueprint.

Recovery: return to an earlier level of responding after it is been reduced by an operation such as extinction or punishment. The vocabulary does not distinguish between recovery during maintained conditions and recovery after the conditions are discontinued. For example, *recovery during punishment* refers to a return of responding toward prepunishment levels while punishment continues, and *recovery after punishment* refers to a return toward those levels after punishment is discontinued.

Reflex: see UNCONDITIONED REFLEX, CONDITIONED REFLEX.

Reflexive relation or **reflexivity:** the identity relation. See EQUIVALENCE CLASS.

Regression: the reappearance of previously extinguished behavior during the extinction of more recently reinforced behavior.

Regular reinforcement: see CONTINUOUS REINFORCEMENT.

Rehearsal: behavior that occurs between storage and retrieval. In most usages, rehearsal includes encoding; occasionally it refers only to behavior that follows encoding. The initial response to a verbal stimulus, which is directly occasioned by that stimulus, may be distinguished from subsequent repetitions by calling it *hearsal*. Some usages also recognize different kinds of rehearsal: *maintenance rehearsal* involves repetitions of encoded items and *elaborative rehearsal* involves further encoding and/or processing. See also REMEMBERING; cf. MEDIATING BEHAVIOR.

Reinforcement: the response-produced presentation of positive reinforcers or termination of negative reinforcers, or the increase or maintenance of responding resulting from this operation. *Reinforcers* are stimuli (e.g., food); *reinforcement* is an operation (e.g., presentation of food given a response) or a process. The operation reinforces responses, not organisms; organisms are sometimes said to be *rewarded*, but this term often implies effects of stimuli other than reinforcing effects. Earlier in its history, reinforcement was also applied to presentations of the US in respondent conditioning, but that usage is now unusual.

A stimulus is a *positive reinforcer* if its presentation increases responding that produces it, or a *negative reinforcer* if its removal increases responding that terminates or postpones it. The distinction matters mainly when responses produced by the reinforcer can compete with the reinforced response. For example, reinforcement by heat of a rat's lever presses in cold is more likely to be called negative reinforcement by removal of cold than positive reinforcement by presentation of heat because cold produces huddling and shivering that may compete with lever pressing.

Reinforcers can also be defined independently of their behavioral consequences. The effectiveness of a reinforcer depends on the relative probabilities of the responses it occasions and the responses to be reinforced; these can be altered by limiting the organism's opportunities to engage in one or the other response (*response deprivation*: cf. MOTIVATING OPERATION). If a less probable response produces a stimulus that occasions a more probable response, then the stimulus will reinforce the less probable response. This definition takes into account the *relativity of reinforcers*; the reinforcement relation is reversible. For example, if water deprivation makes drinking more probable than wheel running, the opportunity to drink will reinforce running, but if limited access to the wheel makes running more probable than drinking, the opportunity to run will reinforce drinking. Cf. OPERANT, PUNISHMENT.

Reinforcer: see REINFORCEMENT.

Reinforcer class: the reinforcer analog of classes of discriminative stimuli or of operants. Reinforcers are the third member of many three-term contingencies. Instances of reinforcers vary like those of other events (e.g., grains in a pigeon feeder may differ in size and shape), and so must be treated as classes. Behavior may be affected by whether the consequences it produces are all members of a single reinforcer class or are members of different classes (Steinman, 1969).

Reinforcer-specific effects: effects that depend on which reinforcers are produced by which responses. If common reinforcers create operant classes, then complex tasks that involve different responses all of which produce the same reinforcer may delay the differentiation of the separate components of the task. Acquisition of matching-to-sample may be accelerated if correct responses for different comparisons each produce a different reinforcer instead of all producing the same reinforcer (Urcuioli, 2005). In an equivalence class experiment with children, for example, all reinforcers of correct matches for one class may consist of one visual display accompanied by music, and those for the other class may consist of a different display accompanied by different music.

Relational autoclitic: see AUTOCLITIC.

Relational discrimination or **relational learning:** discrimination based on relational rather than absolute properties of stimuli: for example, to the left of or to the right of; same or different; greater than or less than (see also MATCHING-TO-SAMPLE).

Relational frame: a description of the relations characterizing a higher-order class, especially in cases of complex stimulus control. For example, if relations AB and AC both satisfy the frame *opposite of*, then the relation BC is one of sameness. Cf. EQUIVALENCE CLASS.

Relative rate: the rate of one event (especially a response or a reinforcer) as a proportion of the summed rates of that and other events. For example, given rates *a* and *b*, the relative rate of *a* is *a* divided by the sum of *a* plus *b*.

Relativity, of language: see LANGUAGE RELATIVITY.

Relativity, of reinforcers: see REINFORCEMENT.

Releaser: an ethological term for a stimulus that elicits a stereotyped pattern of behavior (cf. FIXED ACTION PATTERN). Releasers are often USs provided by the behavior or physical features of another organism, but the behavior produced by releasers typically occurs full blown; unlike the effects of US, it is not likely to vary with stimulus magnitude. In some usages, releasers have some properties of discriminative stimuli that occasion operant behavior. The comparison is complicated because the functions of releasers are usually analyzed differently from those of USs and discriminative stimuli. For example, releasers are usually presented for extended time periods and may vary during those times, particularly when they depend on another organism's behavior, whereas USs are more often presented briefly, in discrete trials.

An artificial releaser that is more likely to produce a fixed action pattern than its natural counterpart is called a *supernormal* stimulus. Cf. DISPLACEMENT ACTIVITY, RESPONDENT BEHAVIOR, VACUUM ACTIVITY.

Remembering: a response occasioned by a stimulus no longer present, perhaps directly or perhaps through the mediation of other behavior with respect to that stimulus. Remembering is often discussed in terms of a metaphor of storage and retrieval, where *storage* occurs when the stimulus is presented and *retrieval* when it is recalled. The time between storage and retrieval is sometimes called the *retention* interval. See HEARSAL, REHEARSAL, RETRIEVAL, STORAGE.

Types of remembering are sometimes distinguished by their time courses. *Iconic memory* and *echoic memory* refer respectively to the brief persistence of effects of visual and auditory stimuli, without an opportunity for a response occasioned by the stimulus (neither hearsal nor rehearsal). *Short-term memory* (*STM*) is remembering based on a single presentation of items and without coding and/or rehearsal beyond the initial response to the stimulus (hearsal but not rehearsal); it is of short duration (e.g., 10 to 20 s) and limited to roughly 5 to 9 items (historically, the *span of immediate memory*). *Long-term memory* (*LTM*) occurs after coding or rehearsal and/or multiple presentations of items (both hearsal and rehearsal), and is therefore of unlimited duration and capacity; subsequent recall after a period of retention involves rehearsal but not hearsal.

Remembering is also classified in terms of what is remembered. Examples include: *procedural memory* (remembering operations or ways of doing things), often contrasted with *declarative memory* (remembering facts); *autobiographical* or *episodic memory* (remembering specific events in one's life); *semantic memory* (remembering aspects of one's language); *spatial memory* (remembering paths and things located on them); and *retrospective memory* (remembering past events), often contrasted with *prospective memory* (remembering things one has to do in the future). See also METAMEMORY.

Reminiscence: an increase in recall as time passes since learning. Reminiscence is an occasional phenomenon and usually appears, if at all, shortly after learning.

Repeated acquisition: a procedure that examines acquisition as steady-state performance. For example, assume a monkey must emit a sequence of presses on four levers to produce a reinforcer and the required sequence changes each session. After many sessions, the monkey has had enough contact with correction procedures and other experimental details that all it has to learn within a session is the new sequence of presses (Boren & Devine, 1968. The repeated acquisition of new sequences may then be used as a baseline for studying how acquisition is affected by different variables (e.g., drugs). The consistent way in which the monkey masters each new sequence in steady-state performance may be called a *strategy*. Cf. HIGHER-ORDER CLASS OF BEHAVIOR, LEARNING SET.

Repertory or **Repertoire:** the behavior an organism can emit, in the sense that the behavior exists at a nonzero level, has been shaped or, if extinguished, may be rapidly reinstated. The organism need not engage in behavior for it to be in its repertory; for example, a rat that has learned a maze has maze running in its repertory even when not in the maze. To the extent that some responses in it are more likely than others, a repertory consists of a hierarchy; operant procedures modify the relative positions of responses in the hierarchy.

Replicative memory: see REPRODUCTION.

Representation: a transformation of stimuli occurring either when an organism responds to them or, as in remembering, later. In some accounts they are copies; in other accounts representations have arbitrary relations to stimuli, as when a visually presented letter is represented by its sound. The latter are more like recipes than like copies and have behavioral dimensions. Cf. CODING, REMEMBERING.

Reproduction: memory interpreted as the production of copies of what is remembered. Current accounts of remembering instead favor RECONSTRUCTION.

Resistance to change: see STRENGTH.

Resistance to extinction: the responses emitted, time elapsed or number of trials until performance meets some extinction criterion, as, for example, the number of responses emitted before 10 min pass with no response. The measure must be specified, because one contingency or schedule may produce more resistance to extinction than another according to one measure but less according to a different one. Cf. STRENGTH.

Respondent: a class of responses defined in terms of stimuli that reliably produce them. For example, salivation elicited by food or acid in the mouth is a member of one respondent class and salivation elicited by a CS is a member of another; spontaneous salivation, in the absence of identifiable stimuli, is not strictly a member of a respondent class, although it is sometimes loosely referred to as such. Cf. OPERANT.

Respondent behavior: behavior elicited by stimuli (cf. RESPONDENT, UNCONDITIONED REFLEX, CONDITIONED REFLEX). Respondent behavior was once considered primarily autonomic (e.g., responses of glands and smooth muscles), but the reflex relation defines respondent behavior regardless of the character of the response. Thus, skeletal responses may have respondent characteristics (see AUTOSHAPING; cf. OPERANT BEHAVIOR).

Respondent conditioning: the modification of respondent behavior by stimulus-stimulus contingencies, also referred to as *classical conditioning* or *Pavlovian conditioning*. Cf. CONDITIONED REFLEX.

Respondent discrimination: differential conditioning, a type of respondent conditioning in which one stimulus is followed by the US but a second is not, as whenfood in the mouth follows bell but not tone. Discrimination has occurred when the CR is elicited by the first stimulus but not the second. The term does not refer to respondent conditioning in general, even though such conditioning entails discrimination between the presence and absence of stimuli.

Response: a unit of behavior, a discrete and usually recurring segment of behavior. Cf. OPERANT, RESPONDENT, SPOND, STRENGTH; see also specific properties: DURATION OF RESPONSE, RATE OF RESPONDING, TOPOGRAPHY OF RESPONSE.

Response competition: the reduction of one response by the time and/or effort involved in concurrent responding. Major interactions among response classes depend mainly on the reinforcers each produces, so such reductions are difficult to demonstrate; with reinforcer effects controlled for, concurrent responses are more likely to show independence than competition. Cf. INHIBITION.

Response cost: any property or consequence of responding that may reduce or punish it. Examples include increases in response effort or force and response-contingent loss or reduction of reinforcers, especially, with humans, point loss superimposed on responding maintained by points; in such cases, however, the effectiveness of points as reinforcers is sometimes assumed rather than confirmed experimentally.

Response deprivation: see MOTIVATING OPERATION, REINFORCEMENT.

Response generalization: an alternative term for *induction*.

Response-independent reinforcer: See FREE REINFORCER, TIME SCHEDULE.

Response induction: see INDUCTION.

Response rate: see RATE OF RESPONDING.

Response strength: see STRENGTH.

Resurgence: see REGRESSION.

Retention: the time between storage and retrieval in the storage-retrieval metaphor of memory.

Retrieval: in the memory metaphor of storage and retrieval, what the learner does at the time something is remembered. Retrieval is typically occasioned by a discriminative stimulus that sets the occasion for it (e.g., a question or an instruction). Cf. DECODING, REMEMBERING.

Retroaction: effects of learning at one time on other learning that occurred earlier. When the earlier learning is impaired, the effect is a variety of negative TRANSFER called *retroactive interference* or *retroactive inhibition*. Cf. PROACTION.

Retrospective memory: see REMEMBERING.

Reversible effects: changes in performance that are eliminated, either immediately or over some time, when the operations that produced them are discontinued. For example, if responding returns to earlier levels after punishment, the effects of punishment are reversible. Effects that are not completely eliminated are sometimes said to be partially reversible.

Reward: see REINFORCEMENT.

Rf: reinforcement or reinforcer. See REINFORCEMENT.

RI: random interval. See INTERVAL SCHEDULE.

Robustness: Resistance to change in the phenotype; a phylogenic analog of operant strength defined in terms of resistance to change.

RR: random ratio. See RATIO SCHEDULE.

RS interval: response-shock interval. See AVOIDANCE.

Rule-governed behavior: verbally governed behavior. Because of varied definitions of *rule* both inside and outside the discipline, this is one of the more problematic expressions in behavior analytic terminology. In many cognitive usages, for example, rules are regarded not as instances of verbal behavior but rather as the internal codification of central processes or concepts, so they have no verbal status. With rules defined as verbal antecedents, any verbal antecedent qualifies as a rule in some usages, as when one is told to do or say something; in others, rules are only those verbal antecedents that specify contingencies, as when one is told what will happen if one does or says something; in still others, the listener's behavior counts only if the stimulus was generated as a verbal stimulus, so that a rescue team's response to a rock formation that accidentally looked like an SOS signal would not count as verbally governed. See VERBALLY GOVERNED BEHAVIOR and cf. CONTINGENCY-SHAPED BEHAVIOR, SPECIFICATION.

Run: a sequence of responses bounded by pauses or some event. For example, an FR run is the response sequence within a single ratio.

Running memory: see METAMEMORY.

Running rate: see RATE OF RESPONDING.

S

s: seconds.

S: stimulus.

S^D, S^+ (“S-dee”): positive stimulus. See DISCRIMINATIVE STIMULUS.

S^Δ, S- (“S-delta”): negative stimulus. See DISCRIMINATIVE STIMULUS.

Sample-specific behavior: in matching-to-sample, differential responding to each sample, usually introduced to ensure stimulus control by the sample. For example, pigeon matching-to-sample may be arranged with a fixed-duration sample after which a peck on the sample turns on the comparisons only if some differential criterion is met (e.g., more than five pecks given green samples or fewer than four given red); the trial ends without comparisons if sample responding does not meet the criterion. Sample-specific responding may guarantee attention to samples but does not necessarily do so for the sample-comparison relation.

Sample stimulus: see MATCHING-TO-SAMPLE.

Satiation: an establishing operation, continued presentation or availability of a reinforcer, that reduces its effectiveness or, as a process, the reduction in effectiveness it produces. Satiation may occur as responses are reinforced or it may be arranged independently of responses. A criterion for satiation with food reinforcers is prefeeding (presenting food for some fixed time or in some fixed amount before a session). Cf. DEPRIVATION.

Scallop: see RATE OF RESPONDING.

Schedule: a specification of the criteria by which responses become eligible to produce reinforcers. The term has been extended to other operations (e.g., schedules of escape, avoidance or punishment). See specific cases: COMPOUND SCHEDULES, DIFFERENTIAL-REINFORCEMENT SCHEDULES, HIGHER-ORDER SCHEDULES, INTERVAL SCHEDULES, LIMITED HOLD, RATIO SCHEDULES, TIME SCHEDULES.

Schedule-induced behavior: see ADJUNCTIVE BEHAVIOR.

Schema: in cognition, an organized representation of events, especially in complex contexts. Examples include spatial schemas relating lengths, areas and volumes and social ones arranged in scripts, scenarios and narratives. Cf. COGNITIVE MAP, REPRESENTATION.

Search: in the storage and retrieval metaphor of memory, search for a match to some target item. The search may be *exhaustive* (all items are checked) or *self-terminating* (the search ends when a match is found).

Secondary reinforcer: see CONDITIONAL REINFORCER.

Second-order: see HIGHER-ORDER CLASS OF BEHAVIOR, HIGHER-ORDER CONDITIONING, HIGHER-ORDER SCHEDULE.

Selection by consequences: operant selection or the ontogenic analogue of phylogenic or Darwinian selection, expressed as an abbreviated form of *the selection of behavior by its consequences*. In a more general sense, all varieties of selection involve consequences. For example, the evolution of the eye depends on the consequences of more finely differentiated seeing. See ARTIFICIAL SELECTION, CULTURAL SELECTION, ONTOGENIC SELECTION, PHYLOGENIC SELECTION, NATURAL SELECTION, SEXUAL SELECTION.

Self-control: a term derived from the colloquial vocabulary that applies to cases in which a relatively immediate small reinforcer is deferred in favor of a later large reinforcer or in favor of avoiding a later large aversive event, or in which a relatively immediate small aversive event is accepted when the acceptance leads to a later large reinforcer or avoids a later large aversive event. Examples include deferring a small purchase to save for a large one, refusing a drink to avoid a hangover, exercising to perform well in a later athletic event, and undergoing preventive dental procedures. The opposite of self-control is called *impulsiveness* or *impulsivity*.

Self-reinforcement: a misnomer for the delivery of a reinforcer to oneself based on one's own behavior. In so-called self-reinforcement, the contingencies and motivating operations that affect the behavior that is purportedly reinforced are confounded with those that affect the delivery of the reinforcer to oneself. The organism that appears to self-reinforce must be able to discriminate behavior that qualifies for the reinforcer from behavior that does not; this behavior is more appropriately described as an example of the discrimination of properties of one's own behavior. When this procedure has effects, it is more appropriately called *self-regulation*.

Semantic memory: see REMEMBERING.

Semantics: see LINGUISTICS.

Sensitivity: the organism's capacity to respond differentially to different stimuli or conditions. In most behavioral usages, sensitivity is measured in terms of *thresholds*. Cf. PSYCHOPHYSICS.

Sensitivity to contingencies: see HIGHER-ORDER CLASS OF BEHAVIOR.

Sensitization: the lowering of a threshold, as when prior delivery of an aversive stimulus lowers the intensity at which a noise elicits a startle response.

Sensory preconditioning: in respondent conditioning, a type of higher-order conditioning in which a contingent relation between two stimuli precedes making one of them a CS. Sensory preconditioning is said to have occurred if the other stimulus elicits the CR solely by virtue of its relation to the first one. Preconditioning procedures have been extended to operant cases. For example, correlating response-independent reinforcers with one stimulus that later signals reinforcement but not with a second one that later signals extinction sometimes facilitates acquisition of an operant discrimination between the stimuli (Morse & Skinner, 1958): see MEDIATION.

Sequential dependencies: conditional probabilities of successive events: for example, given concurrent responses A and B, the probabilities of A followed by A, A followed by B, B followed by A, and B followed by B.

Sequential grammar: see GRAMMARS.

Serial anticipation: see SERIAL LEARNING.

Serial learning: learning an ordered list of items. In *serial recall*, the learner has an opportunity for recall after the entire list is presented. In *serial anticipation*, items are presented one at a time and the learner has an opportunity to say which comes next. In both cases, list presentations continue until the learner meets some learning criterion.

Serial-position effect: differential recall of an item depending on its position in a list, especially in free recall. Early items are more likely to be recalled than later ones (*primacy*), and the most recent items are more likely to be recalled than earlier ones (*recency*). Thus, items at the beginning or end are more likely to be recalled than ones in the middle. Primacy effects are usually stronger than recency effects.

Serial recall: see SERIAL LEARNING.

Set: loosely, a disposition to respond. The term may refer to stereotyped patterns of operant behavior, especially under stimulus control, or to effects of the conditional stimuli of a conditional discrimination. A common usage is provided by instructions at the start of an experiment, often said to produce in humans a *set* to attend to particular features of the situation. Cf. ATTENTION.

Setup: in reinforcement schedules (especially interval schedules), an arrangement that makes a response eligible to produce a reinforcer. Older schedule arrangements allowed only one setup to be available at a time; newer ones sometimes allow setups to accumulate if reinforcers have not been collected.

Sexual selection: phylogenic natural selection based upon female preferences for characteristics of the male, as when the peahen is more likely to mate with the peacock with the most impressive tail. The female preferences have evolved because the quality of the peacock's tail is correlated with its fitness, which is substantial enough to outweigh the disadvantages of its tail.

Shaping: gradually modifying some property of responding, often but not necessarily topography, by differentially reinforcing successive approximations to a target operant class. Shaping is used to produce responses that, because of low operant levels and/or complexity, might not otherwise be emitted or might be emitted only after a considerable time. The variability of the responding that follows reinforcement usually provides opportunities for reinforcing further responses that still more closely approximate the criteria that define the target operant class. Shaping is a variety of operant selection.

Shock: a stimulus sometimes used as an aversive stimulus. Shock is usually delivered through a grid floor on which the organism stands. The parallel bars of the grid are far enough apart that feces or urine cannot short-circuit them. A complication is that the organism's behavior may alter shock level, as when rats contact the shock source through either a bald or a furry body area.

Short-term memory (STM): see REMEMBERING.

SIB: self-injurious behavior.

Side effect: any effect that accompanies the main effect with which an experimenter is concerned. The distinction is often arbitrary, because no stimulus has a single effect. For example, a researcher interested in aggression induced by extinction may regard the decrease in previously reinforced responding during extinction as a side effect, whereas another researcher interested in operant extinction may regard the aggression as the side effect. Cf. MULTIPLE CAUSATION OF BEHAVIOR.

Sidman avoidance: see AVOIDANCE.

Signal: roughly, a discriminative stimulus or an occasion-setting stimulus; a stimulus that sets the occasion on which some contingency operates or on which another stimulus may be presented. Cf. INFORMATIVE and PREDICTIVE STIMULUS.

Signal detection analysis: an analysis of stimulus detectability in terms of conditional probabilities of a response given a signal in noise or noise alone. A response given a signal in noise is a *correct detection* or *hit* and one to noise alone is a *false alarm*; no response given a signal in noise is a *miss* and given noise alone is a *correct rejection*. A measure of *sensitivity* to the signal derived from these measures is called *d'* (d-prime); another measure based on whether false alarms or misses are favored is called *bias*.

Sign-tracking: responding directed toward some feature of a stimulus correlated with reinforcement. Cf. FEATURE-POSITIVE STIMULUS.

Simulation: duplication of some of the properties of events by creating a model; imagining and covert problem solving. Simulation that takes actual contingencies into account may mediate effective behavior in actual environments. Cf. THINKING.

Simultaneous discrimination: a discrimination in which two or more stimuli are presented at the same time rather than successively (e.g., see JUMPING STAND) and which therefore involves two or more alternative responses. The locations of the stimuli are ordinarily at or close to those of the alternative responses (e.g., stimuli on each of two pigeon keys) and the organism is said to respond to one or the other stimulus. Cf. SUCCESSIVE DISCRIMINATION.

Skinner box: a term not in current usage. See CHAMBER.

Span, memory: see MEMORY SPAN.

Spandrel: an incidental byproduct of selection.

Species-specific behavior: behavior observed in all or most members of a species, of only one or of both sexes, and perhaps only over limited times in each organism's life. Different usages may include: emitted behavior before its selection by consequences; unconditioned respondent behavior; and, in fairly consistent environments, stereotyped operant behavior maintained by species-specific primary reinforcers or conditioned reflexes that depend on species-specific unconditioned reflexes. See also specific examples: DISPLACEMENT ACTIVITY, FIXED ACTION PATTERN, RELEASER, VACUUM ACTIVITY.

Species-specific defense reaction (SSDR): avoidance or escape responding that has a phylogenic origin. Such behavior presumably evolved because natural environments do not allow organisms to learn certain types of avoidance or escape responses. For example, a mouse that fails to avoid a predatory cat on its first encounter will probably never have another opportunity to do so.

Specification: the correspondence between a verbal response and what it tacts, when the verbal response occurs outside of the tact relation, as when a mand is said to specify its reinforcer even though the reinforcer may be absent. The term is often used in an informal rather than a technical sense, especially in reference to effects on a listener, as when a response to a word is said to share properties with responses to what the word ordinarily tacts.

Spoken verbal behavior: See VERBAL BEHAVIOR, VOCAL BEHAVIOR

Spond; sponse: coinages suggested as alternatives to *respond* and *response*, on the grounds that they do not carry the implication of responding as being *to* something; the terms thereby capture the distinctive property of emitted behavior (Provine, 1988).

Spontaneous recovery: in operant or respondent extinction, an increment in responding at the start of one session of extinction, relative to the level of responding at the end of the preceding one. Cf. WARMUP.

SS interval: shock-shock interval. See AVOIDANCE.

Stability: session-to-session variability in performance (the lower the variability, the more stable the performance). A performance that can shift from one to another of two or more steady-state baselines maintained by the same conditions is said to be *metastable*. Cf. STEADY STATE.

State-dependent learning: learning that is most likely to be demonstrated when the learner is in the same context as during the original learning. The term is often reserved for learning under specific physiological conditions such as drug states. For example, the learner who learned an item while drunk is more likely to remember it when drunk again than when sober.

Statistics: quantitative methods for summarizing data (descriptive statistics) or evaluating data (statistical inference). Descriptive statistics include measures of *central tendency* or average value (e.g., mean, median, mode); measures of *variability* or dispersion, or the spread of successive measures around an average value (e.g., range, standard deviation, variance); measures of *regression* or the relation between two variables (e.g., the function best describing how two response measures vary together); and measures of *correlation* or how well one variable predicts the value of another (e.g., correlation coefficients, which are positive when two variables vary directly and negative when they vary inversely, and which range from 1.0 when one of the variables is perfectly predicted by the other to zero when one is completely independent of the other). See also DISTRIBUTION, PROBABILITY.

Statistical inference estimates whether an experimental outcome is likely to have been produced by experimental operations or is better regarded as having occurred by chance. It compares an experimental outcome with a theoretical distribution of possible outcomes (e.g., normal, chi-square, or, in analyses of variance, F) based on the assumption that the outcome depended on chance. If the outcome was highly unlikely on this basis (e.g., probability less than .05), the outcome is said to be *statistically significant*. Statistical significance is independent of and unrelated to substantive significance.

Steady state: performance maintained by a set of conditions after systematic session-to-session changes have become negligible. For example, when the rate and pattern of responding within an FI does not vary systematically over sessions, FI performance is said to have reached a steady state. Steady-state performance is a preferred baseline for analyzing effects of variables; if a baseline is unstable, it might be impossible to assess where it would have been if the variable had not been introduced. The decision as to when performance has reached a steady state depends on the criteria for saying that systematic changes have become negligible; such criteria range from informal observation to stringent quantitative assessments.

Stereotyped response: a response with properties (especially topography) that are relatively invariant over successive occurrences.

Stimulus (plural: **stimuli**): any physical event, combination of events or relation among events. A response may be said to occur *in the presence of* a stimulus or *during* a stimulus. The stimulus vocabulary classifies aspects of the environment in much the same way that the response vocabulary classifies aspects of behavior. Like responses, stimuli may be described in terms of physical or behavioral properties and, again like responses, they may be defined in terms of descriptive (nominal) or functional classes (cf. OPERANT). The term may refer to any of the following: specific instances of physical events (e.g., the sound of a bell); combinations of events, sometimes also referred to as *compound stimuli* or *stimulus complexes* (e.g., feeder operation, with accompanying auditory and visual components); the absence of events (e.g., a dark chamber as a stimulus); a relation among events (e.g., matches in a matching-to-sample problem); specific physical properties of events (e.g., green referred to as a stimulus even though it is only one of many properties of a light); classes defined by physical properties (e.g., a stimulus class consisting of all lights within certain limits of intensity and wavelength); and classes defined in terms of behavioral functions (e.g., classes of effective discriminative stimuli or of stimuli effective as reinforcers or punishers).

When *stimulus* is used descriptively, the *continua* or dimensions along which stimuli vary (e.g., intensity, wavelength or frequency, spatial extent, duration) may be discussed in at least two distinct ways: a change in some stimulus property is said to produce a change in the stimulus, or it is said to change one stimulus to another. The usage is typically determined by convenience of exposition rather than convention. For example, *the light was changed from green to blue* is equivalent to *the green light was replaced by blue*. When *stimulus* is used functionally, an event is not a stimulus unless it exerts control over behavior. Functional classes can often be characterized verbally even though their limits cannot be specified adequately in physical terms; for example, red stimuli do not necessarily include wavelengths in the red region of the spectrum). See also ABSTRACTION, CONCEPT, DISCRIMINATION, GENERALIZATION.

Stimulus continuum (plural: **continua**): a stimulus dimension. Stimulus continua that vary along intensive dimensions (e.g., brightness, loudness) are called *prothetic*; those that vary along nonintensive dimensions (e.g., color, pitch) are called *metathetic*. See also STIMULUS.

Stimulus control: the discriminative control of behavior, including control in a respondent discrimination. See CONTROL, DISCRIMINATION, GENERALIZATION, STIMULUS.

Stimulus generalization: see GENERALIZATION.

Stimulus substitution: an account of respondent conditioning, no longer widely accepted, in which the CS becomes a substitute for the US. But a CR is not just a UR now elicited by a new stimulus; one of several problems is that CRs typically differ from URs in many ways. For example, chemical composition may distinguish CS-elicited from US-elicited salivation.

STM: short-term memory. See REMEMBERING.

Storage: in the memory metaphor of storage and retrieval, what the learner does when something to be remembered is presented. Some of the behavior relevant to the stimulus that occurs at or after storage has been called *rehearsal*. Cf. ENCODING, REHEARSAL, REMEMBERING.

Strain: see RATIO STRAIN.

Strategy: a higher-order discriminated operant characterized by relations among different stimuli, responses and/or consequences occurring across trials and/or conditions rather than by specific stimulus or response properties within trials and/or conditions. Different strategies may be appropriate to different settings. For example, if the availability of reinforcers is more likely to alternate between two levers than to remain with one lever, a *win-shift lose-stay* strategy (change levers after each reinforcer) will be more effective than a *win-stay lose-shift* strategy (stick with the lever that produced the last reinforcer).

Strength: as a property of behavior, the resistance of behavior to change, including resistance to extinction, to disruption by added stimuli and/or to effects of reinforcing alternative responses Nevin & Grace, 2000). The term has also been used, in place of specific measures, to describe the general state of a response or reflex, on the assumption that different measures vary together and reflect an underlying disposition to respond. For example, if response latency decreases while magnitude, duration and resistance to extinction increase, response strength is said to have increased. With operants, measures such as rate, latency, force and duration have been used as indices of strength, but each is independently modifiable by differential reinforcement. For brevity without sacrificing generality, processes that might otherwise be described in terms of each of several measures (especially rate, latency and probability of response) are often described just as increments or decrements in responding.

Stroop effect: a demonstration of competition between verbal and nonverbal responses to a visual verbal stimulus. It is difficult to name rapidly the different colors in which color words are printed if the colors and the color words do not correspond (e.g., *blue* printed in red).

Substitution: see STIMULUS SUBSTITUTION.

Successive discrimination: a discrimination in which two or more stimuli are presented one at a time rather than simultaneously and which therefore usually involves only a single response, as in a multiple schedule. In the most accurate usage, the organism responds *in the presence of* each stimulus, but this is often abbreviated to responding *during*, *in* or *to* each. Cf. SIMULTANEOUS DISCRIMINATION.

Summation: the accumulated effect of a repeated stimulus. A stimulus that does not elicit responding if presented only once may do so if presented repeatedly at a high enough rate.

Supernormal stimulus: see RELEASER.

Superstition: the modification or maintenance of behavior by *accidental* (also *adventitious*, *incidental* or *spurious*) relations between responses and reinforcers, as opposed to those either explicitly or implicitly arranged (cf. CONTINGENCY). The term has lost currency with accumulating data demonstrating that behavior is not ordinarily maintained by response-independent reinforcers (see FREE REINFORCER).

Classes of superstitions include: *simple superstitions*, in which responses are maintained, usually unstably, by reinforcers delivered independently of behavior; *concurrent superstitions*, in which one response is maintained by reinforcers produced by a different response; *sensory superstitions*, in which identical contingencies maintain different performances during different stimuli; and *topographical superstitions*, in which reinforcers produce and maintain a response topography that varies over a much narrower range than that specified by the limits of the operant class. Interpretations in terms of superstitious behavior must be drawn with caution, because (i) it is inevitably variable either within or across organisms, and because (ii) performances that superficially appear to be superstitious can sometimes be shown to depend instead on subtle contingencies. Many human superstitions depend on verbally governed behavior rather than, or in addition to, accidental contingencies. For example, to be superstitious about breaking mirrors, one needn't first have seven years of bad luck after doing so.

Suppression: a reduction in responding produced directly or indirectly by an aversive stimulus (e.g., by punishment or by a preaversive stimulus). The term is sometimes extended to any reduction of responding by a stimulus (e.g., an extinction stimulus), but such usages are not always accompanied by a demonstration that the stimulus reduced responding as opposed to having failed to maintain responding. Cf. INHIBITION.

Surface structure: the arrangement of constituents in a particular sentence. Cf. DEEP STRUCTURE.

Symbolic behavior: in some usages, verbal behavior; in a more specialized usage, behavior the function of which has transferred from one stimulus to another by virtue of the membership of both in an *equivalence class*.

Symbolic matching: see MATCHING-TO-SAMPLE.

Symmetrical relation or **symmetry:** see EQUIVALENCE CLASS.

Syntax: see LINGUISTICS.

Synthesis: putting the parts obtained through analysis back together again. Cf. BEHAVIOR ANALYSIS.

T

T, t: usually, time, or an arbitrary number of seconds.

Tact: a verbal discriminative response, as when the verbal response *apple* in the presence of an apple is said to *tact* the apple. The tact captures stimulus control as it enters into verbal behavior. The tact relation includes only responses in the presence of or very shortly after a stimulus, so it is not equivalent to naming or reference.

Tandem (tand) schedule: a compound schedule in which a reinforcer is produced by the successive completion of two or more component schedules, all of which operate during a single stimulus. Cf. CHAINED SCHEDULE.

Taste aversion: rejection of substances with a given taste after their ingestion has been followed later by gastrointestinal distress or nausea, for example as produced by x-irradiation. It might be interpreted as operant behavior (punishment of ingestion of substances with this taste) or as respondent conditioning (where gastrointestinal distress is the US and taste becomes a CS). In either case, its special characteristic is the long delay (sometimes hours) between the taste and its aftermath. The procedure is ineffective over such delays if stimuli such as sounds or lights are substituted for taste. For this reason, taste aversion is often cited as an example of *preparedness*.

Taxis (plural: taxes): phylogenically determined movement or orientation toward or away from a stimulus. For example, negative phototaxis is movement away from light. Cf. KINESIS.

Teaching: arranging contingencies that change the behavior of a student. Educational systems too often emphasis what the teacher does, but the most important part of teaching is what the teacher can bring the student to do. Behavior analytic educational systems include Direct Instruction, Precision Teaching, the Personalized System of Instruction (PSI), and the Verbal Behavior Analysis of the Keller School, among many others.

Temporal conditioning: respondent conditioning in which a US is presented at regular intervals, such as every 10 min. Conditioning is said to have occurred when the CR tends to occur shortly before each US.

Temporal discrimination: discrimination based on temporal properties of stimuli (i.e., duration), often appealed to in accounts of spaced responding. For example, if a response is more likely to be emitted at 10 s than at 5 s since the last one, the two durations may be said to be discriminated. When changes in contingencies alter response rate, the temporal spacing of responses necessarily changes too. Thus, it is preferable to study temporal discrimination directly, as by reinforcing one response after one stimulus duration and a second after another. Duration as a discriminable property of stimuli has some unique features: It is not determined until time has passed, so a discriminative response cannot occur in its presence; and durations cannot be changed discontinuously, unlike other stimulus properties such as intensity.

Temporal integration: control of behavior by the distribution of events in time. Behavior can be affected by events extended over some time, and recent events may weigh more heavily than those further in the past. The way in which the events combine to affect current behavior is called temporal integration. When events are so far removed in time that they no longer contribute, they are said to be beyond the organism's *time horizon*.

Temporally extended unit or sequence: see CHUNKING.

Terminal behavior: stereotyped behavior that reliably occurs late in interstimulus intervals in fixed-time procedures or temporal conditioning, and usually related topographically to the behavior produced by the reinforcer or the CS (e.g., with pigeons, pecking given food presentations). Cf. INTERIM BEHAVIOR.

Terminal link: see CONCURRENT-CHAIN SCHEDULES.

Terminal rate: see RATE OF RESPONDING.

Textual behavior: A formal verbal class in which a written stimulus occasions a corresponding spoken verbal response. The correspondence is defined by the one-to-one relation of verbal units (e.g., letters or words). Textual behavior is not equivalent to reading, because it does not include the additional behavior called understanding or reading for meaning. Cf. DICTATION-TAKING, ECHOIC BEHAVIOR, TRANSCRIPTION.

Thinking: behavior, especially covert and/or verbal behavior. Thinking is not something that produces behavior: "It is not some mysterious process responsible for behavior but the very behavior itself in all the complexity of its controlling relations" (Skinner, 1957, p. 449). Productive thinking occurs when verbal responses are reinforced by specific consequences, as in solving a mathematical problem.

Three-term contingency: see CONTINGENCY.

Threshold: see PSYCHOPHYSICS.

Timeout (TO): a period of nonreinforcement arranged either by extinction during a stimulus or by removal of an opportunity to respond, as, with pigeons, which only rarely peck keys in darkness, by turning off all lights in the chamber). The term is occasionally extended to other cases (e.g., *timeout from avoidance*, during which no shocks are delivered) and so is more precisely specified as *timeout from positive reinforcement*. Timeout as used with children was derived from the procedure, but the practices following from such extensions typically deviate from these technical specifications in various ways.

Time schedule: a schedule of response-independent reinforcer deliveries. Aside from the absence of a required response, time schedules are classified like interval schedules. In *fixed-time* (*FT*) schedules, the time between reinforcers is constant (cf. TEMPORAL CONDITIONING); in *variable-time* (*VT*) schedules, it varies from one delivery to the next. A *random-time* (*RT*) schedule arranges a constant probability of reinforcer delivery at the end of constant recycling time periods. Cf. INTERVAL SCHEDULE.

Timing behavior: see MEDIATING BEHAVIOR, TEMPORAL DISCRIMINATION.

Titration schedule: a schedule in which one response changes a variable in one direction and either a second response or nonoccurrence of the first changes it in the other, as when one response raises stimulus intensity while another lowers it in a psychophysical procedure (Blough, 1958a), or each response produces an increment while each 5-s period of no responding produces a decrement.

TO: see TIMEOUT.

Token reinforcer: a conditional reinforcer (e.g., a coin) that the organism may accumulate and later exchange for other reinforcers.

Topographical drift: gradual changes over time in the topography of responses maintained by a free reinforcers.

Topographical tagging: the identification of different functional properties of responding by correlating each with a different topography, especially a different spatial location. For example, shock avoidance in rats often consists of moderate rates of lever pressing interrupted by occasional high-rate bursts after shock. If an escape lever is added, so that presses on the original lever continue to avoid shock but the rat can terminate shock once delivered only by pressing the escape lever, the high-rate bursts move to that lever. Thus, the moderate rates on the original lever are tagged by their location as depending on the avoidance contingency, whereas the high-rate bursts are tagged as depending on shock deliveries and the escape contingency (Boren, 1961).

Topography of response: spatial configuration or form, such as how an organism operates an operandum or moves from one place to another, sometimes also including a specified location, such as the place on a key struck by the pigeon's beak. Topographies can be complex and are more often described verbally than quantitatively, as in the specification of the limb with which a rat presses a lever.

Trace conditioning: respondent conditioning in which a brief CS presentation is followed by the US after some fixed, extended time period (according to general usage not less than 5 s, but usually considerably longer). Cf. TEMPORAL CONDITIONING.

Tracking: instruction-following based on a history of correspondences between verbal behavior and environmental events. Cf. PLIANCE, VERBALLY GOVERNED BEHAVIOR.

Transcription: A formal verbal class in which a written stimulus occasions a corresponding written response. Some usages have also included dictation-taking as transcription (cf. Skinner, 1957). Correspondence is defined by the one-to-one relation of verbal units (e.g., letters, words). It is not equivalent to visual copying because the units are not defined by form. For example, it may involve print to script, or lowercase to uppercase. Cf. DICTATION-TAKING, ECHOIC BEHAVIOR, TEXTUAL BEHAVIOR.

Transfer: substituting one set of discriminative stimuli for another or, as a process, the stimulus control created or maintained after such a substitution. Transfer may be based on common properties of two sets of stimuli or on similar correlations of the two sets of stimuli with differential contingencies. In verbal learning, transfer from one task to another is usually assessed with reference to a control group that did not learn the first task; it is positive if the first task enhances performance on the second and negative if it does the opposite. Cf. GENERALIZATION, LEARNING SET, PROACTION, RETROACTION, TRANSPOSITION.

Transformational grammar: see GRAMMARS.

Transitive relation or **transitivity:** see EQUIVALENCE CLASS.

Transposition: in transfer experiments, a reversal of stimulus function depending on control by relations among stimuli on a continuum rather than by absolute values. For example, a rat learns to choose the larger of two circles and the smaller one is then replaced by a new one larger than either of the others; transposition with respect to size is shown if the rat chooses the new larger circle rather than the circle, now smaller, it had previously chosen.

Trial: a discrete period, usually stimulus-correlated, during which an organism has an opportunity to respond. Trials are separated by intertrial intervals that may consist of any of the following: a stimulus condition (e.g., a dark chamber); removal of the operandum or operanda; or removal of the organism from the chamber, especially when the organism, after emitting a response such as running an alley, is no longer in a position to respond again. Trials distinguish discrete-operant procedures from free-operant procedures. Cf. OPERANT, PSEUDOTRIAL.

Tropism: directed movement. This term is most commonly applied to plants, as when a flowering plant turns toward the light in a *phototropism*.

Two-factor theory: in general, any behavioral theory involving the interaction of operant and respondent processes; more specifically, an avoidance theory stating that avoidance responses are operants reinforced by termination of conditioned aversive stimuli established through a respondent process.

U

Unconditioned reflex or **unconditional reflex:** a relation between a stimulus and a response that does not depend on prior conditioning. A reflex is the reliable production of a response by a stimulus. The stimulus is an *unconditioned stimulus* (*US*) and the response is an *unconditioned response* (*UR*). The stimulus is said to elicit the response. Examples of unconditioned reflexes are the salivary reflex (salivation elicited by food or acid in the mouth) and the patellar reflex (a knee jerk elicited by a blow on the patellar tendon). In each case, the elicitation of the response by the stimulus, not the response alone or the stimulus alone, defines the reflex. Cf. RESPONDENT.

Unconditioned or **unconditional response (UR)** or **stimulus (US):** see UNCONDITIONED REFLEX.

Undermatching: see MATCHING LAW.

Universals of language: see LANGUAGE UNIVERSALS.

UR: unconditioned response. See UNCONDITIONED REFLEX.

US: unconditioned stimulus. See UNCONDITIONED REFLEX.

V

Vacuum activity: an ethological term referring to responding (see FIXED ACTION PATTERN) in the absence of the stimulus (see RELEASER) that ordinarily produces it. Cf. DISPLACEMENT ACTIVITY.

Variability: the raw material on which selection operates. Variability is also a property for which contingencies can be arranged, but it cannot be attributed to single responses because it can only be a property of populations of responses. See also STATISTICS

Variable-interval schedule: see INTERVAL SCHEDULE.

Variable-ratio schedule: see RATIO SCHEDULE.

Variable-time schedule: see TIME SCHEDULE.

Variation: see VARIABILITY; cf. SELECTION BY CONSEQUENCES.

Verbal behavior: any behavior involving words, without regard to modality (e.g., spoken, written, gestural). Verbal behavior involves both listener behavior shaped by its effects on speaker behavior and speaker behavior shaped by its effects on listener behavior. The field of verbal behavior is concerned with the behavior of individuals, and the functional units of their verbal behavior are determined by the practices of a verbal community. Cf. LANGUAGE, VOCAL BEHAVIOR.

Verbal discrimination: any discrimination among verbal stimuli, as in discriminating among nouns and verbs in a sentence. Discrimination among items on the basis of whether they appeared in a given context, *verbal recognition*, is a special case of verbal discrimination.

Verbal learning: see FREE RECALL, PAIRED-ASSOCIATES LEARNING, SERIAL LEARNING, VERBAL DISCRIMINATION, VERBAL RECOGNITION.

Verbal recognition: see VERBAL DISCRIMINATION.

Verbally governed behavior: behavior, either verbal or nonverbal, under the control of verbal antecedents. Verbally governed behavior has also been called RULE-GOVERNED BEHAVIOR and *instruction-following*. Contingencies operate for the following of instructions, so instruction-following is likely to become a higher-order class. Verbal antecedents may alter the functions of other stimuli, as when something neutral becomes a reinforcer after one is told that it is worth having. They may also produce instruction-following; they do not qualify as discriminative stimuli if they do so even when they are no longer present. The verbal behavior of one individual may provide verbal antecedents for another, but verbal antecedents may also be shaped or self-generated. Once verbal contingencies have created correspondences between saying and doing so that saying is often accompanied by doing, other behavior may be affected by such shaped or self-generated verbal behavior. Cf. CONTINGENCY-SHAPED BEHAVIOR, HIGHER-ORDER CLASS OF BEHAVIOR.

VI: variable interval. See INTERVAL SCHEDULE.

Vicarious learning: see OBSERVATIONAL LEARNING.

Vocal behavior: Behavior of lips, tongue, etc., that modulates air flow and produces sound. Vocal behavior is not necessarily verbal, though it typically is in humans. Cf. VERBAL BEHAVIOR.

Volition: voluntary behavior; the verbal report of contingency-shaped behavior that is jointly determined by the interaction of responses and their consequences and by levels of variability that change as a function of the reinforcement contingencies (Neuringer & Jensen, 2010). When these properties fall within certain ranges, either one's own behavior or the behavior of others is likely to be called voluntary; in other words, reports of volition depend on discriminations of the way in which behavior interacts with the environment,

Von Restorff effect: the enhanced likelihood of recall of a distinctive item in a list.

VR: variable ratio. See RATIO SCHEDULE.

VT: variable time. See TIME SCHEDULE.

W

Warmup: a low or zero response rate at the start of a session followed by an increase to the rate maintained later, especially in avoidance.

Warning stimulus: a stimulus that precedes an avoidable aversive stimulus. See AVOIDANCE.

Wheel running: sometimes taken as an index of level of activity, especially in rats. The rat runs inside the wheel, which usually turns in only one direction to simplify recording of revolutions or distance run. Wheel running has a high baseline level and is relatively continuous compared to such discrete responses as lever presses.

Win-shift lose-stay or **win-stay lose-shift:** see STRATEGY.

Working memory: see METAMEMORY.

Y

Yoking: connecting chambers so that the performance of an organism in one determines the stimuli and/or schedules for an organism in the other, as in equating VR and VI reinforcement rates by letting the times between reinforcers in one organism's VR performance determine the intervals of another's VI schedule. In within-organism yoking, an experimental condition is yoked to some property of the organism's own performance in an earlier condition.

Conclusions from yoking must be cautiously drawn. For example, assume two types of rats equally distributed among groups in a yoking experiment on the role of avoidance in shock-induced ulcers. Sensitive types are prone to ulcers when shocked; they respond rapidly at low shock levels and receive few shocks but sporadically at higher levels and receive many shocks. Insensitive types are resistant to ulcers when shocked; they also respond slowly at low shock levels and receive many shocks but rapidly at higher levels and receive few shocks. For each shock received by an avoidance rat, an unavoidable shock is delivered to its yoked partner. At low shock levels, only yoked rats develop ulcers: only insensitive avoidance rats respond slowly and receive frequent shocks, and they do not develop ulcers; but all their yoked partners also receive frequent shocks and half of those are sensitive. At higher levels, more avoidance rats develop ulcers than do their yoked partners: sensitive avoidance rats respond sporadically, thereby receiving frequent shocks and developing ulcers; all their yoked partners also receive frequent shocks but only half of those are sensitive and develop ulcers. Thus, a yoking experiment done at one shock level would yield a different conclusion about avoidance and shock-induced ulcers than one done at another level.

Z

Zeigarnik effect: the greater likelihood of remembering an unfinished than a finished task, as when you can't figure out how to end a

;-)

References

Ader, R., & Cohen, N. (1985). CNS-immune system interactions: Conditioning phenomena. *Behavioral and Brain Sciences, 8*, 379–394. [ch 17]

American Heritage. (2011). *The American Heritage Dictionary of the English Language* (5th ed.). Boston: Houghton Mifflin Harcourt. [etym, gloss]

American Psychiatric Association. (2000). *Diagnostic and Statistical Manual of Mental Disorders (DSM-IV-TR)* (4th, Text Revision ed.). Washington, DC: American Psychiatric Association. [ch 9]

Anderson, J. R. (1984). Monkeys with mirrors: Some questions for primate psychology. *International Journal of Primatology, 5*, 81–98. [ch 19]

Anderson, J. R. (1993). Problem solving and learning. *American Psychologist, 48*, 35–44. [ch 29]

Anderson, J. R., & Bower, G. H. (1972). Recognition and retrieval processes in free recall. *Psychological Review, 79*, 97–123. [ch 21]

Anderson, J. R., & Ross, B. H. (1980). Evidence against a semantic-episodic distinction. *Journal of Experimental Psychology: Human Learning and Memory, 6*, 441–466. [ch 27]

Andresen, J. T. (1990). Skinner and Chomsky thirty years later. *Historiographia Linguistica, 17*, 145–166. [ch 25]

Andronis, P. T., Layng, T. V. J., & Goldiamond, I. (1997). Contingency adduction of "symbolic aggression" by pigeons. *Analysis of Verbal Behavior, 14*, 5–17. [ch 13]

Anger, D. (1956). The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology, 52*, 145–161. [ch 15]

Anger, D. (1963). The role of temporal discriminations in the reinforcement of Sidman avoidance behavior. *Journal of Experimental Psychology, 6*, 477–506. [ch 8]

Antonitis, J. J. (1951). Response variability in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology, 42*, 273–281. [ch 9]

Appiah, K. A. (2010). *The honor code: How moral revolutions happen.* New York: Norton. [ch 11, 30]

Arbib, M. A. (2005). From monkey-like action recognition to human language: An evolutionary framework for neurolinguistics. *Behavioral and Brain Sciences, 28*, 105–167. [ch 19]

Asen, Y., & Cook, R. G. (2012). Discrimination and categorization of actions by pigeons. *Psychological Science, online early.* [ch 19]

Ashcraft, M. H. (1982). The development of mental arithmetic: A chronometric approach. *Developmental Review, 2*, 213–236. [ch 29]

Austin, J. L. (1962). *How to do things with words.* Cambridge, MA: Harvard University Press. [ch 22, 25]

Ayer, A. J. (1946). *Language, truth and logic* (2nd ed.). New York: Dover. [ch 30]

Ayllon, T., & Azrin, N. H. (1968). *The token economy.* New York: Appleton-Century-Crofts. [ch 16]

Ayllon, T., & Michael, J. (1959). The psychiatric nurse as a behavioral engineer. *Journal of the Experimental Analysis of Behavior, 2*, 323–334. [ch 9]

Azrin, N. H. (1956). Some effects of two intermittent schedules of immediate and nonimmediate punishment. Journal of Psychology. *Journal of Psychology, 42*, 3–21. [ch 7, 15]

Azrin, N. H., & Hake, D. F. (1969). Positive conditioned suppression: Conditioned suppression using positive reinforcers as the unconditioned stimuli. *Journal of the Experimental Analysis of Behavior, 12*, 167–173. [ch 18]

Azrin, N. H., & Holz, W. C. (1966). Punishment. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 380–447). New York: Appleton-Century-Crofts. [ch 7]

Azrin, N. H., Holz, W. C., Ulrich, R., & Goldiamond, I. (1961). The control of the content of conversation by reinforcement. *Journal of the Experimental Analysis of Behavior, 4*, 4–25. [ch 23]

Azrin, N. H., Hutchinson, R. R., & Hake, D. F. (1966). Extinction-induced aggression. *Journal of the Experimental Analysis of Behavior, 9*, 191–204. [ch 5]

Azrin, N. H., Hutchinson, R. R., & Hake, D. F. (1967). Attack, avoidance, and escape reactions to aversive shock. *Journal of the Experimental Analysis of Behavior, 10*, 131–148. [ch 5, 8]

Azrin, N. H., Hutchinson, R. R., & McLaughlin, R. (1965). The opportunity for aggression as an operant reinforcer. *Journal of the Experimental Analysis of Behavior, 8*, 171–180. [ch 5]

Baddeley, A. D. (1976). *The psychology of memory*. New York: Basic Books. [ch 27, 28]

Badia, P., Suter, S., & Lewis, P. (1966). Rat vocalization to shock with and without a CS. *Psychonomic Science, 4*, 117–118. [ch 4]

Baer, D. M., Peterson, R. F., & Sherman, J. A. (1967). The development of imitation by reinforcing behavioral similarity to a model. *Journal of the Experimental Analysis of Behavior, 10*, 405–416. [ch 19]

Baer, R. A., Detrich, R., & Weninger, J. M. (1988). On the functional role of the verbalization in correspondence training procedures. *Journal of Applied Behavior Analysis, 21*, 345–356. [ch 23]

Balda, R. P., Kamil, A. C., & Grim, K. (1986). Revisits to empty cache sites by nutcrackers. *Animal Behaviour, 34*, 1289–1298. [ch 12]

Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist, 30*, 441–451. [ch 3]

Ballard, P. B. (1913). Obliviscence and reminiscence. *British Journal of Psychology Monograph Supplements, 1(No. 2)*. [ch 27]

Bandura, A. (1976). Self-reinforcement: Theoretical and methodological considerations. *Behaviorism, 4*, 135–155. [ch 19]

Bandura, A. (1986). *Social foundations of thought and action*. Englewood Cliffs, NJ: Prentice-Hall. [ch 19]

Bandura, A. (1995). Comments on the crusade against the causal efficacy of human thought. *Journal of Behavior Therapy and Experimental Psychiatry, 26*, 179–190. [ch 19]

Bandura, A., Adams, N. E., & Beyer, J. (1977). Cognitive processes mediating behavioral change. *Journal of Personality and Social Psychology, 35*. [ch 23]

Baron, A., & Leinenweber, A. (1995). Effects of a variable-ratio conditioning history on sensitivity to fixed-interval contingencies in rats. *Journal of the Experimental Analysis of Behavior, 63*, 97–110. [ch 23]

Barrett, B. H. (2002). *The Technology of Teaching revisited*. Concord, MA: Cambridge Center for Behavioral Studies. [Ch 13]

Barrett, J. E., & Stanley, J. A. (1980). Maintenance of responding by squirrel monkeys under a concurrent shock-postponement, fixed-interval shock-presentation schedule. *Journal of the Experimental Analysis of Behavior, 34*, 117–129. [ch 7]

Bartlett, F. C. (1932). *Remembering*. Cambridge: Cambridge University Press. [ch 25, 26]

Bauer, D. W., & Miller, J. L. (1982). Stimulus-response compatibility and the motor system. *Quarterly Journal of Experimental Psychology, 34A*, 367–380. [ch 18]

Bauer, P. J. (1996). What do infants recall of their lives? Memory for specific events by one- to two-year-olds. *American Psychologist, 51*, 29–41. [ch 26]

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior, 22*, 231–242. [ch 16]

Baum, W. M. (2002). From molecular to molar: a paradigm shift in behavior analysis. *Journal of the Experimental Analysis of Behavior, 78*, 95–116. [ch 16]

Baum, W. M. (2011). Behaviorism, private events, and the molar view of behavior. *Behavior Analyst, 34*, 185–200. [ch 22]

Bechterev, V. M. (1933). *General principles of human reflexology (tr. E. Murphy & W. Murphy)*. London: Jarrolds. [ch 17]

Bellezza, F. S. (1982). Updating memory using mnemonic devices. *Cognitive Psychology, 14*, 301–327. [ch 26]

Bem, D. J. (1967). Self perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review, 74*, 183–200. [ch 19, 22]

Berg, J. P., & McDowell, J. J. (2011). Quantitative, steady-state properties of Catania's computational model of the operant reserve. *Behavioural Processes, 75*, 71–83. [ch 15]

Berko, J. (1958). The child's learning of English morphology. *Word, 14*, 150–177. [ch 13]

Bernstein, D. J., & Ebbesen, E. B. (1978). Reinforcement and substitution in humans: A multiple-response analysis. *Journal of the Experimental Analysis of Behavior, 30*, 243–253. [ch 6]

Bersh, P. J., & Lambert, J. V. (1975). The discriminative control of free-operant avoidance despite exposure to shock during the stimulus correlated with nonreinforcement. *Journal of the Experimental Analysis of Behavior, 23*, 111–120. [ch 8, 14, 16]

Bickel, W. K., Green, L., & Vuchinich, R. E. (1995). Behavioral economics. *Journal of the Experimental Analysis of Behavior. 64*, 257–262. [ch 16]

Binder, C. (1996). Behavioral fluency: Evolution of a new paradigm. *The Behavior Analyst, 1996*, 163–197. [ch 6, 13, 30]

Bisanz, G. L., Vesonder, G. T., & Voss, J. F. (1978). Knowledge of one's own responding and the relation of such knowledge to learning. *Journal of Experimental Psychology, 25*, 116–128. [ch 26]

Bix, H. P. (2000). *Hirohito and the making of modern Japan*. New York: Harper Collins. [ch 24]

Bjork, R. A. (1978). The updating of human memory. In G. H. Bower (Ed.), *The psychology of learning and motivation. Volume 12* (pp. 235–259). New York: Academic Press. [ch 26]

Blackman, D. E. (1968). Conditioned suppression or facilitation as a function of the behavioral baseline. *Journal of the Experimental Analysis of Behavior, 11*, 53–61. [ch 18]

Blackman, D. E. (1977). Conditioned suppression and the effects of classical conditioning on operant behavior. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 340–363). Englewood Cliffs, NJ: Prentice-Hall. [ch 18]

Blaisdell, A. P., & Cook, R. G. (2005). Two-item same-different concept learning in pigeons. *Learning and Behavior, 33*, 67–77. [ch 14]

Blakemore, C., & Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature, 228*, 477–478. [ch 6]

Bloom, H. S., Criswell, E. L., Pennypacker, H. S., Catania, A. C., & Adams, C. K. (1982). Major stimulus dimensions determining detection of simulated breast lesions. *Perception and Psychophysics, 32*, 251–260. [ch 11]

Blough, D. S. (1958a). A method for obtaining psychophysical thresholds from the pigeon. *Journal of the Experimental Analysis of Behavior, 1*, 31–43. [glossT, ch 11]

Blough, D. S. (1958b). New test for tranquillizers. *Science, 127*, 586–587. [ch 9]

Blough, D. S. (1959). Delayed matching in the pigeon. *Journal of the Experimental Analysis of Behavior, 2*, 151–160. [ch 26]

Blough, D. S. (1989). Odd-item search in pigeons: Display size and transfer effects. *Journal of Experimental Psychology: Animal Behavior Processes, 15*, 14–22. [ch 12]

Blough, D. S. (1992). Effects of stimulus frequency and reinforcement variables on reaction time. *Journal of the Experimental Analysis of Behavior, 57*, 47–50. [ch 22]

Boakes, R. A. (1973). Response decrements produced by extinction and by response-independent reinforcement. *Journal of the Experimental Analysis of Behavior, 19*, 293–302. [ch 5]

Bolinger, D. (1973). Truth is a linguistic question. *Language, 49*, 539–550. [ch 22]

Bolker, J. (2000). Modularity in development and why it matters to evo-devo. *Amewrican Zoologist, 40*, 770–776. [ch 3, 30]

Bolles, R. C. (1970). Species-specific defense reactions and avoidance learning. *Psychological Review, 77*, 32–48. [ch 8]

Bolles, R. C. (1975). *Theory of motivation* (2nd ed.). New York: Harper & Row. [ch 6]

Bones, R., Keenan, M., Askin, D., Adams, L., Taylor, D., & Nicholas, O. (2001). The effects of response topography on functional equivalence class formation. *Psychological Record, 51*, 89–110. [ch 13]

Boren, J. J. (1961). Isolation of post-shock responding in a free operant avoidance procedure. *Psychological Reports, 9*, 265–266. [gloss]

Boren, J. J., & Devine, D. D. (1968). The repeated acquisition of behavioral chains. *Journal of the Experimental Analysis of Behavior, 11*, 651–660. [glossR, ch 10, 12]

Boren, M. C. P., & Gollub, L. R. (1972). Accuracy of performance on a matching-to-sample procedure under interval schedules. *Journal of the Experimental Analysis of Behavior, 18*, 65–77. [ch 16]

Bottoms, B. L., Najdowski, C. J., & Goodman, G. S. (Eds.). (2009). *Children as victims, witnesses, and offenders*. New York: Guilford. [ch 26]

Bousfield, W. A. (1953). The occurrence of clustering in the recall of randomly arranged associates. *Journal of General Psychology, 49*, 229–240. [ch 21]

Bower, G. H. (1970). Analysis of a mnemonic device. *American Scientist, 58*, 496–510. [ch 26]

Bower, G. H. (1981). Mood and memory. *American Psychologist, 36*, 129–148. [ch 26]

Bowler, P. J. (1983). *The eclipse of Darwinism*. Baltimore: Johns Hopkins University Press. [ch 3]

Brady, J. V. (2007). Behavior analysis in the space age. *The Behavior Analyst Today, 8*, 398–412. [ch 9]

Braine, M. D. S., & Rumain, B. (1981). Development of comprehension of "or": Evidence for a sequence of competencies. *Journal of Experimental Child Psychology, 31*, 46–70. [ch 25]

Brainerd, C. J., Reyna, V. F., & Brandse, E. (1995). Are children's false memories more persistent than their true memories? *Psychological Science, 6*, 359–364. [ch 26]

Brandauer, C. (1958). *The effects of uniform probabilities of reinforcement on the response rate of the pigeon*. Unpublished doctoral dissertation, Columbia University. [ch 15]

Bransford, J. D., & Franks, J. J. (1971). The abstraction of linguistic ideas. *Cognitive Psychology, 2*, 331–350. [ch 25]

Breland, K., & Breland, M. (1951). A field of applied animal psychology. *American Psychologist, 6*, 202–204. [ch 9]

Breland, K., & Breland, M. (1961). The misbehavior of organisms. *American Psychologist, 16*, 681–684. [ch 18]

Bridgman, P. W. (1927). *The logic of modern physics*. New York: Macmillan. [ch 30]

Bright, M. (1985). *Animal language*. Ithaca, NY: Cornell. [ch 24]

Brogden, W. J. (1939). Sensory preconditioning. *Journal of Experimental Psychology, 25*, 323–332. [ch 17]

Brown, P. L., & Jenkins, H. M. (1968). Auto-shaping of the pigeon's key-peck. *Journal of the Experimental Analysis of Behavior, 11*, 1–8. [ch 17]

Brown, R. (1973). *A first language*. Cambridge, MA: Harvard University Press. [ch 25]

Brown, R., & McNeill, D. (1966). The "tip of the tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior, 5*, 325–337. [ch 22, 26]

Bryan, W. L., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review, 6*, 345–375. [ch 30]

Burns, G. L., & Staats, A. W. (1991). Rule-governed behavior: Unifying radical and paradigmatic behaviorism. *Analysis of Verbal Behavior, 9*, 127–143. [ch 30]

Butler, R. A. (1957). The effect of deprivation of visual incentives on visual exploration motivation in monkeys. *Journal of Comparative and Physiological Psychology, 50*, 177–179. [ch 6]

Bykov, K. M. (1957). *The cerebral cortex and the internal organs (tr. W. H. Gantt).* New York: Chemical Publishing. [ch 17]

Calkins, M. W. (1894). Association. *Psychological Review, 1*, 476–483. [ch 21]

Calkins, M. W. (1896). Association. II. *Psychological Review, 3*, 32–49. [ch 21]

Cameron, J., Banko, K. M., & Pierce, W. D. (2001). Pervasive negative effects of rewards on intrinsic motivation: The myth continues. *The Behavior Analyst, 24*, 1–44. [ch 14]

Cameron, J., & Pierce, W. D. (1994). Reinforcement, reward, and intrinsic motivation: A meta-analysis. *Review of Educational Research, 64*, 363–423. [ch 14]

Camp, D. S., Raymond, G. A., & Church, R. M. (1967). Temporal relationship between response and punishment. *Journal of Experimental Psychology, 74*, 114–123. [ch 7]

Camões-Costa, V., Erjavec, M., & Horne, P. J. (2011). The impact of body-part-naming training on the accuracy of imitative performances in 2– to 3–year-old children. *Journal of the Experimental Analysis of Behavior, 96*, 291–315. [ch 19, 22]

Capaldi, E. D., & Davidson, T. L. (1979). Control of instrumental behavior by deprivation stimuli. *Journal of Experimental Psychology: Animal Behavior Processes, 5*, 355–367. [ch 19]

Caramazza, A., & Brones, I. (1980). Semantic classification by bilinguals. *Canadian Journal of Psychology, 34*, 77–81. [ch 20]

Carew, T. J. (1992). Aplysia: Development of processes underlying learning. In L. R. Squire (Ed.), *The encyclopedia of learning and memory* (pp. 51–56). New York: Macmillan. [ch 30]

Carmichael, L. C., Hogan, H. P., & Walter, A. A. (1932). An experimental study of the effect of language on the reproduction of visually perceived form. *Journal of Experimental Psychology, 15*, 73–86. [ch 26]

Carr, E. G., Newsom, C. D., & Binkoff, J. A. (1980). Escape as a factor in the aggressive behavior of two retarded children. *Journal of Applied Behavior Analysis, 13*, 101–117. [ch 8]

Carr, J. E., Severtson, J. M., & Lepper, T. L. (2009). Noncontingent reinforcement is an empirically supported treatment for problem behavior exhibited by individuals with developmental disabilities. *Research in Developmental Disabilities, 30*, 44–57. [ch 5]

Carrigan, P. F., Jr., & Sidman, M. (1992). Conditional discrimination and equivalence relations: A theoretical analysis of control by negative stimuli. *Journal of the Experimental Analysis of Behavior, 58*, 183–204. [ch 12]

Carroll, S. B. (2005). *Endless forms most beautiful: The new science of evo devo.* New York: Norton. [ch 3, 13, 16, 19, 29, 30]

Carroll, S. B. (2006). *The making of the fittest.* New York: Norton. [glossG, ch 3, 4, 7, 9, 11, 30]

Carter, D. E., & Werner, T. J. (1978). Complex learning and information processing by pigeons: A critical analysis. *Journal of the Experimental Analysis of Behavior, 29*, 565–601. [ch 12]

Catania, A. C. (1963a). Concurrent performances: A baseline for the study of reinforcement magnitude. *Journal of the Experimental Analysis of Behavior, 6*, 299–300. [ch 16]

Catania, A. C. (1963b). Concurrent performances: Reinforcement interaction and response independence. *Journal of the Experimental Analysis of Behavior, 6*, 252–263. [ch 16]

Catania, A. C. (1966). Concurrent operants. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 213–270). New York: Appleton-Century-Crofts. [ch 16]

Catania, A. C. (1968). Glossary. In A. C. Catania (Ed.), *Contemporary research in operant behavior* (pp. 327–349). Glenview, IL: Scott, Foresman. [gloss]

Catania, A. C. (1969). Concurrent performances: Inhibition of one response by reinforcement of another. *Journal of the Experimental Analysis of Behavior, 12*, 731–744. [ch 14, 16]

Catania, A. C. (1970). Reinforcement schedules and psychophysical judgments: A study of some temporal properties of behavior. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 1–42). New York: Appleton-Century-Crofts. [ch 10, 11, 15]

Catania, A. C. (1971). Reinforcement schedules: The role of responses preceding the one that produces the reinforcer. *Journal of the Experimental Analysis of Behavior, 15*, 271–287. [ch 5, 15]

Catania, A. C. (1972). Chomsky's formal analysis of natural languages: A behavioral translation. *Behaviorism, 1*, 1–15. [ch 20, 25]

Catania, A. C. (1973a). The concept of the operant in the analysis of behavior. *Behaviorism, 1*, 103–116. [ch 9]

Catania, A. C. (1973b). The psychologies of structure, function, and development. *American Psychologist, 28*, 434–443. [ch 20, 30]

Catania, A. C. (1975a). Freedom and knowledge: An experimental analysis of preference in pigeons. *Journal of the Experimental Analysis of Behavior, 24*, 89–106. [ch 14]

Catania, A. C. (1975b). The myth of self-reinforcement. *Behaviorism, 3*, 192–199. [ch 19]

Catania, A. C. (1978). The psychology of learning: some lessons from the Darwinian revolution. *Annals of the New York Academy of Sciences, 309*, 18–28. [ch 3, 30]

Catania, A. C. (1980a). Autoclitic processes and the structure of behavior. *Behaviorism, 8*, 175–186. [ch 13]

Catania, A. C. (1980b). Freedom of choice: A behavioral analysis. In G. H. Bower (Ed.), *The psychology of learning and motivation. Volume 14* (pp. 97–145). New York: Academic Press. [ch 14]

Catania, A. C. (1981). The flight from experimental analysis. In C. M. Bradshaw, E. Szabadi & C. F. Lowe (Eds.), *Quantification of steady-state operant behaviour* (pp. 49–64). Amsterdam: Elsevier/North-Holland. [ch 16]

Catania, A. C. (1985). The two psychologies of learning: Blind alleys and nonsense syllables. In S. Koch & D. E. Leary (Eds.), *A century of psychology as science* (pp. 322–335). New York: McGraw-Hill. [ch 30]

Catania, A. C. (1987). Some Darwinian lessons for behavior analysis. A review of Peter J. Bowler's "The eclipse of Darwinism." *Journal of the Experimental Analysis of Behavior, 47*, 249–257. [ch 3, 10, 30]

Catania, A. C. (1989). Speaking of behavior. *Journal of the Experimental Analysis of Behavior, 52*, 193–196. [gloss]

Catania, A. C. (1991). Glossary. In I. H. Iversen & K. A. Lattal (Eds.), *Experimental analysis of behavior. Part 2* (pp. G1–G44). Amsterdam, The Netherlands: Elsevier. [gloss]

Catania, A. C. (1993a). The unconventional philosophy of science of behavior analysis. *Journal of the Experimental Analysis of Behavior, 60*, 449–452. [ch 29, 30]

Catania, A. C. (1993b). What John B. Watson left out of his behaviorism. *Mexican Journal of Behavior Analysis, 19*, 133–146. [ch 30]

Catania, A. C. (1994). The natural and artificial selection of verbal behavior. In S. C. Hayes, L. Hayes, M. Sato & K. Ono (Eds.), *Behavior analysis of language and cognition* (pp. 31–49). Reno, NV: Context. [ch 3]

Catania, A. C. (1995a). Higher-order behavior classes: Contingencies, beliefs, and verbal behavior. *Journal of Behavior Therapy and Experimental Psychiatry, 26*, 191–200. [ch 19, 23]

Catania, A. C. (1995b). Selection in biology and behavior. In J. T. Todd & E. K. Morris (Eds.), *Modern perspectives on B. F. Skinner and contemporary behaviorism* (pp. 185–194). Westport, CT: Greenwood. [ch 3]

Catania, A. C. (1996). *Responses and time in the quantitative analysis of the delay-of-reinforcement gradient.* Paper presented at the Society for Quantitative Analyses of Behavior, San Francisco. [ch 15]

Catania, A. C. (1997a). An orderly arrangement of well-known facts: Retrospective review of B. F. Skinner's "Verbal behavior". *Contemporary Psychology, 42*, 967–970. [ch 20, 25]

Catania, A. C. (1997b). Remembering Nat Schoenfeld. *Behavior Analyst, 20*, 31–36. [ch 20]

Catania, A. C. (2000). Ten points every behavior analyst needs to remember about reinforcement. In J. C. Leslie & D. E. Blackman (Eds.), *Experimental and applied analyses of human behavior* (pp. 23–37). Reno, NV: Context.

Catania, A. C. (2001a). Positive psychology and positive reinforcement. *American Psychologist, 56*, 86–87. [ch 5]

Catania, A. C. (2001b). Three varieties of selection and their implications for the origins of language. In G. Györi (Ed.), *Language evolution: Biological, linguistic and philosophical perspectives* (pp. 55–71). New York/Bern: Peter Lang. [ch 3, 19, 24, 30]

Catania, A. C. (2002). The watershed years of 1958–1962 in the Harvard Pigeon Laboratory. *Journal of the Experimental Analysis of Behavior, 77*, 327–345. [ch 10,15]

Catania, A. C. (2003a). Verbal governance, verbal shaping, and attention to verbal stimuli. In K. A. Lattal & P. N. Chase (Eds.), *Behavior theory and philosophy* (pp. 301–321). New York: Kluwer/Academic Press. [ch 24]

Catania, A. C. (2003b). Why behavior should matter to linguists. *Behavioral and Brain Sciences, 26*, 670–672. [ch 13, 25]

Catania, A. C. (2004). Antecedents and consequences of words. *European Journal of Behavior Analysis, 5*, 55–64. [ch 24]

Catania, A. C. (2005a). Attention-deficit/hyperactivity disorder (ADHD): Delay-of-reinforcement gradients and other behavioral mechanisms. *Behavioral and Brain Sciences, 28*, 419–424. [ch 16]

Catania, A. C. (2005b). Attention-deficit/hyperactivity disorder (ADHD): One process or many? *Behavioral and Brain Sciences, 28*, 446–450. [ch 16]

Catania, A. C. (2005c). The nonmaintenance of behavior by noncontingent reinforcement. *European Journal of Behavior Analysis, 6*, 89–94. [ch 15]

Catania, A. C. (2005d). The operant reserve: A computer simulation in (accelerated) real time. *Behavioural Processes, 69*, 257–278. [ch 5, 15, 29]

Catania, A. C. (2006). Operant contingencies: Responses and their consequences. *European Journal of Behavior Analysis, 7*, 99–102. [ch 15]

Catania, A. C. (2011a). Model versus model or model versus data? A commentary on Berg and McDowell's "Quantitative, steady-state properties of Catania's computational model of the operant reserve." *Behavioural Processes, 87*, 84–87. [ch 15]

Catania, A. C. (2011b). On Baum's public claim that he has no significant private events: A commentary on "Behaviorism, private events, and the molar view of behavior". *Behavior Analyst, 34*, 227–236. [ch 22]

Catania, A. C. (2012). *Behavior Side Story program.* Cornwall-on-Hudson: Sloan. [ch 20]

Catania, A. C., & Brigham, T. A. (Eds.). (1978). *Handbook of applied behavior analysis: Social and instructional processes*. New York: Irvington. [ch 30]

Catania, A. C., & Cerutti, D. (1986). Some nonverbal properties of verbal behavior. In T. Thompson & M. D. Zeiler (Eds.), *Analysis and integration of behavioral units* (pp. 185–211). Hillsdale, NJ: Erlbaum. [ch 13]

Catania, A. C., DeLeon, I. G., & Cataldo, M. F. (2001). A master's program in applied behavior analysis: Contingencies for initiation and maintenance. *Behavior Analysis Today, 1*(4), 58–63. [ch 5]

Catania, A. C., & Dobson, R. (1972). Concurrent performances: Rate and accuracy of free-operant oddity responding. *Journal of the Experimental Analysis of Behavior, 17*, 25–35. [ch 13]

Catania, A. C., & Gill, C. A. (1964). Inhibition and behavioral contrast. *Psychonomic Science, 1*, 257–258. [ch 16]

Catania, A. C., & Harnad, S. (Eds.). (1988). *The selection of behavior: The operant behaviorism of B. F. Skinner*. New York: Cambridge University Press. [ch 9]

Catania, A. C., Horne, P. J., & Lowe, C. F. (1989a). Transfer of function across members of an equivalence class. *Analysis of Verbal Behavior, 7*, 99–110. [ch 12]

Catania, A. C., & Keller, K. J. (1981). Contingency, contiguity, correlation, and the concept of causation. In P. Harzem & M. D. Zeiler (Eds.), *Predictability, correlation, and contiguity. Advances in analysis of behaviour, Volume 2* (pp. 125–167). Chichester: Wiley. [ch 5, 15]

Catania, A. C., Lowe, C. F., & Horne, P. J. (1990a). Nonverbal behavior correlated with the shaped verbal behavior of children. *Analysis of Verbal Behavior, 8*(43–55). [ch 23]

Catania, A. C., Matthews, B. A., & Shimoff, E. (1982). Instructed versus shaped human verbal behavior: Interactions with nonverbal responding. *Journal of the Experimental Analysis of Behavior, 38*, 233–248. [ch 23]

Catania, A. C., Matthews, B. A., & Shimoff, E. (1989b). *Behavior on a disk: Computer programs in the psychology of learning*. Ellicott City, MD: CMS Software. [ch 29]

Catania, A. C., Matthews, B. A., & Shimoff, E. (1990b). Properties of rule-governed behaviour and their implications. In D. E. Blackman & H. Lejeune (Eds.), *Behaviour analysis in theory and practice: Contributions and controversies* (pp. 215–230). Hove & London, UK: Erlbaum. [ch 23]

Catania, A. C., Matthews, T. J., Silverman, P. J., & Yohalem, R. (1977). Yoked variable-interval and variable-ratio responding in pigeons. *Journal of the Experimental Analysis of Behavior, 28*, 155–161. [ch 15]

Catania, A. C., Ono, K., & de Souza, D. (2000). Sources of novel behavior: Stimulus control arranged for different response dimensions. *European Journal of Behavior Analysis, 1*, 23–32. [ch 11, 13]

Catania, A. C., & Reynolds, G. S. (1968). A quantitative analysis of the behavior maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior, 11*, 327–383. [ch 15]

Catania, A. C., & Sagvolden, T. (1980). Preference for free choice over forced choice in pigeons. *Journal of the Experimental Analysis of Behavior, 34*, 77–86. [ch 14]

Catania, A. C., Sagvolden, T., & Keller, K. J. (1988). Reinforcement schedules: retroactive and proactive effects of reinforcers inserted into fixed-interval performances. *Journal of the Experimental Analysis of Behavior, 49*, 49–73. [ch 16]

Catania, A. C., & Shimoff, E. (1998). The experimental analysis of verbal behavior. *Analysis of Verbal Behavior, 15*, 97–100. [ch 23]

Catania, A. C., Shimoff, E., & Matthews, B. A. (1989c). An experimental analysis of rule-governed behavior. In S. C. Hayes (Ed.), *Rule-governed behavior: Cognition, contingencies, and instructional control* (pp. 119–150). New York: Plenum. [ch 23]

Catania, A. C., Silverman, P. J., & Stubbs, D. A. (1974). Concurrent performances: Stimulus-control gradients during signalled and unsignalled schedules of concurrent reinforcement. *Journal of the Experimental Analysis of Behavior, 21*, 99–107. [ch 11]

Catania, A. C., Yohalem, R., & Silverman, P. J. (1980). Contingency and stimulus change in chained schedules of reinforcement. *Journal of the Experimental Analysis of Behavior, 33*, 213–219. [ch 16]

Catania, K. C., & Kaas, J. H. (1996). The unusual nose and brain of the star-nosed mole. *BioScience, 46*, 578–586. [ch 30]

Catania, K. C., & Kaas, J. H. (1997). The mole nose instructs the brain. *Somatosensory and Motor Research, 14*, 56–58. [ch 20]

Cermak, L. S., & Craik, F. I. M. (Eds.). (1979). *Levels of processing in human memory*. Hillsdale, NJ: Erlbaum. [ch 26]

Chabris, C. F., & Hearst, E. S. (2003). Visualization, pattern recognition, and forward search: Effects of playing speed and sight of the position on grandmaster chess errors. *Cognitive Science, 27*, 637–648. [ch 29]

Chadwick, P. D. J., Lowe, C. F., Horne, P. J., & Higson, P. J. (1994). Modifying delusions: The role of empirical testing. *Behavior Therapy, 25*, 35–49. [ch 23]

Chang, J., & Halliday, J. (2005). *Mao: The unknown story*. New York: Knopf. [ch 24]

Charney, R. (1980). Speech roles and the development of personal pronouns. *Journal of Child Language, 7*, 509–528. [ch 25]

Chase, S. (1938). *The tyranny of words*. New York: Harcourt, Brace & World. [ch 23]

Chenoweth, E., & Stephan, M. J. (2011). *Why civil resistance works: The strategic logic of nonviolent conflict*. New York: Columbia University Press. [ch 30]

Chiat, S. (1982). If I were you and you were me: The analysis of pronouns in a pronoun-reversing child. *Journal of Child Language, 9*, 359–379. [ch 25]

Chomsky, N. (1959). Review of B. F. Skinner's *Verbal behavior*. *Language, 35*, 26–58. [ch 13, 20, 25]

Chomsky, N., & Miller, G. A. (1963). Introduction to the formal analysis of natural languages. In R. D. Luce, R. R. Bush & E. Galanter (Eds.), *Handbook of mathematical psychology. Volume II* (pp. 269–321). New York: Wiley. [ch 25]

Chung, S.-H. (1965). Effects of effort on response rate. *Journal of the Experimental Analysis of Behavior, 8*, 1–7. [ch 16]

Church, R. M. (1963). The varied effects of punishment on behavior. *Psychological Review, 70*, 369–402. [ch 7]

Church, R. M. (1969). Response suppression. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior*. (pp. 111–156). New York: Appleton-Century-Crofts. [ch 7]

Chédotal, A., & Richards, L. J. (2010). Wiring the brain: The biology of neuronal guidance. *Cold Spring Harbor Perspectives in Biology, 2*, a001917. [ch 3, 30]

Clark, H. H., & Sengul, C. J. (1979). In search of referents for nouns and pronouns. *Memory and Cognition, 7*, 35–41. [ch 25]

Coe, M. (1967). "Necking" behaviour in the giraffe. *Journal of Zoology, 151*, 313–321. [ch 3]

Coleman, L., & Clancy, P. E. (1999). *Has a child been molested?* Walnut Creek, CA: Berkeley Creek. [ch 26]

Collier, G. H., & Rovee-Collier, C. K. (1981). A comparative analysis of optimal foraging behavior: Laboratory simulations. In A. C. Kamil & T. D. Sargent (Eds.), *Foraging behavior* (pp. 39–76). New York: Garland. [ch 11]

Conrad, R. (1964). Acoustic confusions in immediate memory. *British Journal of Psychology Monograph Supplements, 55*, 75–84. [ch 26, 28]

Cook, R. G., Wright, A. A., & Sands, S. F. (1991). Interstimulus interval and viewing time effects in monkey list memory. *Animal Learning and Behavior, 19*, 153–163. [ch 26]

Cooper, J. O., Heron, T. E., & Heward, W. L. (2007). *Applied behavior analysis* (2nd ed.). Upper Saddle River, NJ: Pearson. [ch 30]

Cott, A., Pavlovski, R. P., & Black, A. H. (1981). Operant conditioning and discrimination of alpha: Some methodological limitations inherent in response-discrimination experiments. *Journal of Experimental Psychology: General, 110*, 398–414. [ch 19]

Craik, F. I. M. (1985). Paradigms in human memory research. In L. Nilsson & T. Archer (Eds.), *Perspectives on learning and memory* (pp. 197–221). Hillsdale, NJ: Erlbaum. [ch 26]

Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11*, 671–684. [ch 27]

Crain, S. (1991). Language acquisition in the absence of experience. *Behavioral and Brain Sciences, 14*, 597–650. [ch 13, 25]

Crider, A., Schwartz, G. E., & Shnidman, S. (1969). On the criteria for instrumental autonomic conditioning: A reply to Katkin and Murray. *Psychological Bulletin, 71*, 455–461. [ch 10]

Critchfield, T. S. (1993). Signal-detection properties of verbal self-reports. *Journal of the Experimental Analysis of Behavior, 60*, 495–514. [ch 19]

Crowder, R. G. (1976). *Principles of learning and memory*. Hillsdale, NJ: Erlbaum. [ch 27, 28]

Culicover, P. W. (1992). Language learning: Humans. In L. R. Squire (Ed.), *The encyclopedia of learning and memory* (pp. 327–331). New York: Macmillan. [ch 25]

Cuvo, A. (2000). Development and function of consequence classes in operant behavior. *The Behavior Analyst, 23*, 57–68. [ch 14]

Czubaroff, J. (1988). Criticism and response in the Skinner controversies. *Journal of the Experimental Analysis of Behavior, 49*, 321–329. [ch 25]

D'Amato, M. R., Salmon, D. P., Loukas, E., & Tomie, A. (1985). Symmetry and transitivity of conditional relations in monkeys (Cebus apella) and pigeons (Columba livia). *Journal of the Experimental Analysis of Behavior, 44*, 35–47. [ch 12]

Damon, W., & Hart, D. (1982). The development of self-understanding from infancy through adolescence. *Child Development, 53*, 841–864. [ch 25]

Daniels, A. C. (2009). *Oops! 13 management practices that waste time and money*. Atlanta, GA: Performance Management Publications. [ch 14]

Daniels, A. H. (1895). The memory after-image and attention. *American Journal of Psychology, 6*, 558–564. [ch 27]

Darwin, C. (1859). *On the origin of species by means of natural selection*. London: John Murray. [ch 3, 30]

Darwin, C. (1871). *The descent of man*. London: John Murray. [ch 3, 20]

Darwin, C. (1872). *The expression of the emotions in man and animals*. London: John Murray. [ch 18]

Darwin, C. J., Turvey, M. T., & Crowder, R. G. (1972). An auditory analogue of the Sperling partial report procedure: Evidence for brief auditory storage. *Cognitive Psychology, 3*, 255–267. [ch 27]

Davis, H., & Pérusse, R. (1988). Numerical competence in animals: Definitional issues, current evidence, and a new research agenda. *Behavioral and Brain Sciences, 11*, 561–615. [ch 12]

Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum. [ch 16]

Dawkins, R. (1976). *The selfish gene*. New York: Oxford University Press. [ch 3, 19, 22, 29, 30]

Dawkins, R. (1982). *The extended phenotype.* San Francisco: Freeman. [ch 3, 19]

Dawkins, R. (1986). *The blind watchmaker.* New York: Norton. [ch 3]

Dawkins, R., & Krebs, J. R. (1978). Animal signals: Information or manipulation? In J. R. Krebs & N. B. Davies (Eds.), *Behavioral ecology* (pp. 282–309). Sunderland, MA: Sinauer. [ch 19]

Day, W. F. (1969). On certain similarities between the philosophical investigations of Ludwig Wittgenstein and the operationism of B. F. Skinner. *Journal of the Experimental Analysis of Behavior, 12*, 489–506. [ch 22, 25]

de Rose, J. C., McIlvane, W. J., Dube, W. V., Galpin, V. C., & Stoddard, L. T. (1988). Emergent simple discrimination established by indirect relation to differential consequences. *Journal of the Experimental Analysis of Behavior, 50*, 1–20. [ch 12]

De Villiers, J. G., & De Villiers, P. A. (1978). *Language acquisition.* Cambridge, MA: Harvard University Press. [ch 25]

De Villiers, P. A., & De Villiers, J. G. (1974). On this, that, and the other: Nonegocentrism in very young children. *Journal of Experimental Child Psychology, 18*, 438–447. [ch 25]

de Waal, F. (1989). *Peacemaking among primates.* Cambridge, MA: Harvard University Press. [ch 19]

DeCasper, A. J., & Fifer, W. P. (1980). Of human bonding: Newborns prefer their mothers' voices. *Science, 208*, 1174–1176. [ch 20]

degli Espinosa, F. (2011). *Verbal behaviour development for children with autism.* University of Southampton, Southampton. [ch 19, 22]

Delgado, L. E., & Lutzger, J. R. (1988). Training young parents to identify and report their children's illnesses. *Journal of Applied Behavior Analysis, 21*, 311–319. [ch 11]

Delius, J. D., & Nowak, B. (1982). Visual symmetry recognition by pigeons. *Psychological Research, 44*, 199–212. [ch 12]

Deluty, M. Z. (1978). Self-control and impulsiveness involving aversive events. *Journal of Experimental Psychology: Animal Behavior Processes, 4*, 250–266. [ch 16]

Dennett, D. C. (1987). *The intentional stance.* Cambridge, MA: MIT Press. [ch 19]

Derrickson, J. G., Neef, N. A., & Cataldo, M. F. (1993). Effects of signalling invasive procedures on a hospitalized infant's affective behavior. *Journal of Applied Behavior Analysis, 26*, 133–134. [ch 18]

Deutsch, R. (1974). Conditioned hypoglycemia: A mechanism for saccharin-induced sensitivity to insulin in the rat. *Journal of Comparative and Physiological Psychology, 86*, 350–358. [ch 17]

Deutscher, G. (2005). *The unfolding of language.* New York: Holt. [ch 25]

Deutscher, G. (2010). *Through the language glass: Why the world looks different in other languages.* New York: Holt. [ch 25]

Dews, P. B. (1960). Free-operant behavior under conditions of delayed reinforcement: I. CRF-type schedules. *Journal of the Experimental Analysis of Behavior, 3*, 221–234. [ch 15]

Dews, P. B. (1962). The effect of multiple S^{Δ} periods on responding on a fixed-interval schedule. *Journal of the Experimental Analysis of Behavior, 5*, 369–374. [ch 15]

Dews, P. B. (1970). Drugs in psychology. A commentary on T. Thompson & C. R. Schuster's Behavioral Pharmacology. *Journal of the Experimental Analysis of Behavior, 13*, 395–406. [ch 16]

Dill, L. M. (1974). The escape response of the zebra danio (*Brachydanio rerio*) II. The effect of experience. *Animal Behavior, 22*, 723–730. [ch 4]

Dinsmoor, J. A. (1983). Observing and conditioned reinforcement. *Behavioral and Brain Sciences, 6*, 693–728. [ch 16]

Dinsmoor, J. A. (1985). The role of observing and attention in establishing stimulus control. *Journal of the Experimental Analysis of Behavior, 43*(3), 365–381. [ch 16]

Dinsmoor, J. A. (1989). Selective observing: Pigeons turn discriminative stimuli on and off by pecking separate keys. In N. W. Bond & D. A. T. Siddle (Eds.), *Psychobiology: Issues and Applications* (pp. 443–456). North-Holland: Elsevier. [ch 16]

Dinsmoor, J. A. (1995). Stimulus control. *The Behavior Analyst, 18*, 51–68 and 253–269. [ch 11, 16]

Donahoe, J. W., Burgos, J. E., & Palmer, D. C. (1993). A selectionist approach to reinforcement. *Journal of the Experimental Analysis of Behavior, 60*, 17–40. [ch 9, 30]

Donahoe, J. W., & Palmer, D. C. (1989). The interpretation of complex human behavior: Some reactions to Parallel Distributed Processing. *Journal of the Experimental Analysis of Behavior, 51*, 399–416. [ch 28, 29]

Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior.* Boston: Allyn and Bacon. [ch 11, 17, 30]

Donahoe, J. W., & Vegas, R. (2004). Pavlovian Conditioning: The CS-UR Relation. *Journal of Experimental Psychology: Animal Behavior Processes, 30*, 17–33. [ch 17]

Dorsey, M. F., Iwata, B. A., Ong, P., & McSween, T. E. (1980). Treatment of self-injurious behavior using a water mist: Initial response suppression and generalization. *Journal of Applied Behavior Analysis, 13*, 343–353. [ch 7]

Dreyfus, H. L. (1992). *What computers still can't do.* Cambridge, MA: MIT Press. [ch 25, 29]

Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over machine.* New York: Macmillan. [ch 29]

Dube, W. V., MacDonald, R. P. F., Mansfield, R. C., Holcomb, W. L., & Ahearn, W. H. (2004). Toward a behavioral analysis of joint attention. *The Behavior Analyst, 27*, 197–207. [ch 11, 19, 22]

Dube, W. V., & McIlvane, W. J. (1995). Stimulus-reinforcer relations and emergent matching to sample. *Psychological Record, 45*, 591–612. [ch 14]

Dube, W. V., McIlvane, W. J., Callahan, T. D., & Stoddard, L. T. (1993). The search for stimulus equivalence in nonverbal organisms. *Psychological Record, 43*, 761–778. [ch 12, 14]

Dube, W. V., McIlvane, W. J., Mackay, H. A., & Stoddard, L. T. (1987). Stimulus class membership established via stimulus-reinforcer relations. *Journal of the Experimental Analysis of Behavior, 47*, 159–175. [ch 14]

Dube, W. V., McIlvane, W. J., Maguire, R. W., Mackay, H. A., & Stoddard, L. T. (1989). Stimulus class formation and stimulus-reinforcer relations. *Journal of the Experimental Analysis of Behavior, 51*, 51–65. [ch 14]

Duncker, K. (1945). On problem solving. *Psychological Monographs, 58*(5, Whole No. 270). [ch 29]

Dunham, P. J. (1977). The nature of reinforcing stimuli. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 98–124). Englewood Cliffs, NJ: Prentice-Hall. [ch 6]

Dworkin, B. R., & Miller, N. E. (1986). Failure to replicate visceral learning in the acute curarized rat preparation. *Behavioral Neuroscience, 100*, 299–314. [ch 10]

Ebbinghaus, H. (1885/1964). *Über das Gedächtnis (Memory, tr. H. A. Ruger & C. E. Bussenius).* Leipzig/New York: Duncker & Humblot/Dover. [ch 21, 27, 30]

Eckerman, D. A., Hienz, R. D., Stern, S., & Kowlowitz, V. (1980). Shaping the location of a pigeon's peck: Effect of rate and size of shaping steps. *Journal of the Experimental Analysis of Behavior, 33*, 299–310. [ch 9]

Edelman, G. M. (1987). *Neural Darwinism.* New York: Basic Books. [ch 3, 30]

Eimas, P. D., & Miller, J. L. (1992). Organization in the perception of speech by young infants. *Psychological Science, 3*, 340–345. [ch 4]

Eimas, P. D., Siqueland, E. R., Jusczyk, P., & Vigorito, J. (1971). Speech perception in early infancy. *Science, 171*, 303–306. [ch 4, 20]

Eisenberger, R., & Cameron, J. (1996). Detrimental effects of reward: Reality or myth? *American Psychologist, 51*, 1153–1166. [ch 14]

Eisenberger, R., Karpman, M., & Trattner, T. (1967). What is the necessary and sufficient condition for reinforcement in the contingency situation? *Journal of Experimental Psychology, 74*, 342–350. [ch 6]

Epstein, R. (1981). On pigeons and people: A preliminary look at the Columban Simulation Project. *Behavior Analyst, 4*, 43–55. [ch 29]

Epstein, R., Lanza, R. P., & Skinner, B. F. (1980). Symbolic communication between two pigeons (Columba livia domestica). *Science, 207*, 543–545. [ch 24]

Epstein, R., Lanza, R. P., & Skinner, B. F. (1981). "Self-awareness" in the pigeon. *Science, 212*, 695–696. [ch 19]

Epstein, R., & Skinner, B. F. (1980). Resurgence of responding after the cessation of response-independent reinforcement. *Proceedings of the National Academy of Sciences, 77*, 6251–6253. [ch 5]

Erdelyi, M. H., & Kleinbard, J. (1978). Has Ebbinghaus decayed with time? The growth of recall (hypermnesia) over days. *Journal of Experimental Psychology: Human Learning and Memory, 4*, 275–289. [ch 27]

Erjavec, M., & Horne, P. J. (2008). Determinants of imitation of hand-to-body gestures in 2– and 3–year-old children. *Journal of the Experimental Analysis of Behavior, 89*, 183–207. [ch 19, 22]

Esper, E. A. (1973). *Analogy and association in linguistics and psychology.* Athens, GA: University of Georgia Press. [ch 13, 22, 25]

Estes, W. K. (1940). A visual form of the verbal summator. *Psychological Record, 4*, 174–180. [ch 20]

Estes, W. K. (1944). An experimental study of punishment. *Psychological Monographs, 57, No. 263.* [ch 7]

Estes, W. K. (1964). All-or-none processes in learning and retention. *American Psychologist, 19*, 16–25. [ch 21]

Estes, W. K. (1971). Reward in human learning: Theoretical issues and strategic choice points. In R. Glaser (Ed.), *The nature of reinforcement* (pp. 16–36). New York: Academic Press. [ch 12]

Estes, W. K., & Skinner, B. F. (1941). Some quantitative properties of anxiety. *Journal of Experimental Psychology, 29*, 390–400. [ch 18]

Evans, R. J. (2004). *The coming of the Third Reich.* New York: Penguin. [ch 24]

Evans, R. J. (2005). *The Third Reich in power.* New York: Penguin. [ch 24]

Evans, R. J. (2009). *The Third Reich at war.* New York: Penguin. [ch 24]

Falk, J. L. (1977). The origin and functions of adjunctive behavior. *Animal Learning and Behavior, 5*, 325–335. [ch 4]

Fantino, E., & Abarca, N. (1985). Choice, optimal foraging, and the delay-reduction hypothesis. *Behavioral and Brain Sciences, 8*, 315–330. [ch 16]

Favell, J. E., McGimsey, J. F., & Jones, M. L. (1978). The use of physical restraint in the treatment of self-injury and as positive reinforcement. *Journal of Applied Behavior Analysis, 11*, 225–241. [ch 6]

Fearing, F. (1930). *Reflex action.* Baltimore: Williams and Wilkins. [ch 4]

Fedorchak, P. M., & Bolles, R. C. (1986). Differential outcome effect using a biologically neutral outcome difference. *Journal of Experimental Psychology: Animal Behavior Processes, 12*, 125–130. [ch 14]

Felton, M., & Lyon, D. O. (1966). The post-reinforcement pause. *Journal of the Experimental Analysis of Behavior, 9*, 131–164. [ch 15]

Fennelly, M., Galletly, D. C., & Purdie, G. I. (1991). Is caffeine withdrawal the mechanism of postoperative headache? *Anesthesia and Analgesia, 72*, 449–453. [ch 18]

Ferrari, M., & Harris, S. L. (1981). The limits and motivating potential of sensory stimuli as reinforcers for autistic children. *Journal of Applied Behavior Analysis, 14*, 339–343. [ch 6]

Ferrell, J. E., Jr. (2012). Bistability, bifurcations, and Waddington's epigenetic landscape. *Current Biology, 22*, R458–R466. [ch 3]

Ferster, C. B. (1958). Control of behavior in chimpanzees and pigeons by time out from positive reinforcement. *Psychological Monographs, 72*(8, Whole No. 461). [ch 7, 8]

Ferster, C. B. (1960). Intermittent reinforcement of matching to sample in the pigeon. *Journal of the Experimental Analysis of Behavior, 3*, 259–272. [ch 12]

Ferster, C. B. (1978). Is operant conditioning getting bored with behavior? *Journal of the Experimental Analysis of Behavior, 29*, 347–349. [ch 15]

Ferster, C. B., & DeMyer, M. K. (1961). The development of performances in autistic children in an automatically controlled environment. *Journal of Chronic Diseases, 13*, 312–345. [ch 9]

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts. [ch 3, 5, 10, 15, 16, 30]

Fetterman, J. G. (1996). Dimensions of stimulus complexity. *Journal of Experimental Psychology: Animal Behavior Processes, 22*, 3–18. [ch 12]

Field, T. M., Woodson, R., Greenberg, R., & Cohen, D. (1982). Discrimination and imitation of facial expressions by neonates. *Science, 218*, 179–181. [ch 19]

Fields, L., Adams, B. J., Verhave, T., & Newman, S. (1990). The effects of nodality on the formation of equivalence classes. *Journal of the Experimental Analysis of Behavior, 53*, 345–358. [ch 13]

Fields, L., Arntzen, E., Nartey, R. K., & Eilifsen, C. (2012). Effects of a meaningful, a discriminative, and a meaningless stimulus on equivalence class formation. *Journal of the Experimental Analysis of Behavior, 97*, 163–181. [ch 13]

Fields, L., Landon-Jimenez, D. V., Buffington, D., M., & Adams, B. J. (1995). Maintained nodal-distance effects in equivalence classes. *Journal of the Experimental Analysis of Behavior, 64*, 129–145. [ch 13]

Fields, L., Tittelbach, D., Shamoun, K., Watanabe, M., Fitzer, A., & Matneja, P. (2007). The effects of differential training procedures on linked perceptual class formation. *Journal of the Experimental Analysis of Behavior, 87*, 97–119. [ch 13]

Fields, L., & Verhave, T. (1987). The structure of equivalence classes. *Journal of the Experimental Analysis of Behavior, 48*, 317–332. [ch 13]

Fillenbaum, S. (1966). Memory for gist: Some relevant variables. *Language and Speech, 9*, 217–227. [ch 25]

Findley, J. D. (1962). An experimental outline for building and exploring multioperant behavior repertoires. *Journal of the Experimental Analysis of Behavior, 5*, 113–166. [ch 15]

Findley, J. D., & Brady, J. V. (1965). Facilitation of large ratio performance by use of conditioned reinforcement. *Journal of the Experimental Analysis of Behavior, 8*, 125–129. [ch 16]

Fischer, K. W. (1980). A theory of cognitive development: The control and construction of hierarchies of skills. *Psychological Review, 87*, 477–531. [ch 28]

Fisher, W. W., Piazza, C. C., & Roane, H. S. (Eds.). (2011). *Handbook of applied behavior analysis*. New York: Guilford. [ch 30]

Fisher, W. W., Piazza, C. C., Bowman, L. G., Hagopian, L. P., Owens, J. C., & Slevin, I. (1992). A comparison of two approaches for identifying reinforcers for persons with severe and profound disabilities. *Journal of Applied Behavior Analysis, 25*, 491–498. [ch 6]

Fisher, W. W., & Bouxsein, K. (2011). Developing function-based reinforcement procedures for problem behavior. In W. W. Fisher, C. C. Piazza & H. S. Roane (Eds.), *Handbook of applied behavior analysis* (pp. 335–347). New York: Guikford. [ch 7, 8]

Flanagan, B., Goldiamond, I., & Azrin, N. H. (1958). Operant stuttering: The control of stuttering behavior through reponse-contingent consequences. *Journal of the Experimental Analysis of Behavior, 1*, 173–177. [ch 14]

Flavell, J. H., Friedrichs, A. G., & Hoyt, J. D. (1970). Developmental changes in memorization processes. *Cognitive Psychology, 1*, 324–340. [ch 26]

Fodor, J. A., & Bever, T. G. (1965). The psychological reality of linguistic segments. *Journal of Verbal Learning and Verbal Behavior, 4*, 414–420. [ch 25]

Fortey, R. (1998). *Life: A natural history of the first four billion years of life on earth*. New York: Knopf. [ch 3]

Foss, D. J. (1988). Experimental psycholinguistics. *Annual Review of Psychology, 39*, 301–348. [ch 22]

Fowler, C. A., Wolford, G., Slade, R., & Tassinary, L. (1981). Lexical access with and without awareness. *Journal of Experimental Psychology: General, 110*, 341–362. [ch 20]

Fowler, H., & Trapold, M. A. (1962). Escape performance as a function of delay of reinforcement. *Journal of Experimental Psychology, 63*, 464–467. [ch 8]

Fraenkel, G. S., & Gunn, D. L. (1961). *The orientation of animals*. New York: Dover. [ch 4]

Francisco, M. T., Borrero, J. C., & Sy, J. R. (2008). Evaluation of absolute and relative reinforcer value using progressive-ratio schedules. *Journal of Applied Behavior Analysis, 41*, 189–202. [ch 14]

Freeman, R. D., Mitchell, D. E., & Millidot, M. A. (1972). Neural effect of partial visual deprivation in humans. *Science, 175*, 1384–1386. [ch 6]

Freud, S. (1917). *Wit and its relation to the unconscious (trans. A. A. Brill)*. London: Allen & Unwin. [ch 30]

Frey, P. W. (Ed.). (1977). *Chess skill in man and machine*. New York: Springer-Verlag. [ch 29]

Friedlander, B. Z. (1968). The effect of speaker identity, voice inflection, vocabulary, and message redundancy on infants' selection of vocal reinforcement. *Journal of Experimental Child Psychology, 6*, 443–459. [ch 20]

Friman, P. C., Hayes, S. C., & Wilson, K. G. (1998). Why behavior analysts should study emotion: The example of anxiety. *Journal of Applied Behavior Analysis, 31*, 137–156.

Friman, P. C., & Piazza, C. C. (2011). Behavioral pediatrics: Integrating applied behaivor analysis with pediatric medicine. In W. W. Fisher, C. C. Piazza & H. S. Roane (Eds.), *Handbook of applied behavior analysis* (pp. 433–450). New York: Guilford. [ch 14]

Fromkin, V. A. (1971). The non-anomalous nature of anomalous utterances. *Language, 47*, 27–52. [ch 20]

Fujita, K. (1983). Formation of the sameness-difference concept by Japanese monkeys from a small number of color stimuli. *Journal of the Experimental Analysis of Behavior, 40*, 289–300. [ch 12]

Galbicka, G., Kautz, M. A., & Jagers, T. (1993). Response acquisition under targeted percentile schedules: A continuing quandary for molar models of operant behavior. *Journal of the Experimental Analysis of Behavior, 60*, 171–184. [ch 15]

Galef, B. G., Jr., & Stein, M. (1985). Demonstrator influence on observer diet preference: Analysis of critical social interactions and olfactory signals. *Animal Learning and Behavior, 13*, 31–38. [ch 19]

Gallistel, C. R. (1980). *The organization of action*. Hillsdale, NJ: Erlbaum. [ch 4, 18]

Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press. [ch 6, 12, 29]

Gallup, G. G., Jr. (1979). Self-awareness in primates. *American Scientist, 67*, 417–419. [ch 19]

Galton, F. (1879). Psychometric experiments. *Brain, 2*, 149–162. [ch 21]

Garcia, E., Baer, D. M., & Firestone, I. (1971). The development of generalized imitation within topographically determined boundaries. *Journal of Applied Behavior Analysis, 4*, 101–112. [ch 19]

Garcia, J., & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science, 4*, 123–124. [ch 18]

Gardner, R. A., & Gardner, B. T. (1969). Teaching sign language to a chimpanzee. *Science, 165*, 664–672. [ch 24]

Garner, W. R. (1974). *The processing of information and structure*. Hillsdale, NJ: Erlbaum. [ch 28]

Garrett, M., Bever, T. G., & Fodor, J. (1966). The active use of grammar in speech perception. *Perception and Psychophysics, 1*, 30–32. [ch 25]

Garwood, C. (2007). *Flat earth*. New York: St. Martin's Press. [ch 24]

Gellately, R. (2007). *Lenin, Stalin, and Hitler: The age of social catastrophe*. New York: Knopf. [ch 24]

Geller, I. (1960). The acquisition and extinction of conditioned suppression as a function of the base-line reinforcer. *Journal of the Experimental Analysis of Behavior, 3*, 235–240. [ch 18]

Gewirtz, J. L., & Pelaez-Nogueras, M. (1992). B. F. Skinner's legacy to human infant behavior and development. *American Psychologist, 47*, 1411–1422. [ch 6]

Gibbon, J., Farrell, L., Locurto, C. M., Duncan, H. J., & Terrace, H. S. (1980). Partial reinforcement in autoshaping with pigeons. *Animal Learning and Behavior, 8*, 45–59. [ch 17]

Gibson, E. J. (1965). Learning to read. *Science, 148*, 1066–1072. [ch 11, 20, 26, 29]

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin. [ch 1, 11]

Gilbert, T. F. (1978). *Human competence*. New York: McGraw-Hill.

Gillan, D. J. (1981). Reasoning in the chimpanzee: II. Transitive inference. *Journal of Experimental Psychology: Animal Behavior Processes, 7*, 150–164. [ch 24]

Gleick, J. (1987). *Chaos*. New York: Viking. [ch 3]

Gleitman, L. R., & Rozin, P. (1973). Teaching reading by use of a syllabary. *Reading Research Quarterly, 8*, 447–483. [ch 20, 21]

Glencross, D. J. (1977). Control of skilled movements. *Psychological Bulletin, 84*, 14–29. [ch 18]

Glucksberg, S., Gildea, P., & Bookin, H. B. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior, 21*, 85–98. [ch 22]

Goh, H.-L., Iwata, B. A., & DeLeon, I. G. (2000). Competition between noncontingent and contingent reinforcement schedules during response acquisition. *Journal of Applied Behavior Analysis, 33*, 195–205. [ch 5]

Goldiamond, I. (1966). Perception, language, and conceptualization rules. In B. Kleinmuntz (Ed.), *Problem solving* (pp. 183–224). New York: Wiley. [ch 21]

Goldiamond, I. (1976). Self-reinforcement. *Journal of Applied Behavior Analysis, 9*, 509–514. [ch 19]

Goldiamond, I., & Thompson, D. M. (2002). A functional analysis of behavior and its extensions. In P. T. Andronis (Ed.), *The Blue Books: Goldiamond's functional analysis of behavior (CD)*. Concord, MA: Cambridge Center for Behavioral Studies. [ch 13, gloss]

Gollub, L. R. (1977). Conditioned reinforcement: Schedule effects. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 288–312). Englewood Cliffs, NJ: Prentice Hall. [ch 16]

Goodman, G. S. (2006). Children's eyewitness memory: A modern history and contemporary commentary. *Journal of Social Issues, 62*, 811–832. [ch 26]

Goodsell, D. S. (2009). *The machinery of life* (2nd ed.). New York: Springer. [ch 3]

Gormezano, I. (1972). Investigations of defense and reward conditioning in the rabbit. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II* (pp. 151–181). New York: Appleton-Century-Crofts. [ch 17]

Gould, S. J. (1975). Darwin's "Big Book." *Science, 188*, 824–827. [ch 3, 30]

Gould, S. J. (1977). *Ontogeny and phylogeny*. Cambridge, MA: Harvard University Press. [ch 3, 19, 30]

Gould, S. J. (1996a). *Full house*. New York: Harmony Books. [ch 3]

Gould, S. J. (1996b). The tallest tale. *Natural History, 105*(5, May), 18–23, 54–57. [ch 3]

Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London B, 205*, 581–598. [ch 3, 25]

Gouzoules, S., Gouzoules, H., & Marler, P. (1984). Rhesus monkey (Macaca mulatta) screams: Representational signalling in the recruitment of agonistic aid. *Animal Behaviour, 32*, 182–193. [ch 19, 24]

Graefe, T. M., & Watkins, M. J. (1980). Picture rehearsal: An effect of selectively attending to pictures no longer in view. *Journal of Experimental Psychology: Human Learning and Memory, 6*, 156–162. [ch 28]

Gray, J. (1953). *How animals move*. London: Cambridge University Press. [ch 18]

Green, C. W., Reid, D. H., White, L. K., Halford, R. C., Brittain, D. P., & Gardner, S. M. (1988). Identifying reinforcers for persons with profound handicaps: Staff opinion versus systematic assessment of preferences. *Journal of Applied Behavior Analysis, 21*, 31–43. [ch 6]

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley. [ch 9, 11, 23]

Green, L., & Freed, D. E. (1998). Behavioral economics. In W. Donohue (Ed.), *Learning and behavior therapy* (pp. 274–300). Boston: Allyn & Bacon. [ch 16]

Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin, 130*, 769–792. [ch 16]

Greenberg, J. H. (Ed.). (1966). *Universals of language*. Cambridge, MA: MIT Press. [ch 25]

Greenspoon, J. (1955). The reinforcing effect of two spoken sounds on the frequency of two responses. *American Journal of Psychology, 68*, 409–416. [ch 23]

Greer, R. D., & McDonough, S. H. (1999). Is the learn unit a fundamental measure of pedagogy? *Behavior Analyst, 22*, 5–16. [ch 6, 13, 30]

Greer, R. D., & Ross, D. E. (2004). Verbal behavior analysis: A program of research in the induction and expansion of complex verbal behavior. *Journal of Early and Intensive Behavior Intervention, 1*, 141–165. [ch 6, 13, 30]

Grosch, J., & Neuringer, A. (1981). Self-control in pigeons under the Mischel paradigm. *Journal of the Experimental Analysis of Behavior, 35*, 3–21. [ch 16]

Gurteen, P. M., Horne, P. J., & Erjavec, M. (2011). Rapid word learning in 13– and 17–month-olds in a naturalistic two-word procedure: Looking versus reaching measures. *Journal of Experimental Child Psychology, 109*, 201–217. [ch 22]

Guthrie, E. R. (1935). *The psychology of learning*. New York: Harper. [ch 30]

Guthrie, E. R., & Horton, G. P. (1946). *Cats in a puzzle box*. New York: Rinehart. [ch 5, 17]

Guttman, N. (1959). Generalization gradients around stimuli associated with different reinforcement schedules. *Journal of Experimental Psychology, 58*, 335–340. [ch 11]

Guttman, N., & Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology, 51*, 79–88. [ch 11]

Haber, R. N. (1969). Eidetic images. *Scientific American, 220*, 36–44. [ch 28]

Hackenberg, T. D. (1992). Choice in situations of time-based diminishing returns: Immediate versus delayed consequences of action. *Journal of the Experimental Analysis of Behavior, 57*, 67–80. [ch 16]

Hadamard, J. (1949). *The psychology of invention in the mathematical field*. Princeton, NJ: Princeton University Press. [ch 26]

Hagopian, L. P., Crockett, J. L., van Stone, M., DeLeon, I. G., & Bowman, L. G. (2000). Effects of noncontingent reinforcement on problem behavior and stimulus engagement: The role of satiation, extinction, and alternative reinforcement. *Journal of Applied Behavior Analysis, 33*, 433–449. [ch 5]

Hailman, J. P. (1969). How an instinct is learned. *Scientific American, 221*(6), 98–106. [ch 2]

Hall, R. V., Lund, D., & Jackson, D. (1968). Effects of teacher attention on study behavior. *Journal of Applied Behavior Analysis, 1*, 1–12. [ch 6]

Hall, W. G., & Oppenheim, R. W. (1987). Developmental psychobiology: Prenatal, perinatal, and early postnatal aspects of behavioral development. *Annual Review of Psychology, 38*, 91–128. [ch 3, 4]

Hampl, K. F., Schneider, M. C., Rüttiman, U., Ummenhofer, W., & Drewe, J. (1995). Perioperative administration of caffeine tablets for prevention of postoperative headaches. *Canadian Journal of Anesthesia, 42*, 789–792. [ch 18]

Hanson, H. M. (1959). Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology, 58*, 321–334. [ch 11]

Harlow, H. F. (1949). The formation of learning sets. *Psychological Review, 56*, 51–65. [ch 12]

Harlow, H. F., & Harlow, M. K. (1966). Learning to love. *American Scientist, 54*, 244–272. [ch 18]

Harlow, H. F., & Suomi, S. J. (1970). Nature of love: Simplified. *American Psychologist, 25*, 161–168. [ch 18]

Harnad, S. (2003). Categorical perception. In L. Nadel (Ed.), *Encyclopedia of cognitive science*. New York: Macmillan. [ch 20]

Harris, A. H., & Turkkan, J. S. (1981). Generalization of conditioned blood pressure elevations: Schedule and stimulus control effects. *Physiology and Behavior, 26*, 935–940. [ch 10]

Harris, C. S. (1965). Perceptual adaptation to inverted, reversed, and displaced vision. *Psychological Review, 72*, 419–444. [ch 6]

Harris, M. (1977). *Cannibals and kings*. New York: Random House. [ch 19, 30]

Harris, R. A. (1993). *The linguistic wars*. New York: Oxford. [ch 25]

Hart, B. M., Reynolds, N. J., Baer, D. M., Brawley, E. R., & Harris, F. R. (1968). Effect of contingent and non-contingent social reinforcement on the cooperative play of a preschool child. *Journal of Applied Behavior Analysis, 1*, 73–76. [ch 5]

Hart, B., & Risley, T. R. (1980). In vivo language intervention: Unanticipated general effects. *Journal of Applied Behavior Analysis, 13*, 407–432. [ch 25]

Hart, B., & Risley, T. R. (1995). *Meaningful differences in the everyday experience of young American children*. Baltimore: Paul H. Brookes. [ch 25]

Hart, B. L. (1973). Reflexive behavior. In G. Berman (Ed.), *Perspectives on animal behavior* (pp. 171–193). Glenview, IL: Scott, Foresman. [ch 4]

Hartline, H. K., & Ratliff, F. (1957). Inhibitory interaction of receptor units in the eye of Limulus. *Journal of General Physiology, 40*, 357–673. [ch 16]

Hartline, H. K., Wagner, H. G., & Ratliff, F. (1956). Inhibition in the eye of Limulus. *Journal of General Physiology, 39*, 651–673. [ch 16]

Hartline, H. K., Wagner, H. G., & Ratliff, F. (1958). Spatial summation of inhibitory influences in the eye of Limulus, and the mutual interaction of receptor units. *Journal of General Physiology, 41*, 1049–1066. [ch 16]

Hartman, A. M. (1975). Analysis of conditions leading to the regulation of water flow by a beaver. *Psychological Record, 25*, 427–431. [ch 8]

Harzem, P., Lee, I., & Miles, T. R. (1976). The effects of pictures on learning to read. *British Journal of Educational Psychology, 46*, 318–322. [ch 30]

Hasher, L., & Griffin, M. (1978). Reconstructive and reproductive processes in memory. *Journal of Experimental Psychology, 4*, 318–330. [ch 26]

Hayes, K., & Hayes, C. (1951). The intellectual development of a home-raised chimpanzee. *Proceedings of the American Philosophical Society, 95*, 105–109. [ch 20, 24]

Hayes, S. C. (1994). Relational frame theory: A functional approach to verbal events. In S. C. Hayes, L. J. Hayes, M. Sato & K. Ono (Eds.), *Behavior analysis of language and cognition* (pp. 9–30). Reno, NV: Context Press. [ch 12, 13]

Hayes, S. C., Barnes-Holmes, D., & Roche, B. (2001). *Relational frame theory: A post-Skinnerian account of human language and cognition*. New York: Kluwer Academic/Plenum. [ch 12, 13]

Hayes, S. C., & Wilson, K. G. (1994). Acceptance and commitment therapy: Altering the verbal support for experiential avoidance. *The Behavior Analyst, 17*, 289–303. [ch 13]

Hayes, S. C., Zettle, R. D., & Rosenfarb, I. (1989). Rule-following. In S. C. Hayes (Ed.), *Rule-governed behavior: Cognition, contingencies, and instructional control* (pp. 191–220). New York: Plenum. [ch 23]

Healy, A. F. (1976). Detection errors on the word *the*: Evidence for reading units larger than letters. *Journal of Experimental Psychology: Human Perception and Performance, 2*, 235–242. [ch 21]

Hearst, E. (1958). The behavioral effects of some temporally defined schedules of reinforcement. *Journal of the Experimental Analysis of Behavior, 1*, 45–55. [ch 15]

Hearst, E., & Jenkins, H. M. (1974). *Sign-tracking: The stimulus-reinforcer relation and directed action*. Austin, TX: Psychonomic Society. [ch 17]

Hearst, E., Koresko, M. B., & Poppen, R. (1964). Stimulus generalization and the response-reinforcement contingency. *Journal of the Experimental Analysis of Behavior, 7*, 369–379. [ch 11]

Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley. [ch 3, 30]

Hebb, D. O. (1956). The distinction between "classical" and "instrumental." *Canadian Journal of Psychology, 10*, 165–166. [ch 17]

Hefferline, R. F. (1958). The role of proprioception in the control of behavior. *Transactions of the New York Academy of Sciences, 20*, 739–764. [ch 19]

Hein, A., Vital-Durand, F., Salinger, W., & Diamond, R. (1979). Eye movements initiate visual-motor development in the cat. *Science, 22*, 1321–1322. [ch 6]

Hein, H. (1972). The endurance of the mechanism-vitalism controversy. *Journal of the History of Biology, 5*, 159–188. [ch 30]

Held, R., & Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal* *of Comparative and Physiological Psychology, 56*, 872–876. [ch 6]

Hemmes, N. S. (1973). Behavioral contrast in pigeons depends upon the operant. *Journal of Comparative and Physiological Psychology, 85*, 171–178. [ch 16]

Henig, R. M. (2000). *The monk in the garden*. NY: Houghton Mifflin. [ch 3]

Henton, W. W., & Iversen, I. H. (1978). *Classical conditioning and operant conditioning*. New York: Springer-Verlag. [ch 16]

Herbranson, W. T., & Schroeder, J. (2010). Are birds smarter than mathematicians? Pigeons (Columba livia) perform optimally on a version of the Monty Hall dilemma. *Journal of Comparative Psychology, 124*, 1–13. [ch 29]

Herman, L. M., & Forestell, P. H. (1985). Reporting presence or absence of named objects by a language-trained dolphin. *Neuroscience and Biobehavioral Reviews, 9*, 667–681. [ch 24]

Herrick, R. M., Myers, J. L., & Korotkin, A. L. (1959). Changes in S^D and in S^Δ rates during the development of an operant discrimination. *Journal of Comparative and Physiological Psychology, 52*, 359–364. [ch 11]

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior, 4*, 267–272. [ch 16]

Herrnstein, R. J. (1964a). Aperiodicity as a factor in choice. *Journal of the Experimental Analysis of Behavior, 7*, 179–182. [ch 16]

Herrnstein, R. J. (1964b). Secondary reinforcement and rate of primary reinforcement. *Journal of the Experimental Analysis of Behavior, 7*, 27–36. [ch 14, 16]

Herrnstein, R. J. (1966). Superstition: A corollary of the principles of operant conditioning. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 33–51). New York: Appleton-Century-Crofts. [ch 5, 15]

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13*, 243–266. [ch 16]

Herrnstein, R. J., & Hineline, P. N. (1966). Negative reinforcement as shock-frequency reduction. *Journal of the Experimental Analysis of Behavior, 9*, 421–430. [ch 8]

Herrnstein, R. J., Loveland, D. H., & Cable, C. (1976). Natural concepts in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 2*, 285–311. [ch 11]

Herrnstein, R. J., & Morse, W. H. (1957). Some effects of response-independent positive reinforcement on maintained operant behavior. *Journal of Comparative and Physiological Psychology, 51*, 461–467. [ch 18]

Herrnstein, R. J., & Sidman, M. (1958). Avoidance conditioning as a factor in the effects of unavoidable shocks on food-reinforced behavior. *Journal of Comparative and Physiological Psychology, 51*, 380–385. [ch 18]

Herz, R. S., & Engen, T. (1996). Odor memory: Review and analysis. *Psychonomic Bulletin and Review, 3*, 300–313. [ch 27]

Hilgard, E. R. (1951). Method and procedures in the study of learning. In S. S. Stevens (Ed.), *Handbook of experimental psychology* (pp. 517–567). New York: Wiley. [ch 5]

Hilgard, E. R. (1986). *Divided consciousness: Multiple controls in human thought and action* (Expanded ed.). New York: Wiley. [ch 23]

Hilgard, E. R., & Marquis, D. G. (1940). *Conditioning and learning*. New York: Appleton-Century-Crofts. [ch 30]

Hineline, P. N. (1970). Negative reinforcement without shock reduction. *Journal of the Experimental Analysis of Behavior, 14*, 259–268. [ch 8, 16]

Hineline, P. N. (1977). Negative reinforcement and avoidance. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 364–414). Englewood Cliffs, NJ: Prentice-Hall. [ch 8]

Hineline, P. N. (1981). The several roles of stimuli in negative reinforcement. In P. Harzem & M. D. Zeiler (Eds.), *Predictability, correlation, and contiguity* (pp. 203–246). New York: Wiley. [ch 8]

Hineline, P. N. (1990). The origins of environment-based psychological theory. *Journal of the Experimental Analysis of Behavior, 53*, 305–320. [glossA, ch 29]

Hineline, P. N. (2001). Beyond the molar-molecular distinction: We need multiscaled analyses. *Journal of the Experimental Analysis of Behavior, 75*, 342–347. [ch 8, 16]

Hinson, J. M., & Staddon, J. E. R. (1981). Maximizing on interval schedules. In C. M. Bradshaw, E. Szabadi & C. F. Lowe (Eds.), *Quantification of steady-state operant behavior* (pp. 35–47). Amsterdam: Elsevier/North-Holland. [ch 16]

Hirsh, M. (2010). *The liberators: America's witnesses to the Holocaust*. New York: Bantam Books. [ch 3, 24]

Hodges, A. (1985). *Alan Turing: The enigma of intelligence*. London: Unwin. [ch 25]

Hoffman, H. S. (1996). *Amorous turkeys and addicted ducklings: A search for the causes of social attachment*. Boston: Authors Cooperative. [ch 4]

Hoffman, H. S., & Fleshler, M. (1959). Aversive control with the pigeon. *Journal of the Experimental Analysis of Behavior, 2*, 213–218. [ch 8]

Hoffman, H. S., & Fleshler, M. (1962). The course of emotionality in the development of avoidance. *Journal of Experimental Psychology, 64*, 288–294. [ch 18]

Hoffman, M. L. (1975). Developmental synthesis of affect and cognition and its implications for altruistic motivation. *Developmental Psychology, 11*, 607–622. [ch 19]

Holland, J. G. (1958). Human vigilance. *Science, 128*, 61–67. [ch 15]

Holland, P. C., & Ross, R. T. (1981). Within compound associations in serial compound conditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 7*, 228–241. [ch 17]

Hollard, V., & Davison, M. C. (1971). Preference for qualitatively different reinforcers. *Journal of the Experimental Analysis of Behavior, 16*, 375–380. [ch 16]

Hollard, V. D., & Delius, J. D. (1982). Rotational invariance in visual pattern recognition by pigeons and humans. *Science, 218*, 804–806. [ch 28]

Holt, G. L., & Shafer, J. N. (1973). Function of intertrial interval in matching-to-sample. *Journal of the Experimental Analysis of Behavior, 19*, 181–186. [ch 12]

Holz, W. C., & Azrin, N. H. (1961). Discriminative properties of punishment. *Journal of the Experimental Analysis of Behavior, 4*, 225–232. [ch 7]

Honig, W. K., Boneau, C. A., Burstein, K. R., & Pennypacker, H. S. (1963). Positive and negative generalization gradients obtained after equivalent training conditions. *Journal of Comparative and Physiological Psychology, 56*, 111–116. [ch 11]

Hooper, J. (2002). *Of moths and men: An evolutionary tale*. New York: Norton. [ch 10]

Horne, P. J., & Erjavec, M. (2007). Do infants show generalized imitation of gestures? *Journal of the Experimental Analysis of Behavior, 87*, 63–87. [ch 19, 22]

Horne, P. J., Hughes, J. C., & Lowe, C. F. (2006). Naming and categorization in young children: IV. Listener behavior training and transfer of function. *Journal of the Experimental Analysis of Behavior, 85*, 247–373. [ch 22]

Horne, P. J., & Lowe, C. F. (1996). On the origins of naming and other symbolic behavior. *Journal of the Experimental Analysis of Behavior, 65*, 185–241. [ch 12, 13, 22]

Horne, P. J., Lowe, C. F., & Harris, F. D. A. (2007). Naming and categorization in young children: V. Manual sign training. *Journal of the Experimental Analysis of Behavior, 87*, 367–381. [ch 22]

Horne, P. J., Lowe, C. F., & Randle, V. R. L. (2004). Naming and categorization in young children: II. Listener behavior training. *Journal of the Experimental Analysis of Behavior, 81*, 267–288. [ch 22]

Hull, C. L. (1920). Quantitative aspects of the evolution of concepts: An experimental study. *Psychological Monographs, 28, No. 123*. [ch 22]

Hull, C. L. (1934). Learning: II. The factor of the conditioned reflex. In C. Murchison (Ed.), *Handbook of general experimental psychology*. (pp. 382–455). Worchester, MA: Clark University Press. [ch 17]

Hull, C. L. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts. [ch 1, 8, 17, 30]

Hulse, S. H., Cynx, J., & Humpal, J. (1984). Cognitive processing of pitch and rhythm structures by birds. In H. L. Roitblat, T. G. Bever & H. S. Terrace (Eds.), *Animal cognition* (pp. 183–198). Hillsdale, NJ: Erlbaum. [ch 24]

Hursh, S. (1980). Economic concepts for the analysis of behavior. *Journal of the Experimental Analysis of Behavior, 34*, 219–238. [ch 15, 16]

Hursh, S. (1991). Behavioral economics of drug self-administration. *Journal of the Experimental Analysis of Behavior, 56*, 377–393. [ch 16]

Huttenlocher, J. (1968). Constructing spatial images: A strategy in reasoning. *Psychological Review, 75*, 550–560. [ch 28]

Huxley, R. (1970). The development of the correct use of subject personal pronouns in two children. In G. B. F. d'Arcais & W. J. M. Levelt (Eds.), *Advances in psycholinguistics* (pp. 141–165). Amsterdam: Elsevier/North-Holland. [ch 25]

Hübner, M. M. C., Austin, J. L., & Miguel, C. F. (2008). Effects of praising qualifying autoclitics on the frequency of reading. *Analysis of Verbal Behavior, 24*, 55–62. [ch 23]

Intraub, H. (1979). The role of implicit naming in pictorial encoding. *Journal of Experimental Psychology: Human Learning and Memory, 5*, 78–87. [ch 26]

Irvine, K. D., & Rauskolb, C. (2001). Boundaries in development: Formation and function. *Annual Review of Cell Developmental Biology, 17*, 189–214. [ch 3, 20, 30]

Ison, J. R., & Hoffman, H. S. (1983). Reflex modification in the domain of startle: II. The anomalous history of a robust and ubiquitous phenomenon. *Psychological Bulletin, 94*, 3–17. [ch 4, 17]

Iwata, B. A., Pace, G. M., Cowdery, G. E., & Miltenberger, R. G. (1994). What makes extinction work: An analysis of procedural form and function. *Journal of Applied Behavior Analysis, 27*, 131–144. [ch 9]

Iwata, B. A., Pace, G. M., Kalsher, M. J., Cowdery, G. E., & Cataldo, M. F. (1990). Experimental analysis and extinction of self-injurious escape behavior. *Journal of Applied Behavior Analysis, 23*, 11–27. [ch 9]

Jackendoff, R. (2002). *Foundations of Language*. Oxford: Oxford University Press. [ch 25]

Jacobs, J. (1887). Experiments on "prehension." *Mind, 12*, 75–79. [ch 27]

James, W. (1890). *The principles of psychology. Volume 1*. New York: Holt. [ch 12]

Jans, J. E., & Catania, A. C. (1980). Short-term remembering of discriminative stimuli in the pigeon. *Journal of the Experimental Analysis of Behavior, 34*, 177–183. [ch 26]

Jaynes, J. (1976). *The origin of consciousness in the breakdown of the bicameral mind*. Boston, MA: Houghton Mifflin. [ch 19, 22, 23, 24, 25]

Jenkins, H. M., & Harrison, R. H. (1960). Effect of discrimination training on auditory generalization. *Journal of Experimental Psychology, 59*, 246–253. [ch 11]

Jenkins, H. M., & Moore, B. R. (1973). The form of the auto-shaped response with food or water reinforcers. *Journal of the Experimental Analysis of Behavior, 20,* 163–181. [ch 17]

Jenkins, H. M., & Sainsbury, R. S. (1970). Discrimination learning with the distinctive feature on positive or negative trials. In D. I. Mostofsky (Ed.), *Attention: Contemporary theory and analysis* (pp. 239–273). New York: Appleton-Century-Crofts.

Jenkins, J. G., & Dallenbach, K. M. (1924). Obliviscence during sleep and waking. *American Journal of Psychology, 35,* 605–612. [ch 27]

Jennings, H. S. (1906). *Behavior of the lower organisms.* New York: Macmillan. [ch 30]

Johanson, I. B., & Hall, W. G. (1979). Appetitive learning in 1-day-old rat pups. *Science, 205,* 419–421. [ch 3]

Johnson, D. F., & Cumming, W. W. (1968). Some determiners of attention. *Journal of the Experimental Analysis of Behavior, 11,* 157–166. [ch 11]

Johnson, K. R., & Layng, T. V. J. (1992). Breaking the structuralist barrier: Literacy and numeracy with fluency. *American Psychologist, 47,* 1475–1490. [ch 5, 6, 13, 30]

Johnson, M. (1987). *The body in the mind.* Chicago: University of Chicago Press. [ch 22]

Johnson, M. K., & Hasher, L. (1987). Human learning and memory. *Annual Review of Psychology, 38,* 631–668. [ch 27]

Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review, 88,* 67–85. [ch 26]

Johnson, P. (1982). The functional equivalence of imagery and movement. *Quarterly Journal of Experimental Psychology, 34A,* 349–365. [ch 28]

Johnston, J. M. (2000). Behavior analysis and the R&D paradigm. *The Behavior Analyst, 23,* 141–148. [ch 9]

Johnston, J. M., & Pennypacker, H. S. (1993). *Strategies and tactics of behavioral research* (2nd ed.). Hillsdale, NJ: Erlbaum. [ch 6]

Jones, E. A., & Carr, E. G. (2004). Joint attention in children with autism: Theory and intervention. *Focus on Autism and Other Developmental Disabilities, 19,* 13–26. [ch 19]

Jones, M. (2004, Sept 19). Who was abused. *New York Times Magazine.* [ch 26]

Jonçich, G. (1968). *The sane positivist. A biography of Edward L. Thorndike.* Middletown, CT: Wesleyan University Press. [ch 26, 29]

Jordanova, L. J. (1984). *Lamarck.* Oxford: Oxford University Press. [ch 3]

Jost, A. (1897). Die Associationsfestigkeit in ihrer Abhängigkeit von der Verteilung der Wiederholungen. *Zeitschrift für Psychologie, 14,* 436–472. [ch 27]

Kaas, J. H. (1991). Plasticity of sensory and motor maps in adult mammals. *Annual Review of Neuroscience, 14,* 137–167. [ch 30]

Kahneman, D. (2011). *Thinking, fast and slow.* New York: Farrar, Straus and Giroux. [ch 16, 29]

Kamil, A. C., Yoerg, S. I., & Clements, K. C. (1988). Rules to leave by: Patch departure in foraging blue jays. *Animal Behaviour, 36,* 843–853. [ch 16]

Kamin, L. J. (1956). The effects of termination of the CS and avoidance of the US on avoidance learning. *Journal of Comparative and Physiological Psychology, 49,* 420–424. [ch 8, 17]

Kamin, L. J. (1957). The retention of an incompletely learned avoidance response. *Journal of Comparative and Physiological Psychology, 50,* 457–460. [ch 27]

Kandel, E. R., & Schwartz, J. H. (1982). Molecular biology of learning: Modulation of transmitter release. *Science, 218,* 433–443. [ch 30]

Kastak, C. R., & Schusterman, R. J. (2002). Sea lions and equivalence: Expanding classes by exclusion. *Journal of the Experimental Analysis of Behavior, 78,* 449–465. [ch 24]

Kastak, C. R., Schusterman, R. J., & Kastak, D. (2001). Equivalence classification by California sea lions using class-specific reinforcers. *Journal of the Experimental Analysis of Behavior, 76,* 131–158. [ch 24]

Katz, I., & Catania, A. C. (2005). Concurrent performances: Extinction, noncontingent reinforcement and variably delayed reinforcement. *European Journal of Behavior Analysis, 6,* 95–108. [ch 5, 15]

Katz, J. S., & Wright, A. A. (2006). Same/different abstract-concept learning by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 32,* 80–86. [ch 12, 14]

Kawamura, S. (1959). The process of sub-culture propagation among Japanese macaques. *Primates, 2,* 43–60. [ch 19]

Kazdin, A. E. (1977). The influence of behavior preceding a reinforced response on behavior change in the classroom. *Journal of Applied Behavior Analysis, 10,* 299–310. [ch 5]

Kelleher, R. T., & Gollub, L. R. (1962). A review of positive conditioned reinforcement. *Journal of the Experimental Analysis of Behavior, 5,* 543–597. [ch 16]

Kelleher, R. T., Riddle, W. C., & Cook, L. (1962). Observing responses in pigeons. *Journal of the Experimental Analysis of Behavior, 5,* 3–13. [ch 16]

Keller, F. S. (1941). Light aversion in the white rat. *Psychological Record, 4,* 235–250. [ch 8]

Keller, F. S. (1958). The phantom plateau. *Journal of the Experimental Analysis of Behavior, 1,* 1–13. [ch 26, 30]

Keller, F. S. (1968). "Goodbye, teacher." *Journal of Applied Behavior Analysis, 1,* 79–89. [ch 13]

Keller, F. S., & Schoenfeld, W. N. (1950). *Principles of psychology.* New York: Appleton-Century-Crofts. [preface, ch 5, 10, 11, gloss]

Keller, K. J. (1974). The role of elicited responding in behavioral contrast. *Journal of the Experimental Analysis of Behavior, 21,* 249–257. [ch 16, 17]

Kendall, S. B. (1965). Spontaneous recovery after extinction with periodic time-outs. *Psychonomic Science, 2,* 117–118. [ch 5]

Keppel, G., & Underwood, B. J. (1962). Proactive inhibition in short-term retention of single items. *Journal of Verbal Learning and Verbal Behavior, 1,* 153–156. [ch 27]

Kershaw, I. (1998). *Hitler 1889–1936: Hubris.* New York: Norton. [ch 24]

Kershaw, I. (2000). *Hitler 1936–1945: Nemesis.* New York: Norton. [ch 24]

Kiecker, C., & Lumsden, A. (2005). Compartments and their boundaries in vertebrate brain development. *Nature Reviews Neuroscience, 6,* 553–564. [ch 3, 30]

Kiecker, C., & Lumsden, A. (2009). Recent advances in neural development. *Biology Reports, 1,* online journal. [ch 20, 30]

Kiesler, C. A., Nisbett, R. E., & Zanna, M. P. (1969). On inferring one's beliefs from one's behavior. *Journal of Personality and Social Psychology, 11,* 321–327. [ch 22]

Killeen, P. R. (1972). A yoked-chamber comparison of concurrent and multiple schedules. *Journal of the Experimental Analysis of Behavior, 18,* 13–22. [ch 16]

Killeen, P. R. (1992). Mechanics of the animate. *Journal of the Experimental Analysis of Behavior, 57,* 429–463. [ch 6]

Killeen, P. R. (1994). Mathematical principles of reinforcement. *Behavioral and Brain Sciences, 17,* 105–172. [glossA, ch 6, 15, 16]

Killeen, P. R. (2011). Markov model of smoking cessation. *Proceedings of the National Academy of Sciences, 108,* 15549–15556. [ch 6, 14]

Killeen, P. R., & Amsel, A. (1987). The kinematics of locomotion toward a goal. *Journal of Experimental Psychology: Animal Behavior Processes, 13,* 92–101. [ch 5]

Killeen, P. R., & Bizo, L. A. (1998). The mechanics of reinforcement. *Psychonomic Bulletin and Review, 5,* 221–238. [ch 6]

Killeen, P. R., & Hall, S. S. (2001). The principle components of response strength. *Journal of the Experimental Analysis of Behavior, 75,* 111–134. [ch 6]

Killeen, P. R., Hall, S. S., Reilly, M. P., & Kettle, L. C. (2002). Molecular analyses of the principal components of response strength. *Journal of the Experimental Analysis of Behavior, 78,* 127–160. [ch 6]

Killeen, P. R., & Sitomer, M. T. (2003). MPR. *Behavioural Processes, 62,* 49–64. [ch 6]

Kim, H., & Bao, S. (2008). Distributed representation of perceptual categories in the auditory cortex. *Journal of Computational Neuroscience, 24,* 277–290. [ch 20]

Kimble, G. A. (1947). Conditioning as a function of the time between conditioned and unconditioned stimuli. *Journal of Experimental Psychology, 37,* 1–15. [ch 17]

Kimble, G. A. (1961). *Hilgard & Marquis' Conditioning and learning* (2nd ed.). New York: Appleton-Century-Crofts. [ch 1, 5, 30]

Kimmel, H. D. (1976). Notes from "Pavlov's Wednesdays": Pavlov's law of effect. *American Journal of Psychology, 89,* 553–556. [ch 17]

Kinchla, R. A., & Wolfe, J. M. (1979). The order of visual processing: "Top-down," "bottom-up," or "middle-out." *Perception and Psychophysics, 25,* 225–231. [ch 28]

King, A. P., & West, M. J. (1985). Social metrics of song learning. *Learning and Motivation, 15,* 441–458. [ch 9]

Kirschner, M. W., & Gerhart, J. C. (2005). *The plausibility of life: Resolving Darwin's dilemma.* New Haven: Yale University Press. [ch 3, 9, 10, 13, 23, 29]

Kish, G. B. (1966). Studies of sensory reinforcement. In W. K. Honig (Ed.), *Operant Behavior: Areas of research and application* (pp. 109–159). New York: Appleton-Century-Crofts. [ch 6]

Kluender, K. R., Diehl, R. L., & Killeen, P. R. (1987). Japanese quail can learn phonetic categories. *Science, 237,* 1195–1197. [ch 24]

Knapp, T. J. (1992). *Verbal Behavior:* The other reviews. *Analysis of Verbal Behavior, 10,* 87–95. [Ch 20, 25]

Kohn, A. (1993). *Punished by rewards.* Boston: Houghton Mifflin.

Kolers, P. A. (1966). Reading and talking bilingually. *American Journal of Psychology, 79,* 357–376. [ch 20, 21, 27]

Kolers, P. A. (1985). Skill in reading and memory. *Canadian Journal of Psychology, 39,* 232–239. [ch 20]

Kolers, P. A., & Palef, R. (1976). Knowing not. *Memory and Cognition, 4,* 553–558. [ch 26]

Kolers, P. A., & Roediger, H. L., III. (1984). Procedures of mind. *Journal of Verbal Learning and Verbal Behavior, 23*(425–449). [ch 28]

Kolers, P. A., & Smythe, W. E. (1979). Images, symbols, and skills. *Canadian Journal of Psychology, 33,* 158–184. [ch 28]

Konorski, J. (1948). *Conditioned reflexes and neuron organization.* New York: Cambridge University Press. [ch 17]

Koriat, A. (1995). Dissociating knowing and the feeling of knowing: Further evidence for the accessibility model. *Journal of Experimental Psychology: General, 124,* 311–333. [ch 19]

Korzybski, A. (1941). *Science and sanity* (2nd ed.). Lancaster, PA: International Non-Aristotelian Library. [ch 23]

Krechevsky, I. (1932). Hypotheses in rats. *Psychological Review, 39,* 516–532. [ch 10]

Kroodsma, D. E., & Miller, E. H. (Eds.). (1982). *Acoustic communication in birds. Volume 2: Song learning and its consequences.* New York: Academic Press. [ch 19]

Kuczaj, S. A., III. (1977). The acquisition of regular and irregular past tense forms. *Journal of Verbal Learning and Verbal Behavior, 16,* 589–600. [ch 13]

Köhler, W. (1927). *The mentality of apes (tr. E. Winter; 2nd revised edition).* London: Routledge & Kegan Paul. [ch 2, 29]

Lachman, J. L., Lachman, R., & Thronesbery, C. (1979). Metamemory through the adult life span. *Developmental Psychology, 15,* 543–551. [ch 26]

Lakoff, G. (1987). *Women, fire, and dangerous things.* Chicago: University of Chicago. [ch 12, 22]

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by.* Chicago: University of Chicago Press. [ch 22]

Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh.* New York: Basic Books. [ch 22]

Lalli, J. S., Mace, F. C., Wohn, T., & Livezey, K. (1995). Identification and modification of a response-class hierarchy. *Journal of Applied Behavior Analysis, 28*, 551–559. [ch 9, 14]

Lalli, J. S., Casey, S. D., & Kates, K. (1997). Noncontingent reinforcement as treatment for severe problem behavior: Some procedural variations. *Journal of Applied Behavior Analysis, 30*, 127–136. [ch 5]

Lamb, M. R., & Riley, D. A. (1981). Effects of element arrangement on the processing of compound stimuli in pigeons (Columba livia). *Journal of Experimental Psychology: Animal Behavior Processes, 7*, 45–58. [ch 12]

Lane, H. (1961). Operant control of vocalizing in the chicken. *Journal of the Experimental Analysis of Behavior, 4*, 171–177. [ch 20)]

Lane, H. (1965). The motor theory of speech perception: A critical review. *Psychological Review, 72*, 275–309. [ch 20]

Lashley, K. S. (1930). The mechanism of vision. I. A method for rapid analysis of pattern vision in the rat. *Journal of Genetic Psychology, 37*, 453–460. [ch 2, 30]

Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 112–146). New York: Wiley. [ch 10]

Laties, V. G., Weiss, B., Clark, R. L., & Reynolds, M. D. (1965). Overt "mediating" behavior during temporally spaced responding. *Journal of the Experimental Analysis of Behavior, 8*, 107–116. [ch 10]

Lattal, K. A. (1974). Combinations of response reinforcer dependence and independence. *Journal of the Experimental Analysis of Behavior, 22*, 357–362. [ch 15]

Lattal, K. A., & Perone, M. (Eds.). (1998). *Handbook of research methods in human operant behavior.* NY: Plenum. [ch 30]

Lea, S. E. G. (1979). Foraging and reinforcement schedules in the pigeon: Optimal and non-optimal aspects of choice. *Animal Behavior, 27*, 875–886. [ch 11]

Lea, S. E. G., & Harrison, S. N. (1978). Discrimination of polymorphous stimulus sets by pigeons. *Quarterly Journal of Experimental Psychology, 30*, 521–537. [ch 11]

Leitenberg, H. (1966). Conditioned acceleration and conditioned suppression in pigeons. *Journal of the Experimental Analysis of Behavior, 9*, 205–212. [ch 18]

Lepper, M. R., & Greene, D. (Eds.). (1978). *The hidden costs of reward.* Hillsdale, NJ: Erlbaum. [ch 14]

Lepper, M. R., Greene, D., & Nisbett, R. E. (1973). Undermining children's intrinsic interest with extrinsic reward: A test of the "overjustification" hypothesis. *Journal of Personality and Social Psychology, 28*, 129–137. [ch 14]

Lerman, D. C., Iwata, B. A., & Wallace, M. D. (1999). Side effects of extinction: Prevalence of bursting and aggression during the treatment of self-injurious behavior. *Journal of Applied Behavior Analysis, 32*, 1–8. [ch 5]

Lerman, D. C., & Toole, L. M. (2011). Developing function-based punishment procedures for problem behavior. In W. W. Fisher, C. C. Piazza & H. S. Roane (Eds.), *Handbook of applied behavior analysis* (pp. 348–369). New York: Guikford. [ch 7]

Levitsky, D., & Collier, G. (1968). Schedule-induced wheel running. *Physiology and Behavior, 3*, 571–573. [ch 4]

Ley, R. (1990). *A whisper of espionage.* Garden City Park, NY: Avery. [ch 2]

Liberman, A. M. (1982). On finding that speech is special. *American Psychologist, 37*, 148–167. [ch 20]

Lindsley, O. R. (1956). Operant conditioning methods applied to research in chronic schizophrenia. *Psychiatric Research Reports, 6*, 118–139. [ch 9]

Lindsley, O. R. (1992). Precision teaching: Discoveries and effects. *Journal of Applied Behavior Analysis, 25*, 51–57. [ch 6, 30]

Loeb, J. (1900). *Comparative physiology of the brain and comparative psychology.* New York: Putnam. [ch 30]

Loeb, J. (1918/1973). *Forced movements, tropisms, and animal conduct.* New York: Dover. [ch 4]

Loftus, E. F. (1993). The reality of repressed memories. *American Psychologist, 48*, 518–537. [ch 26]

Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior, 13*, 585–589. [ch 26]

Loftus, E. F., & Zanni, G. (1975). Eyewitness testimony: The influence of the wording of a question. *Bulletin of the Psychonomic Society, 5*, 86–88. [ch 26]

Logan, F. A. (1960). *Incentive.* New Haven, CT: Yale University Press. [ch 6]

Loisette, A. (1899). *Assimilative memory, or, How to attend and never forget.* New York: Funk & Wagnalls. [ch 26]

Long, G. M. (1980). Iconic memory: A review and critique of the study of short-term visual storage. *Psychological Bulletin, 88*, 785–820. [ch 27]

Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences, 20*, 130–141. [ch 3]

Lorenz, K. (1937). The companion in the bird's world. *Auk, 54*, 245–273. [ch 4]

Lovaas, O. I. (1964). Cue properties of words: The control of operant responding by rate and content of verbal operants. *Child Development, 35*, 245–256. [ch 23]

Lovaas, O. I., & Simmons, J. Q. (1969). Manipulation of self-destruction in three retarded children. *Journal of Applied Behavior Analysis, 2*, 143–157. [ch 9]

Lowe, C. F. (1980). Determinants of human operant behavior. In P. Harzem & M. D. Zeiler (Eds.), *Advances in the analysis of behavior. Volume 1. Reinforcement and the organization of behavior* (pp. 159–192). Chichester, England: Wiley. [ch 23]

Lowe, C. F., & Chadwick, P. D. J. (1990). Verbal control of delusions. *Behavior Therapy, 21*, 461–479. [ch 23]

Lowe, C. F., Horne, P. J., Harris, F. D. A., & Randle, V. R. L. (2002). Naming and categorization in young children: Vocal tract training. *Journal of the Experimental Analysis of Behavior, 78*, 527–549. [ch 22]

Lowe, C. F., Horne, P. J., & Hughes, J. C. (2005). Naming and categorization in young children: III. Vocal tract training and transfer of function. *Journal of the Experimental Analysis of Behavior, 83*, 47–65. [ch 22]

Lowenkron, B. (1997). The role of naming in the development of joint control. *Journal of the Experimental Analysis of Behavior, 68*, 244–247. [ch 13]

Lowenkron, B. (1998). Some logical functions of joint control. *Journal of the Experimental Analysis of Behavior, 69*, 327–354. [gloss], ch 13]

Lowenkron, B. (2004). Meaning: A verbal behavior account. *Analysis of Verbal Behavior, 20*, 77–97. [ch 13, 22]

Lubbock, J. (1882). *Ants, bees, and wasps.* New York: D. Appleton. [ch 30]

Lubinski, D., & Thompson, T. (1987). An animal model of the interpersonal communication of interoceptive (private) states. *Journal of the Experimental Analysis of Behavior, 48*, 1–15. [ch 19]

Luchins, A. S., & Luchins, E. H. (1950). New experimental attempts at preventing mechanization in problem solving. *Journal of General Psychology, 42*, 279–297. [ch 29]

MacCorquodale, K. (1970). On Chomsky's review of Skinner's "Verbal Behavior". *Journal of the Experimental Analysis of Behavior, 13*, 83–99. [ch 25]

MacLeod, C. M. (1976). Bilingual episodic memory: Acquisition and forgetting. *Journal of Verbal Learning and Verbal Behavior, 15*, 347–364. [ch 27]

MacPhail, E. M. (1968). Avoidance responding in pigeons. *Journal of the Experimental Analysis of Behavior, 11*, 629–632. [ch 8]

Madden, G. J. (Ed.). (2012). *APA Handbook of Behavior Analysis.* Washington, DC: American Psychological Association. [ch 30]

Mahoney, M. J., & Bandura, A. (1972). Self-reinforcement in pigeons. *Learning and Motivation, 3*, 293–303. [ch 19]

Maier, S. F., Albin, R. W., & Testa, T. J. (1973). Failure to learn to escape in rats previously exposed to inescapable shock depends on nature of escape response. *Journal of Comparative and Physiological Psychology, 85*, 581–592. [ch 12]

Maier, S. F., Seligman, M. E. P., & Solomon, R. L. (1969). Pavlovian fear conditioning and learned helplessness: Effects on escape and avoidance behavior of (a) the CS-US contingency and (b) the independence of the US and voluntary responding. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 299–342). New York: Appleton-Century-Crofts. [ch 12]

Maki, R. H. (1979). Right-left and up-down are equally discriminable in the absence of directional words. *Bulletin of the Psychonomic Society, 14*, 181–184. [ch 13, 23]

Malcolm, N. (1971). The myth of cognitive processes and structure. In T. Mischel (Ed.), *Cognitive development and epistemology* (pp. 385–392). New York: Academic Press. [ch 28]

Malone, J. C. (1990). *Theories of learning.* Belmont, CA: Wadsworth. [ch 9, 16]

Malott, R. W., & Cumming, W. W. (1964). Schedules of interresponse time reinforcement. *Psychological Record, 14*, 221–252. [ch 10, 15]

Manabe, K., Kawashima, T., & Staddon, J. E. R. (1995). Differential vocalization in budgerigars: Towards an experimental analysis of naming. *Journal of the Experimental Analysis of Behavior, 63*, 111–126. [ch 24]

Marholin, D., II, & Steinman, W. M. (1977). Stimulus control in the classroom as a function of the behavior reinforced. *Journal of Applied Behavior Analysis, 10*, 465–478. [ch 26]

Markman, E. M., Horton, M. S., & McLanahan, A. G. (1980). Classes and collections: Principles of organization in the learning of hierarchical relations. *Cognition, 8*, 227–241. [ch 28]

Markoff, J. (2011). Computer wins on 'Jeopardy!' Trivial, it's not. *New York Times.* [ch 29]

Marler, P., & Peters, S. (1982). Long-term storage of learned birdsongs prior to production. *Animal Behaviour, 30*, 479–482. [ch 24]

Marochnik, L. S., & Suchkov, A. A. (1995). *The Milky Way Galaxy* (2nd revised ed.). Luxembourg: Gordon and Breach. [ch 3]

Matthews, B. A., Catania, A. C., & Shimoff, E. (1985). Effects of uninstructed verbal behavior on nonverbal responding: Contingency descriptions versus performance descriptions. *Journal of the Experimental Analysis of Behavior, 43*, 155–164. [ch 23]

Matthews, B. A., Shimoff, E., & Catania, A. C. (1987a). Saying and doing: A contingency-space analysis. *Journal of Applied Behavior Analysis, 20*, 69–74. [ch 23]

Matthews, B. A., Shimoff, E., & Catania, A. C. (1987b). *Sociology on a disk: Computer programs in introductory sociology.* Ellicott City, MD: CMS Software. [ch 29]

Matthews, B. A., Shimoff, E., Catania, A. C., & Sagvolden, T. (1977). Uninstructed human responding: Sensitivity to ratio and interval contingencies. *Journal of the Experimental Analysis of Behavior, 27*, 453–467. [ch 23]

Max, L. W. (1934). An experimental study of the motor theory of consciousness. I. History and critique. *Journal of General Psychology, 11*, 112–125. [ch 20]

Mayer, G. R., Sulzer-Azaroff, B., & Wallace, M. (2012). *Behavior analysis for lasting change* (2nd ed.). Cornwall-on-Hudson, NY: Sloan. [ch 30]

Mayr, E. (1982). *The growth of biological thought.* Cambridge, MA: Belknap. [ch 3]

Mazur, J. E. (1986). *Learning and behavior.* Englewood Cliffs, NJ: Prentice-Hall. [ch 18]

Mazur, J. E. (1991). Choice. In I. H. Iversen & K. A. Lattal (Eds.), *Experimental analysis of behavior. Part 1* (pp. 219–250). Amsterdam: Elsevier/North-Holland. [ch 16]

Mazur, J. E. (1996). Procrastination by pigeons: Preference for larger, more delayed work requirements. *Journal of the Experimental Analysis of Behavior, 65*, 159–171. [ch 16]

McCloskey, M., & Santee, J. (1981). Are semantic memory and episodic memory distinct systems? *Journal of Experimental Psychology: Human Learning and Memory, 7*, 66–71. [ch 27]

McGeoch, J. A. (1942). *The psychology of human learning.* New York: Longmans, Green. [ch 26, 29, 30]

McGeoch, J. A., & Irion, A. L. (1952). *The psychology of human learning* (2nd ed.). New York: Longmans, Green. [ch 30]

McGlinchey, A., & Keenan, M. (1997). Stimulus equivalence and social categorization in Northern Ireland. *Behavior and Social Issues, 7*, 113–128. [ch 13]

McGraw, M. B. (1945). *The neuromuscular maturation of the human infant.* New York: Columbia University Press. [ch 4]

Meehl, P. E. (1950). On the circularity of the Law of Effect. *Psychological Bulletin, 47*, 52–75. [ch 5]

Mehler, J., Jusczyk, P., Lambert, G., Halsted, N., Bertoncini, J., & Amiel-Tison, C. (1988). A precursor of language acquisition in young infants. *Cognition, 29*, 143–178. [ch 20]

Meisch, R. A., & Thompson, T. (1971). Ethanol intake in the absence of concurrent food reinforcement. *Psychopharmacologia, 22*, 72–79. [ch 4]

Melton, A. W. (1963). Implications of short-term memory for a general theory of memory. *Journal of Verbal Learning and Verbal Behavior, 2*, 1–21. [ch 27]

Melton, A. W., & Martin, E. (Eds.). (1972). *Coding processes in human memory.* Washington, DC: Winston. [ch 26]

Meyer, K. E., & Brysac, S. B. (2012). *Pax ethnica: Where and how diversity succeeds.* Philadelphia: PublicAffairs. [ch 11, 30]

Michael, J. (1975). Positive and negative reinforcement, a distinction that is no longer necessary; or a better way to talk about bad things. *Behaviorism, 3*, 33–44. [ch 8]

Michael, J. (1982). Distinguishing between discriminative and motivational functions of stimuli. *Journal of the Experimental Analysis of Behavior, 37*, 149–155. [ch 2, 6]

Michael, J. (1989). Establishing operations. In *Concepts and principles of behavior analysis* (pp. 57–72). Kalamazoo, MI: Association for Behavior Analysis.

Milgram, S. (1963). Behavioral study of obedience. *Journal of Abnormal and Social Psychology, 67*, 371–378. [ch 23]

Miller, G. A. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63*, 81–97. [ch 26, 27]

Miller, G. A., & Selfridge, J. A. (1950). Verbal context and the recall of meaningful materia. *American Journal of Psychology, 63*, 176–185. [ch 21]

Miller, J. D., & Bowe, C. A. (1982). Roles of the qualities and locations of stimuli and responses in simple associative learning. *Pavlovian Journal of Biological Science, 17*, 129–139. [ch 18]

Miller, N. E., & Carmona, A. (1967). Modification of a visceral response, salivation in thirsty dogs, by instrumental training with water reward. *Journal of Comparative and Physiological Psychology,* 63, 1–6. [ch 9, 10, 17]

Mineka, S., Davidson, M., Cook, M., & Keir, R. (1984). Observational learning of snake fear in Rhesus monkeys. *Journal of Abnormal Psychology, 93*, 355–372. [ch 19]

Mineka, S., & Zinbarg, R. (2006). A contemporary learning theory perspective on the etiology of anxiety disorders: It's not what you thought it was. *American Psychologist, 61*, 10–26. [ch 17, 18, 19]

Mitchell, D., Scott, D. W., & Williams, K. D. (1973). Container neophobia and the rat's preference for earned food. *Behaviorial Biology, 9*, 613–624. [ch 18]

Miyake, N., & Norman, D. A. (1979). To ask a question, one must know enough to know what is not known. *Journal of Verbal Learning and Verbal Behavior, 18*, 357–364. [ch 26]

Miyashita, Y., Nakajima, S., & Imada, H. (2000). Differential outcome effect in the horse. *Journal of the Experimental Analysis of Behavior, 74*, 245–253. [ch 14]

Moerk, E. L. (1980). Relationships between parental input frequencies and children's language acquisition: A reanalysis of Brown's data. *Journal of Child Language, 7*, 1–14. [ch 25]

Moerk, E. L. (1983). A behavioral analysis of controversial topics in first language acquisition: Reinforcements, corrections, modeling, input frequencies, and the three-term contingency pattern. *Journal of Psycholinguistic Research, 12*, 129–155. [ch 25]

Moerk, E. L. (1992). *First language: Taught and learned.* Baltimore: Paul H. Brookes. [ch 25]

Montee, B. B., Miltenberger, R. G., & Wittrock, D. (1995). An experimental analysis of facilitated communication. *Journal of Applied Behavior Analysis, 28*, 189–200. [ch 2]

Moore, B. R., & Stuttard, S. (1979). Dr. Guthrie and Felis domesticus or: Tripping over the cat. *Science, 205*, 1031–1033. [ch 5]

Moore, C., & Frye, D. (1986). Context, conservation and the meanings of more. *British Journal of Developmental Psychology, 4*, 169–178. [ch 25]

Morgan, C. L. (1920). *Animal behaviour.* London: Edward Arnold. [ch 2]

Morgan, M. J., Fitch, M. D., Holman, J. G., & Lea, S. E. G. (1976). Pigeons learn the concept of an "A." *Perception and Psychophysics, 5*, 57–66. [ch 11]

Morse, W. H., & Kelleher, R. T. (1977). Determinants of reinforcement and punishment. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 174–200). Englewood Cliffs, NJ: Prentice-Hall. [ch 7, 16]

Morse, W. H., & Skinner, B. F. (1958). Some factors involved in the stimulus control of operant behavior. *Journal of the Experimental Analysis of Behavior, 1*, 103–107. [glossM, gloss S]

Motokawa, T. (1989). Sushi science and hamburger science. *Perspectives on Biology and Medicine, 32*, 489–504. [ch 11]

Mowrer, O. H. (1960). *Learning theory and behavior.* New York: Wiley. [ch 17]

Mowrer, O. H., & Jones, H. M. (1943). Extinction and behavior variability as a function of effortfulness of task. *Journal of Experimental Psychology, 33*, 369–385. [ch 5]

Moyer, R. S., & Dumais, S. T. (1978). Mental comparison. In G. H. Bower (Ed.), *The psychology of learning and motivation. Volume 12* (pp. 117–155). New York: Academic Press. [ch 28]

Müller, G. E., & Pilzecker, A. (1900). Experimentelle Beiträge zür Lehre vom Gedächtnis. *Zeitschrift für Psychologie, Ergänzungsband 1.* [ch 26, 29]

Neisser, U. H., N. (1992). Phantom flashbulbs: False recollections of hearing the news about Challenger. In E. Winograd & U. Neisser (Eds.), *Affect and accuracy in recall: Studies of "flashbulb memories"* (pp. 9–31). New York: Cambridge University Press. [ch 27]

Neiworth, J. J., & Rilling, M. E. (1987). A method for studying imagery in animals. *Journal of Experimental Psychology: Animal Behavior Processes, 13*, 203–214. [ch 12]

Nelson, A. (1951). Narapoia. *The Magazine of Fantasy and Science Fiction, 2*(2, April), 114–118. [ch 29]

Nelson, T. O. (1992). Metamemory. In L. R. Squire (Ed.), *The encyclopedia of learning and memory* (pp. 412–415). New York: Macmillan. [ch 26]

Neuman, P., Ahearn, W. H., & Hineline, P. N. (2000). Pigeons' choices between fixed-ratio and linear or geometric escalating schedules. *Journal of the Experimental Analysis of Behavior, 73*, 93–102. [ch 16]

Neuringer, A. (1986). Can people behave randomly? The role of feedback. *Journal of Experimental Psychology: General, 115*, 62–75. [ch 10, 13]

Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin and Review, 9*, 672–705. [ch 9, 10, 13]

Neuringer, A. (2004). Reinforced variability in animals and people. *American Psychologist, 59*, 891–906. [ch 9, 10, 13]

Neuringer, A., & Jensen, G. (2010). Operant variability and voluntary action. *Psychological Review, 117*, 972–993. [glossV, ch 22]

Neuringer, A., Jensen, G., & Piff, P. (2007). Stochastic matching and the voluntary nature of choice. *Journal of the Experimental Analysis of Behavior, 88*, 1–28. [ch 22]

Neuringer, A., & Neuringer, M. (1974). Learning by following a food source. *Science, 184*, 1005–1008. [ch 19]

Neuringer, A., & Schneider, B. A. (1968). Separating the effects of interreinforcement time and number of interreinforcement responses. *Journal of the Experimental Analysis of Behavior, 11*, 661–667. [ch 16]

Nevin, J. A. (1969). Signal detection theory and operant behavior: A review of David M. Green & John A. Swets' *Signal detection theory and psychophysics. Journal of the Experimental Analysis of Behavior, 12*, 475–480. [ch 9, 11, 23]

Nevin, J. A. (1992). An integrative model for the study of behavioral momentum. *Journal of the Experimental Analysis of Behavior, 57*, 301–316. [ch 5, 6, 10, 13]

Nevin, J. A., & Grace, R. C. (2000). Behavioral momentum and the Law of Effect. *Behavioral and Brain Sciences, 23*, 73–130. [glossS, ch 6, 15, 16]

Nevin, J. A., & Grace, R. C. (2005). Resistance to extinction in the steady state and in transition. *Journal of Experimental Psychology: Animal Behavior Processes, 31*, 199–212. [ch 5, 6]

Nevin, J. A., Grace, R. C., Holland, S., & McLean, A. P. (2001a). Variable-ratio versus variable-interval schedules: Response rate, resistance to change, and preference. *Journal of the Experimental Analysis of Behavior, 76*, 43–74. [ch 15, 16]

Nevin, J. A., McLean, A. P., & Grace, R. C. (2001b). Resistance to extinction: Contingency termination and generalization decrement. *Animal Learning and Behavior, 29*, 176–191. [ch 5]

Nevin, J. A., Milo, J., Odum, A. L., & Shahan, T. A. (2003). Accuracy of discrimination, rate of responding, and resistance to change. *Journal of the Experimental Analysis of Behavior, 79*, 307–321. [ch 11, 12]

Newell, A., Shaw, J. C., & Simon, H. A. (1958). Elements of a theory of human problem solving. *Psychological Review, 65*, 151–166. [ch 29]

Newman, S. E. (1972). In search of associative symmetry. In C. P. Duncan, L. Sechrest & A. W. Melton (Eds.), *Human memory* (pp. 133–153). New York: Appleton-Century-Crofts. [ch 21]

Nisbett, N. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84*, 231–259. [ch 19]

Olton, D. S. (1979). Mazes, maps, and memory. *American Psychologist, 34*, 583–596. [ch 11, 12]

Olton, D. S., & Samuelson, R. J. (1976). Remembrance of places passed: Spatial memory in rats. *Journal of Experimental Psychology: Animal Behavior Processes, 2*, 97–116. [ch 11]

Orne, M. T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist, 17*, 776–783. [ch 21]

Ortony, A., & Turner, J. (1990). What's basic about basic emotions? *Psychological Review, 97*, 315–331. [ch 18]

Osgood, C. E. (1949). The similarity paradox in human learning: A resolution. *Psychological Review, 56*, 132–143. [ch 26]

Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The measurement of meaning*. Urbana, IL: University of Illinois. [ch 25]

Overton, D. A. (1964). State-dependent or "dissociated" learning produced with phenobarbital. *Journal of Comparative and Physiological Psychology, 57*, 3–12. [ch 26]

Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes, 11*, 429–452. [ch 10, 13]

Paivio, A. (1971). *Imagery and verbal processes*. New York: Holt, Rinehart & Winston. [ch 26, 28]

Paivio, A. (1975). Neomentalism. *Canadian Journal of Psychology, 29*, 263–291. [ch 28]

Palmer, D. C. (1996). Achieving parity: The role of automatic reinforcement. *Journal of the Experimental Analysis of Behavior, 65*, 289–290. [glossP, ch 13, 25]

Palmer, D. C. (1998). The speaker as listener: The interpretation of structural regularities in verbal behavior. *Analysis of Verbal Behavior, 15*, 3–16. [ch 13, 25]

Paniagua, F. A., & Baer, D. M. (1982). The analysis of correspondence training as a chain reinforceable at any point. *Child Development, 53*, 786–798. [ch 23]

Parker, A. (2003). *In the blink of an eye: How vision kick-started the big bang of evolution*. London: Free Press. [ch 3]

Parsons, H. M. (1974). What happened at Hawthorne? *Science, 183*, 922–932. [ch 6]

Pashler, H., McDaniel, M., Rohrer, D., & Bjork, R. A. (2008). Learning styles: Concepts and evidence. *Psychological Science in the Public Interest, 9*, 105–119. [ch 13, 30]

Pavlov, I. P. (1927). *Conditioned reflexes* (G. V. Anrep, Trans.). London: Oxford University Press. [ch 2, 4, 17, 30]

Pavlov, I. P. (1957). *Experimental psychology and other essays*. New York: Philosophical Library. [ch 30]

Peele, D. B., & Ferster, C. B. (1982). Autoshaped key pecking maintained by access to a social space. *Journal of the Experimental Analysis of Behavior, 38*, 181–189. [ch 17]

Pennypacker, H., & Iwata, M. M. (1990). Mammacare: A case history in behavioral medicine. In D. E. Blackman & H. Lejeune (Eds.), *Behaviour Analysis in Theory and Practice* (pp. 259–288). Hove: Erlbaum. [ch 11]

Pepperberg, I. M. (1988). The importance of social interaction and observation in the acquisition of communicative competence: Possible parallels between avian and human learning. In T. R. Zentall & J. B. G. Galef (Eds.), *Social learning* (pp. 279–299). Hillsdale, NJ: Erlbaum. [ch 24]

Perone, M. (2003). Negative effects of positive reinforcement. *The Behavior Analyst, 26*, 1–14. [ch 8, 16]

Peterson, G. B. (2004). A day of great illumination: B. F. Skinner's discovery of shaping. *Journal of the Experimental Analysis of Behavior, 82*, 317–328. [ch 9]

Peterson, L. R., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology, 58*, 193–198. [ch 27]

Peterson, N. (1960). Control of behavior by presentation of an imprinted stimulus. *Science, 132*, 1395–1396. [ch 4]

Peterson, N. (1962). Effect of monochromatic rearing on the control of responding by wavelength. *Science, 136*, 774–775. [ch 4, 11]

Peterson, R. F. (1968). Some experiments on the organization of a class of imitative behaviors. *Journal of Applied Behavior Analysis, 1*, 225–235. [ch 19]

Petroski, H. (1992). *The evolution of useful things*. New York: Knopf. [ch 30]

Pfungst, O. (1911). *Clever Hans (the horse of Mr. Von Osten): A contribution to experimental animal and human psychology (tr. C. L. Rahn)*. New York: Holt. [ch 2, 24]

Piaget, J., & Inhelder, B. (1969). *The psychology of the child (tr. H. Weaver)*. New York: Basic Books. [ch 28]

Pilgrim, C. (2004). *Stimulus equivalence and class-specific reinforcement*. Paper presented at the American Psychological Association, Honolulu, Hawaii. [ch 14]

Pilgrim, C., Jackson, J., & Galizio, M. (2000). Acquisition of arbitrary conditional discriminations by young normally developing children. *Journal of the Experimental Analysis of Behavior, 73*, 177–193. [ch 14]

Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press. [ch 25]

Pinker, S. (1994). *The language instinct*. New York: Morrow. [ch 25]

Pinker, S., & Bloom, P. (1990). Natural language and natural selection. *Behavioral and Brain Sciences, 13*, 707–784. [ch 25]

Pitts, R. C., & Malagodi, E. F. (1991). Preference for less frequent shock under fixed-interval schedules of electric-shock presentation. *Journal of the Experimental Analysis of Behavior, 56*, 21–32. [ch 7]

Platt, J. R. (1973). Percentile reinforcement: Paradigms for experimental analysis of response shaping. In G. H. Bower (Ed.), *The psychology of learning and motivation. Volume 7* (pp. 271–296). New York: Academic Press. [ch 9]

Pliskoff, S. S., & Goldiamond, I. (1966). Some discriminative properties of fixed ratio performance in the pigeon. *Journal of the Experimental Analysis of Behavior, 9*, 1–9. [ch 19]

Poling, A., Weetjens, B., Cox, C., Beyene, N. W., Bach, H., & Sully, A. (2011). Using trained pouched rats to detect land mines: Another victory for operant conditioning. *Journal of Applied Behavior Analysis, 44*, 351–355. [ch 9]

Polson, D. A. D., Grabavac, D. M., & Parsons, J. A. (1997). Intraverbal stimulus-response reversibility: Fluency, familiarity effects, and implications for stimulus equivalence. *Analysis of Verbal Behavior, 14*, 19–40. [ch 13, 21]

Posner, M. I. (1978). *Chronometric explorations of mind*. Hillsdale, NJ: Erlbaum. [ch 28]

Posner, M. I. (1982). Cumulative development of attentional theory. *American Psychologist, 37*, 168–179. [ch 28]

Potts, L., Eshleman, J. W., & Cooper, J. O. (1993a). Ogden R. Lindsley and the historical development of precision teaching. *The Behavior Analyst, 16*, 177–189. [ch 13, 30]

Poulson, C. L. (1984). Operant theory and methodology in infant vocal conditioning. *Journal of Experimental Child Psychology, 38*, 103–113. [ch 5, 19, 20]

Poulson, C. L., & Kymissis, E. (1988). Generalized imitation in infants. *Journal of Experimental Child Psychology, 46*, 324–336. [ch 19]

Poulson, C. L., Kymissis, E., Reeve, K. F., Andreatos, M., & Reeve, L. (1991). Generalized vocal imitation in infants. *Journal of Experimental Child Psychology, 5*, 267–279. [ch 19, 20]

Powers, R. B., & Osborne, J. G. (1976). *Fundamentals of behavior*. St. Paul, MN: West Publishing. [ch 2]

Pratkanis, A. R. (1992). The cargo-cult science of subliminal persuasion. *Skeptical Inquirer, 16*, 260–272. [ch 26]

Premack, D. (1959). Toward empirical behavior laws: I. Positive reinforcement. *Psychological Review, 66*, 219–233. [ch 6]

Premack, D. (1962). Reversibility of the reinforcement relation. *Science, 136*, 255–257. [ch 6]

Premack, D. (1970). A functional analysis of language. *Journal of the Experimental Analysis of Behavior, 14*, 107–125. [ch 24]

Premack, D. (1971). Catching up with common sense or two sides of a generalization: Reinforcement and punishment. In R. Glaser (Ed.), *The nature of reinforcement* (pp. 121–150). New York: Academic Press. [ch 6, 7]

Provine, R. R. (1976). Development of function in nerve nets. In J. Fentress (Ed.), *Simpler networks and behavior* (pp. 203–220). Sunderland, MA: Sinauer. [ch 1, 4]

Provine, R. R. (1981). Development of wing-flapping and flight in normal and flap-deprived domestic chicks. *Developmental Psychobiology, 14*, 279–291. [ch 18]

Provine, R. R. (1984). Wing-flapping during development and evolution. *American Scientist, 72*, 448–455. [ch 3]

Provine, R. R. (1988). A hierarchy of developmental contingencies: A review of Purves and Lichtman's Principles of Neural Development. *Journal of the Experimental Analysis of Behavior, 50*, 565–569. [glossS, ch 4]

Provine, R. R. (1989). Faces as releasers of contagious yawning: An approach to face detection using normal human subjects. *Bulletin of the Psychonomic Society, 27*, 211–214. [ch 18, 19]

Provine, R. R. (2000). *Laughter: A scientific investigation*. New York: Viking. [ch 4, 6, 7]

Provine, R. R. (2004). Laughing, tickling, and the evolution of speech and self. *Current Directions in Psychological Science, 13*, 215–218. [ch 19]

Provine, R. R. (2012). *Curious behavior: Yawning, laughing, hiccupping, and beyond*. Cambridge, MA: Belknap/Harvard University Press. [ch 19]

Provine, R. R., & Fischer, K. R. (1989). Laughing, smiling, and talking: Relation to sleeping and social context in humans. *Ethology, 83*, 295–305. [ch 18]

Provine, R. R., Krosnowski, K. A., & Brocato, N. W. (2009). Tearing: Breakthrough in human emotional signaling. *Evolutionary Psychology, 7*, 52–56. [ch 4]

Provine, R. R., Tate, B. C., & Geldmacher, L. L. (1987). Yawning: No effect of 3–5% CO_2, 100% O_2, and exercise. *Behavioral and Neural Biology, 48*, 382–393. [ch 14]

Pryor, K. W. (1999). *Don't shoot the dog!* (rev. ed.). New York: Bantam. [ch 9, 14]

Pryor, K. W. (2009). *Reaching the animal mind: Clicker training and what it teaches us about animals*. New York: Scribner. [ch 9, 14]

Pryor, K. W., Haag, R., & O'Reilly, J. (1969). The creative porpoise: Training for novel behavior. *Journal of the Experimental Analysis of Behavior, 12*, 653–661. [ch 9, 10, 13]

Pullum, G. K. (1991). The great Eskimo vocabulary hoax. In G. K. Pullum (Ed.), *The great Eskimo vocabulary hoax and other irreverent essays on the study of language* (pp. 159–171). Chicago: University of Chicago Press. [ch 25]

Rachlin, H. (1967). The effect of shock intensity on concurrent and single-key responding in concurrent-chain schedules. *Journal of the Experimental Analysis of Behavior, 10*, 87–93. [ch 7]

Rachlin, H. (1971). On the tautology of the matching law. *Journal of the Experimental Analysis of Behavior, 15*, 249–251. [ch 16]

Rachlin, H. (1974). Self-control. *Behaviorism, 2*, 94–107. [ch 16]

Rachlin, H., & Baum, W. M. (1972). Effect of alternative reinforcement: Does the source matter? *Journal of the Experimental Analysis of Behavior, 18*, 231–241. [ch 16]

Rachlin, H., & Burkhard, B. (1978). The temporal triangle: Response substitution in instrumental conditioning. *Psychological Review, 85*, 22–47. [ch 6]

Rachlin, H., & Green, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior, 17*, 15–22. [ch 16]

Rachlin, H., Green, L., Kagel, J. H., & Battalio, R. C. (1976). Economic demand theory and psychological studies of choice. In G. H. Bower (Ed.), *The psychology of learning and* *motivation* (Vol. 10, pp. 129–154). New York: Academic Press. [ch 16]

Rajala, A. Z., Rwininger, K. R., Lancaster, K. M., & Populin, L. C. (2010). Rhesus monkeys (*Macaca mulatta*) do recognize themselves in the mirror: Implications for the evolution of self-recognition. *PLoS ONE, 5*(9), e12865. [ch 19]

Ramus, F. (2002). Language discrimination by newborns: Teasing apart phonotactic, rhythmic, and intonational cues. *Annual Review of Language Acquisition, 2*, 85–115. [ch 4]

Randell, T., & Remington, B. (1999). Equivalence relations between visual stimuli: The functional role of naming. *Journal of the Experimental Analysis of Behavior, 71*, 395–415. [ch 13, 22]

Ratliff, F., & Hartline, H. K. (1959). The responses of Limulus optic nerve fibers to patterns of illumination on the receptor mosaic. *Journal of General Physiology, 42*, 1241–1255. [ch 16]

Ratner, S. C. (1970). Habituation: Research and theory. In J. H. Reynierse (Ed.), *Current issues in animal learning* (pp. 55–84). Lincoln, NE: University of Nebraska Press. [ch 4]

Reber, A. S. (1976). Implicit learning of synthetic languages: the role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory, 2*, 88–94. [ch 23]

Reber, A. S., Allen, R., & Reber, E. S. (2009). *The Penguin dictionary of psychology* (4th ed.). New York: Penguin Books. [gloss]

Recanzone, G. H., Merzenich, M. M., Jenkins, W. M., Garjski, K. A., & Dinse, H. R. (1992). Topographic reorganization of the hand representation in cortical area 3b of owl monkeys trained in a frequency-discrimination task. *Journal of Neurophysiology, 67*, 1031–1056. [ch 30]

Reddy, M. J. (1979). The conduit metaphor - a case of frame conflict in our language about language. In A. Ortony (Ed.), *Metaphor and thought* (pp. 284–324). New York: Cambridge University Press. [ch 22]

Reed, D. H. (2007). Natural selection and genetic diversity. *Heredity, 99*, 1–2. [ch 9, 13]

Reichle, J., Lindamood, L., & Sigafoos, J. (1986). The match between reinforcer class and response class: Its influence on communication intervention strategies. *Journal of the Association for Persons with Severe Handicaps, 11*, 131–135. [ch 14]

Reid, A. K., & Staddon, J. E. R. (1982). Schedule-induced drinking: Elicitation, anticipation, or behavior interaction? *Journal of the Experimental Analysis of Behavior, 38*, 1–18. [ch 4]

Reid, R. L. (1958). The role of the reinforcer as a stimulus. *British Journal of Psychology Monograph Supplements, 49*, 202–209. [ch 5]

Reinhard, G., & Lachnit, H. (2002). Differential conditioning of anticipatory pupillary dilation responses in humans. *Biological Psychology, 60*, 51–68. [ch 17]

Rendall, D., Owren, M. J., & Rodman, P. S. (1998). The role of vocal tract filtering in identity cueing in rhesus monkey (*Macaca mulatta*) vocalizations. *Journal of the Acoustical Society of America, 103*, 602–614. [ch 11]

Rendall, D., Rodman, P. S., & Emond, R. E. (1996). Vocal recognition of individuals and kin in free-ranging rhesus monkeys. *Animal Behaviour, 51*, 1007–1015. [ch 11, 24]

Rentfro, A. D. (2010, Nov/Dec). Turning fearful cats friendly. *Animal Sheltering*, 51–56. [ch 14]

Repp, A. C., & Deitz, S. M. (1974). Reducing aggressive and self-injurious behavior of institutionalized retarded children through reinforcement of other behaviors. *Journal of Applied Behavior Analysis, 7*, 313–325. [ch 5]

Rescorla, R. A. (1967). Pavlovian conditioning and its proper control procedures. *Psychological Review, 74*, 71–80. [ch 17]

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology, 66*, 1–5. [ch 18]

Rescorla, R. A. (1980). *Pavlovian second-order conditioning*. Hillside, NJ: Erlbaum. [ch 17]

Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist, 43*, 151–160. [ch 17]

Rescorla, R. A., & Solomon, R. L. (1967). Two process learning theory: Relationships between Pavlovian conditioning and instrumental learning. *Psychological Review, 74*, 151–182. [ch 17, 18]

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II* (pp. 64–99). New York: Appleton-Century-Crofts. [ch 17]

Restle, F. (1957). Discrimination of cues in mazes: A resolution of the "place versus response" question. *Psychological Review, 64*, 217–228. [ch 11]

Revusky, S. H., & Garcia, J. (1970). Learned associations over long delays. In G. H. Bower (Ed.), *The psychology of learning and motivation. Volume 4* (pp. 1–44). New York: Academic Press. [ch 18]

Reyna, V. F. (1981). The language of possibility and probability: Effects of negation on meaning. *Memory and Cognition, 9*, 642–650. [ch 25]

Reynolds, G. S. (1961a). Attention in the pigeon. *Journal of the Experimental Analysis of Behavior, 4*, 203–208. [ch 11]

Reynolds, G. S. (1961b). Behavioral contrast. *Journal of the Experimental Analysis of Behavior, 4*, 57–71. [ch 16]

Reynolds, G. S. (1966). Discrimination and emission of temporal intervals by pigeons. *Journal of the Experimental Analysis of Behavior, 9*, 65–68. [ch 19]

Reynolds, G. S., Catania, A. C., & Skinner, B. F. (1963). Conditioned and unconditioned aggression. *Journal of the Experimental Analysis of Behavior, 6*, 73–74. [ch 18]

Reynolds, G. S., & Skinner, B. F. (1962). Technique for reinforcing either of two organisms with a single food magazine. *Journal of the Experimental Analysis of Behavior, 5*, 58. [ch 18]

Riccio, D. C., Rabinowitz, V. C., & Axelrod, S. (1994). Memory: When less is more. *American Psychologist, 49*, 917–926. [ch 26]

Richards, R. W. (1988). The question of bidirectional associations in pigeons' learning of conditional discrimination tasks. *Bulletin of the Psychonomic Society, 26*, 577–579. [ch 12]

Richter, C. P. (1927). Animal behavior and internal drives. *Quarterly Review of Biology, 2*, 307–343. [ch 6]

Riess, B. F. (1946). Genetic changes in semantic conditioning. *Journal of Experimental Psychology, 36*, 143–152. [ch 24]

Risley, T. R. (1977). The development and maintenance of language: An operant model. In B. C. Etzel, J. M. LeBlanc & D. M. Baer (Eds.), *New developments in behavioral research* (pp. 81–101). Hillsdale, NJ: Erlbaum. [ch 20]

Risley, T. R., & Hart, B. (1968). Developing correspondence between the non-verbal and verbal behavior of preschool children. *Journal of Applied Behavior Analysis, 1*, 267–281. [ch 23]

Rizley, R. C., & Rescorla, R. A. (1972). Associations in second-order conditioning and sensory preconditioning. *Journal of Comparative and Physiological Psychology, 81*, 1–11. [ch 17]

Robinson, E. S., & Brown, M. A. (1926). Effect of serial position upon memorization. *American Journal of Psychology, 37*, 538–552. [ch 21]

Roediger, H. L., III. (1980). Memory metaphors in cognitive psychology. *Memory and Cognition, 8*, 231–246. [ch 26]

Roediger, H. L., III. (1990). Implicit memory: Retention without remembering. *American Psychologist, 45*, 1043–1056. [ch 27]

Roediger, H. L., III, & Craik, F. I. M. (Eds.). (1989). *Varieties of memory and consciousness*. Hillsdale, NJ: Erlbaum. [ch 27]

Rogers-Warren, A. R., & Baer, D. M. (1976). Correspondence between saying and doing: Teaching children to share and praise. *Journal of Applied Behavior Analysis, 9*, 335–354. [ch 23]

Rosales-Ruiz, J. (2011). *From fierce or fearful to friendly: A shaping program for dogs, cats, and beyond*. Paper presented at the Berkshire Association for Behavior Analysis and Therapy, Amherst, MA. [ch 8, 14]

Rosch, E. H. (1973). Natural categories. *Cognitive Psychology, 4*, 328–350. [ch 11]

Rosenfarb, I. S., Newland, M. C., Brannon, S. E., & Howey, D. S. (1992). Effects of self-generated rules on the development of schedule-controlled behavior. *Journal of the Experimental Analysis of Behavior, 58*, 107–121. [ch 9, 23]

Rosenfeld, H. M., & Baer, D. M. (1970). Unbiased and unnoticed verbal conditioning: The double agent robot procedure. *Journal of the Experimental Analysis of Behavior, 14*, 99–107. [ch 23]

Rosenhan, D. L. (1973). On being sane in insane places. *Science, 179*, 250–258. [ch 14, 23]

Ross, D., & Spurrett, D. (2004). What to say to a skeptical metaphysician: A defense manual for cognitive and behavioral scientists. *Behavioral and Brain Sciences, 27*, 603–647. [ch 30]

Rozin, P., & Kalat, J. W. (1971). Specific hungers and poison avoidance as adaptive specializations of learning. *Psychological Review, 78*, 459–486. [ch 18]

Rozin, P., Millman, L., & Nemeroff, C. (1986). Operation of the laws of sympathetic magic in disgust and other domains. *Journal of Personality and Social Psychology, 50*, 703–712. [ch 23]

Rudolph, R. L., Honig, W. K., & Gerry, J. E. (1969). Effects of monochromatic rearing on the acquisition of stimulus control. *Journal of Comparative and Physiological Psychology, 67*, 50–57. [ch 11]

Rudy, J. W., Vogt, M. B., & Hyson, R. L. (1984). A developmental analysis of the rat's learned reactions to gustatory and auditory stimulation. In R. Kail & N. E. Spear (Ed.), *Comparative perspectives on the development of memory* (pp. 181–208). Hillsdale, NJ: Erlbaum. [ch 3]

Rumbaugh, D. M., & Gill, T. V. (1976). The mastery of language-type skills by the chimpanzee (Pan). *Annals of the New York Academy of Sciences, 280*, 562–578. [ch 24]

Rundus, D. (1977). Maintenance rehearsal and single-level processing. *Journal of Verbal Learning and Verbal Behavior, 16*, 665–681. [ch 27]

Rundus, D. (1980). Maintenance rehearsal and long-term recency. *Memory and Cognition, 8*, 226–230. [ch 27]

Russell, E. S. (1916). *Form and function*. London: John Murray. [ch 1, 30]

Ryle, G. (1949). *The concept of mind*. New York: Barnes & Noble. [ch 1]

Sabin, P. (2012). *Simulating war*. London: Continuum. [ch 29]

Sachs, J. S. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception and Psychophysics, 2*, 437–442. [ch 6, 25]

Sagvolden, T., Johansen, E. B., Aase, H., & Russell, V. A. (2005). A dynamic developmental theory of Attention-Deficit/Hyperactivity Disorder (ADHD) predominantly hyperactive/impulsive and combined subtypes. *Behavioral and Brain Sciences, 28*, 397–468. [ch 16]

Salapatek, P., & Kessen, W. (1966). Visual scanning of triangles by the human newborn. *Journal of Experimental Child Psychology, 3*, 155–167. [ch 6]

Saugstad, P., & Raaheim, K. (1960). Problem solving, past experience and availability of functions. *British Journal of Psychology Monograph Supplements, 51*, 97–104. [ch 29]

Saunders, R. R., & Green, G. (1992). The nonequivalence of behavioral and mathematical equivalence. *Journal of the Experimental Analysis of Behavior, 57*, 227–241. [ch 12]

Saunders, R. R., & Green, G. (1999). A discrimination analysis of training-structure effects on stimulus equivalence outcomes. *Journal of the Experimental Analysis of Behavior, 72*, 117–137. [ch 13]

Savage-Rumbaugh, E. S. (1986). *Ape language.* New York: Columbia University Press. [ch 24]

Savage-Rumbaugh, E. S., Murphy, J., Sevcik, R. A., Brakke, K. E., Williams, S. L., & Rumbaugh, D. M. (1993). Language comprehension in ape and child. *Monographs of the Society for Research in Child Development, 58 (3–4, Serial No. 233).* [ch 24]

Savage-Rumbaugh, E. S., Rumbaugh, D. M., & Boysen, S. (1978). Symbolic communication between two chimpanzees (*Pan troglodytes*). *Science, 201*, 641–644. [ch 24]

Savage-Rumbaugh, E. S., Rumbaugh, D. M., Smith, S. T., & Lawson, J. (1980). Reference: The linguistic essential. *Science, 210*, 922–925. [ch 24]

Sawisch, L. P., & Denny, M. R. (1973). Reversing the reinforcement contingencies of eating and keypecking behaviors. *Animal Learning and Behavior, 1*, 189–192. [ch 6]

Schlinger, H., & Blakely, E. (1987). Function-altering effects of contingency-specifying stimuli. *The Behavior Analyst, 10*, 41–45. [ch 23]

Schlosberg, H. (1937). The relationship between success and the laws of conditioning. *Psychological Review, 44*, 379–394. [ch 17]

Schneider, S. M. (2012). *The science of consequences.* Amherst, NY: Prometheus. [ch 30]

Schoenfeld, W. N. (1950). An experimental approach to anxiety, escape and avoidance behavior. In P. H. Hoch (Ed.), *Anxiety* (pp. 70–99). New York: Grune & Stratton. [ch 8]

Schoenfeld, W. N. (1966). Some old work for modern conditioning theory. *Conditional Reflex, 1*, 219–223. [ch 4]

Schoenfeld, W. N. (1969). "Avoidance" in behavior theory. *Journal of the Experimental Analysis of Behavior, 12*, 669–674. [ch 8]

Schoenfeld, W. N. (Ed.). (1970). *The theory of reinforcement schedules.* New York: Appleton-Century-Crofts. [ch 15]

Schoenfeld, W. N., & Cole, B. K. (1972). *Stimulus schedules: The t-t systems.* New York: Harper & Row. [ch 15]

Schomer, H. L. (2002). The impact of class-specific reinforcers on conditional discrimination acquisition and equivalence class performances. *Experimental Analysis of Human Behavior Bulletin, 20*, 24–25. [ch 14]

Schooler, J. W., & Engstler-Schooler, T. Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology, 22*, 36–71. [ch 23]

Schroeder, S. R., & Holland, J. G. (1968). Operant control of eye movements. *Journal of Applied Behavior Analysis, 1*, 161–166. [ch 6]

Schuster, C. R., & Balster, R. L. (1977). The discriminative stimulus properties of drugs. In T. Thompson & P. Dews (Eds.), *Advances in behavioral pharmacology. Volume 1* (pp. 85–138). New York: Academic Press. [ch 11]

Schusterman, R. J., & Kastak, D. (1993). A California sea lion (Zalophus californianus) is capable of forming equivalence relations. *Psychological Record, 43*, 823–839. [ch 12, 24]

Schwartz, B. (1974). On going back to nature: A review of Seligman & Hager's Biological boundaries of learning. *Journal of the Experimental Analysis of Behavior, 21*, 183–198. [ch 18]

Schwartz, B. (1982a). Failure to produce response variability with reinforcement. *Journal of the Experimental Analysis of Behavior, 37*, 171–181. [ch 10]

Schwartz, B. (1982b). Reinforcement-induced stereotypy: How not to teach people to discover rules. *Journal of Experimental Psychology: General, 111*, 23–59. [ch 14]

Schwartz, B. (1986). *The battle for human nature.* New York: Norton. [ch 10]

Schwartz, B., & Williams, D. R. (1972). Two different kinds of key peck in the pigeon: Some properties of responses maintained by negative and positive response-reinforcer contingencies. *Journal of the Experimental Analysis of Behavior, 18*, 201–216. [ch 17]

Sebeok, T. A., & Rosenthal, R. (Eds.). (1981). *The Clever Hans phenomenon: Communication with horses, whales, apes, and people.* New York: New York Academy of Sciences. [ch 24]

Secan, K. E., Egel, A. L., & Tilley, C. S. (1989). Acquisition, generalization, and maintenance of question-answering skills in autistic children. *Journal of Applied Behavior Analysis, 22*, 181–196. [ch 19]

Sechenov, I. M. (1863/1965). *Reflexes of the brain (tr. S. Belsky).* Cambridge, MA: MIT Press. [ch 4]

Seligman, M. E. P. (1970). On the generality of the laws of learning. *Psychological Review, 77*, 406–418. [ch 8, 18]

Seligman, M. E. P., & Csikszentmihalyi, M. (2000). Positive psychology: An introduction. *American Psychologist, 55*, 5–14. [ch 5]

Seligman, M. E. P., & Csikszentmihalyi, M. (2001). Reply to comments. *American Psychologist, 56*, 89–90. [ch 5]

Seyfarth, R. M., Cheney, D. L., & Marler, P. (1980a). Monkey responses to three different alarm calls: Evidence for predator classification and semantic communication. *Science, 210*, 801–803. [ch 19]

Seyfarth, R. M., Cheney, D. L., & Marler, P. (1980b). Vervet monkey alarm calls: Semantic communication in a free-ranging primate. *Animal Behaviour, 28*, 1070–1094. [ch 24]

Shahn, B. B. (1972). *Ben Shahn.* New York: Abrams. [ch 20]

Shannon, C. E., & Weaver, W. (1949). *The mathematical theory of information.* Urbana: University of Illinois Press. [ch 28]

Sheffield, F. D. (1965). Relation between classical conditioning and instrumental learning. In W. F. Prokasy (Ed.), *Classical conditioning* (pp. 302–322). New York: Appleton-Century-Crofts. [ch 8, 17]

Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science, 171*, 701–703. [ch 28]

Sherman, T. M., & Cormier, W. H. (1974). An investigation of the influence of student behavior on teacher behavior. *Journal of Applied Behavior Analysis, 7*, 11–21. [ch 6]

Sherrington, C. (1906). *The integrative action of the nervous system.* New York: Scribners. [ch 4]

Shettleworth, S. J. (1978). Reinforcement and the organization of behavior in golden hamsters: Punishment of three action patterns. *Learning and Motivation, 9*, 99–123. [ch 7]

Shiffrin, R. M., & Atkinson, R. C. (1969). Storage and retrieval processes in long-term memory. *Psychological Review, 76*, 179–193. [ch 27]

Shimoff, E., & Catania, A. C. (1995). Using computers to teach behavior analysis. *Behavior Analyst, 18*, 307–316. [ch 13, 29]

Shimoff, E., & Catania, A. C. (1998). The verbal governance of behavior. In K. A. Lattal & M. Perone (Eds.), *Handbook of research methods in human operant behavior* (pp. 371–404). New York: Plenum. [ch 14, 23]

Shimoff, E., & Catania, A. C. (2001). Effects of recording attendance on grades in introductory psychology. *Teaching of Psychology, 28*, 192–195. [ch 13]

Shimoff, E., Catania, A. C., & Matthews, B. A. (1981). Uninstructed human responding: Sensitivity of low-rate performance to schedule contingencies. *Journal of the Experimental Analysis of Behavior, 36*, 207–220. [ch 23]

Shimoff, E., Catania, A. C., & Matthews, B. A. (1987). Psychology on a disk: Computer programs in introductory psychology. Ellicott City, MD: CMS Software. [ch 13, 29]

Shimp, C. (1966). Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior, 9*, 443–455. [ch 16]

Shimp, C. P., Sabulsky, S. L., & Childers, L. J. (1989). Preference for starting and finishing behavior patterns. *Journal of the Experimental Analysis of Behavior, 52*, 341–352. [ch 12, 19]

Shull, R. L. (1995). Interpreting cognitive phenomena: Review of Donahoe and Palmer's Learning and Complex Behavior. *Journal of the Experimental Analysis of Behavior, 63*, 347–358. [ch 30]

Sidman, M. (1952). A note on functional relations obtained from group data. *Psychological Bulletin, 49*, 263–269. [ch 5]

Sidman, M. (1953). Two temporal parameters in the maintenance of avoidance behavior by the white rat. *Journal of Comparative and Physiological Psychology, 46*, 253–261. [ch 8]

Sidman, M. (1960). *Tactics of scientific research.* New York: Basic Books. [ch 6]

Sidman, M. (1971). The behavioral analysis of aphasia. *Journal of Psychiatric Research, 8*, 413–422. [ch 20]

Sidman, M. (1987). Two choices are not enough. *Behavior Analysis, 22*, 11–18. [ch 12]

Sidman, M. (1994). *Equivalence relations and behavior: A research story.* Boston: Authors Cooperative. [ch 12]

Sidman, M. (2000). Equivalence relations and the reinforcement contingency. *Journal of the Experimental Analysis of Behavior, 74*, 127–146. [ch 13]

Sidman, M., Cresson, O., Jr., & Willson-Morris, M. (1974). Acquisition of matching to sample via mediated transfer. *Journal of the Experimental Analysis of Behavior, 22*, 261–273. [ch 12]

Sidman, M., Herrnstein, R. J., & Conrad, D. G. (1957). Maintenance of avoidance behavior by unavoidable shocks. *Journal of Comparative and Physiological Psychology, 50*, 553–557. [ch 18]

Sidman, M., & Stoddard, L. T. (1967). The effectiveness of fading in programming a simultaneous form discrimination for retarded children. *Journal of the Experimental Analysis of Behavior, 10*, 3–15. [ch 11]

Sidman, M., Wynne, C. K., Maguire, R. W., & Barnes, T. (1989). Functional classes and equivalence relations. *Journal of the Experimental Analysis of Behavior, 52*, 261–274. [ch 12]

Siegel, S. (1977). Morphine tolerance acquisition as an associative process. *Journal of Experimental Psychology: Animal Behavior Processes, 3*, 1–13. [ch 17]

Siegel, S., Hinson, J. M., Krank, M. D., & McCully, J. (1982). Heroin "overdose" death: The contribution of drug-associated environmental cues. *Science, 216*, 436–437. [ch 17]

Silverman, P. J. (1971). Chained and tandem fixed-interval schedules of punishment. *Journal of the Experimental Analysis of Behavior, 16*, 1–13. [ch 16]

Simmons, R. E., & Scheepers, L. (1996). Winning by a neck: Sexual selection in the evolution of the giraffe. *American Naturalist, 148*(5), 771–786. [ch 3]

Simpson, G. G. (1951). *Horses.* New York: Oxford University Press. [ch 3]

Sizemore, O. J., & Lattal, K. A. (1977). Dependency, temporal contiguity, and response-independent reinforcement. *Journal of the Experimental Analysis of Behavior, 27*, 119–125. [ch 15]

Skiba, E. A., Pettigrew, L. E., & Alden, S. E. (1971). A behavioral approach to the control of thumbsucking in the classroom. *Journal of Applied Behavior Analysis, 4*, 121–125. [ch 5]

Skinner, B. F. (1930). On the conditions for elicitation of certain eating reflexes. *Proceedings of the National Academy of Sciences, 16*, 433–438. [ch 5, 30]

Skinner, B. F. (1931). The concept of the reflex in the description of behavior. *Journal of General Psychology, 5*, 427–458. [ch 1, 4]

Skinner, B. F. (1932). Drive and reflex strength. *Journal of General Psychology, 6*, 22–37. [ch 14]

Skinner, B. F. (1933). The rate of establishment of a discrimination. *Journal of General Psychology, 9*, 302–350. [ch 2]

Skinner, B. F. (1934a). The extinction of chained reflexes. *Proceedings of the National Academy of Sciences, 20*, 234–237. [ch 10]

Skinner, B. F. (1934b). Has Gertrude Stein a secret? *Atlantic Monthly*, reprinted in Cumulative Record (4th edition, pp. 405–415). [ch 23]

Skinner, B. F. (1935a). The generic nature of the concepts of stimulus and response. *Journal of General Psychology, 12*, 40–65. [ch 9, 12]

Skinner, B. F. (1935b). Two types of conditioned reflex and a pseudotype. *Journal of General Psychology, 12*, 66–77. [ch 17]

Skinner, B. F. (1936). The verbal summator and a method for the study of latent speech. *Journal of Psychology, 2*, 71–107. [ch 20]

Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York: Appleton-Century-Crofts. [ch 2, 4, 5, 6, 15, 30]

Skinner, B. F. (1945). The operational analysis of psychological terms. *Psychological Review, 52*, 270–277. [ch 22, 25, 28, 29]

Skinner, B. F. (1948). "Superstition" in the pigeon. *Journal of Experimental Psychology: Animal Behavior Processes, 38*, 168–172. [ch 5]

Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review, 57*, 193–216. [ch 5, 12, 15, 16, 29]

Skinner, B. F. (1951). How to teach animals. *Scientific American, 185*(Dec), 26–29. [ch 9]

Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan. [ch 2, 7, 12, 16, 24, 28]

Skinner, B. F. (1956). A case history in scientific method. *American Psychologist, 11*, 221–233. [ch 5, 15, 16, 29]

Skinner, B. F. (1957). *Verbal behavior*. New York: Appleton-Century-Crofts. [glossA, glossT, ch 13, 20, 21, 22, 23, 24, 25, 29, 30, gloss]

Skinner, B. F. (1959a). An experimental analysis of certain emotions. *Journal of the Experimental Analysis of Behavior, 2*, 264. [ch 18]

Skinner, B. F. (1959b). John Broadus Watson, behaviorist. *Science, 129*, 197–198. [ch 28]

Skinner, B. F. (1960). Pigeons in a pelican. *American Psychologist, 15*, 28–37. [ch 9]

Skinner, B. F. (1961). The flight from the laboratory. In W. D. e. al. (Ed.), *Current trends in psychology theory*. Pittsburgh: University of Pittsburgh Press. [ch 29, 30]

Skinner, B. F. (1963). Behaviorism at fifty. *Science, 140*, 951–958. [ch 22, 26]

Skinner, B. F. (1964/1980). 20: Education in 1984. In *The Oxford dictionary of quotations* (3rd ed., pp. 508). New York: Oxford University Press. [ch 13]

Skinner, B. F. (1966). The phylogeny and ontogeny of behavior. *Science, 153*, 1204–1213. [ch 3, 30]

Skinner, B. F. (1968). *The technology of teaching*. New York: Appleton-Century-Crofts. [ch 13, 24, 30]

Skinner, B. F. (1969). An operant analysis of problem solving. In *B. F. Skinner, Contingencies of reinforcement* (pp. 133–157). New York: Appleton-Century-Crofts. [ch 23, 29]

Skinner, B. F. (1972). A lecture on "having" a poem. In *Cumulative record* (3rd ed., pp. 345–355). New York: Appleton-Century-Crofts. [ch 20, 25]

Skinner, B. F. (1975). The shaping of phylogenic behavior. *Journal of the Experimental Analysis of Behavior, 24*, 117–120. [ch 12, 30]

Skinner, B. F. (1976). *Particulars of my life*. New York: Knopf. [ch 26]

Skinner, B. F. (1977a). The force of coincidence. *The Humanist, 37 (3, May/June)*, 10–11. [ch 13]

Skinner, B. F. (1977b). Herrnstein and the evolution of behaviorism. *American Psychologist, 32*, 1006–1012. [ch 18]

Skinner, B. F. (1981). Selection by consequences. *Science, 213*, 501–504. [ch 3, 19, 30]

Skinner, B. F. (1983). Intellectual self-management in old age. [ch 26]

Skinner, B. F. (1984). The evolution of behavior. *Journal of the Experimental Analysis of Behavior, 41*, 217–221. [ch 4]

Skinner, B. F. (1986). The evolution of verbal behavior. *Journal of the Experimental Analysis of Behavior, 45*, 115–122. [ch 20, 24]

Skinner, B. F. (1988). Responses to commentaries. In A. C. Catania & S. Harnad (Eds.), *The Selection of Behavior: The Operant Behaviorism of B. F. Skinner*. New York: Cambridge University Press. [ch 3. 28, 30]

Skinner, B. F. (1989a). The listener. In *Recent issues in the analysis of behavior* (pp. 35–47). Columbus, OH: Merrill. [ch 24]

Skinner, B. F. (1989b). The origins of cognitive thought. *American Psychologist, 44*, 13–18. [ch 22, 25, gloss]

Small, W. S. (1899–1900). Experimental studies of the mental processes of the rat. *American Journal of Psychology, 11*, 1–89. [ch 5, 30]

Smith, B. H. (1968). *Poetic closure*. Chicago, IL: Chicago University Press. [ch 20]

Smith, J. B. (1974). Effects of response rate, reinforcement frequency, and the duration of a stimulus preceding response-independent food. *Journal of the Experimental Analysis of Behavior, 21*, 215–221. [ch 18]

Smith, J. D., Schull, J., Strote, J., McGee, K., Egnor, R., & Erb, L. (1995). The uncertain response in the bottlenosed dolphin (Tursiops truncatus). *Journal of Experimental Psychology: General, 124*, 391–408. [ch 22]

Smith, J. D., & Washburn, D. A. (2005). Uncertainty monitoring and metacognition by animals. *Current Directions in Psychological Science, 14*, 19–24. [ch 22]

Smith, K. (1954). Conditioning as an artifact. *Psychological Review, 61*, 217–225. [ch 17]

Smith, T. L. (1986). Biology as allegory: A review of Elliott Sober's The nature of selection. *Journal of the Experimental Analysis of Behavior, 46*, 105–112. [ch 3]

Smith, W. G. (1895). The relation of attention to memory. *Mind, 4*, 47–73. [ch 27]

Snapper, A. G., Kadden, R. M., & Inglis, G. B. (1982). State notation of behavioral procedures. *Behavior Research Methods and Instrumentation, 14*, 329–342. [ch 15]

Solnick, J. V., Rincover, A., & Peterson, C. R. (1977). Some determinants of the reinforcing and punishing effects of timeout. *Journal of Applied Behavior Analysis, 10*, 415–424. [ch 7]

Solomon, R. L., & Corbit, J. D. (1974). An opponent-process theory of motivation: I. Temporal dynamics of affect. *Psychological Review, 81*, 119–145. [ch 4]

Solomon, R. L., & Turner, L. H. (1962). Discriminative classical conditioning in dogs paralyzed by curare can later control discriminative avoidance responses in the normal state. *Psychological Review, 69*, 202–219. [ch 17]

Spalding, D. (1873/1954). Instinct with original observations on young animals (Reprinted British Journal of Animal Behaviour, 2, 2–11). *Macmillan's Magazine, 27*, 282–293. [ch 3]

Spence, K. W., & Ross, L. E. (1959). A methodological study of the form and latency of eyelid responses in conditioning. *Journal of Experimental Psychology, 58*, 376–381. [ch 4]

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74*(11, Whole No. 498). [ch 27]

Sperling, G., & Reeves, A. (1980). Measuring the reaction time of a shift of visual attention. In R. S. Nickerson (Ed.), *Attention and performance VIII* (pp. 347–360). Hillsdale, NJ: Erlbaum. [ch 28]

Spetch, M. L., Wilkie, D. M., & Pinel, J. P. J. (1981). Backward conditioning: A reevaluation of the empirical evidence. *Psychological Bulletin, 89*, 163–175. [ch 17]

Squier, L. H. (1993). The science and art of training: A review of Pryor's Lads before the wind. *Journal of the Experimental Analysis of Behavior, 59*, 423–431. [ch 9]

Squire, L. R. (Ed.). (1992). *The encyclopedia of learning and memory*. New York: Macmillan. [ch 26]

Sran, S. K., & Borrero, J. C. (2010). Assessing the value of choice in a token system. *Journal of Applied Behavior Analysis, 43*, 553–557. [ch 14]

Staats, A. W. (1986). Behaviorism with a personality: The paradigmatic behavioral assessment approach. In R. O. Nelson & S. C. Hayes (Eds.), *Conceptual foundations of behavioral assessment* (pp. 242–296). New York: Guilford. [ch 30]

Staddon, J. E. R., & Simmelhag, V. L. (1971). The "superstition" experiment: A reexamination of its implications for the principles of adaptive behavior. *Psychological Review, 78*, 3–43. [ch 4, 5]

Stanners, R. F., Coulter, M., Sweet, A. W., & Murphy, P. (1979). The pupillary response as an indicator of arousal and cognition. *Motivation and Emotion, 3*, 319–340. [ch 17]

Stein, L., Xue, B. G., & Belluzzi, J. D. (1993). A cellular analogue of operant conditioning. *Journal of the Experimental Analysis of Behavior, 60*, 41–53. [ch 30]

Steinman, W. M. (1968). Response rate and varied reinforcement: Reinforcers of similar strengths. *Psychonomic Science, 10*, 35–36. [glossR, ch 14]

Sternberg, S. (1969). Memory-scanning: Mental processes revealed by reaction time experiments. *American Scientist, 57*, 421–457. [ch 28]

Stewart, G. R. (1975). *Names on the land*. New York: Oxford University Press. [ch 22]

Stokes, P. D., & Balsam, P. D. (1991). Effects of reinforcing preselected approximations on the topography of the rat's bar press. *Journal of the Experimental Analysis of Behavior, 55*, 213–231. [ch 5]

Stratton, G. M. (1897). Vision without inversion of the retinal image. *Psychological Review, 4*, 341–360; 463–481. [ch 6]

Stratton, G. M. (1917). *Theophrastus and the Greek physiological psychology before Aristotle*. New York: Macmillan. [ch 26]

Straub, R. O., Seidenberg, M. S., Bever, T. G., & Terrace, H. S. (1979). Serial learning in the pigeon. *Journal of the Experimental Analysis of Behavior, 32*, 137–148. [ch 10]

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*, 643–662. [ch 24, 26]

Sturgis, E. T., Tollison, C. D., & Adams, H. E. (1978). Modification of combined migraine-muscle contraction headaches using BVP and EMG feedback. *Journal of Applied Behavior Analysis, 11*, 215–223. [ch 19]

Sutherland, A. (2006). *Kicked, bitten, and scratched*. New York: Viking. [ch 9]

Sutherland, A. (2008). *What Shamu taught me about life, love, and marriage*. New York: Random House. [ch 9]

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press. [ch 29]

Svartdal, F. (1992). Sensitivity to nonverbal operant contingencies: Do limited processing resources affect operant conditioning in humans? *Learning and Motivation, 23*, 383–405. [ch 23]

Swets, J. A., Dawes, R. M., & Monahan, J. (2000). Psychological science can improve diagnostic decisions. *Psychological Science in the Public Interest, 1*, 1–26. [ch 9, 11]

Terrace, H. S. (1963a). Discrimination learning with and without "errors." *Journal of the Experimental Analysis of Behavior, 6*, 1–27. [ch 11]

Terrace, H. S. (1963b). Errorless transfer of a discrimination across two continua. *Journal of the Experimental Analysis of Behavior, 6*, 223–232. [ch 11]

Terrace, H. S. (1975). Evidence of the innate basis of the hue dimension in the duckling. *Journal of the Experimental Analysis of Behavior, 24*, 79–87. [ch 11]

Terrace, H. S., & Chen, S. (1991). Chunking during serial learning by a pigeon: III. what are the necessary conditions for establishing a chunk. *Journal of Experimental Psychology: Animal Behavior Processes, 17*, 107–118. [ch 12]

Terrace, H. S., Petitto, L. A., Sanders, R. J., & Bever, T. G. (1979). Can an ape create a sentence? *Science, 206*, 891–902. [ch 24]

Thelen, E., & Fisher, D. M. (1983). From spontaneous to instrumental behavior: Kinematic analysis of movement changes during very early learning. *Child Development, 54*, 429–440. [ch 18]

Thelen, E., Fisher, D. M., Ridley-Johnson, R., & Griffin, N. J. (1982). Effects of body build and arousal on newborn infant stepping. *Developmental Psychobiology, 15*, 447–453. [ch 4]

Thistlethwaite, D. (1951). A critical review of latent learning and related experiments. *Psychological Bulletin, 48*, 97–129. [ch 6]

Thomas, J. C., Jr. (1974). An analysis of behavior in the hobbits-orcs problem. *Cognitive Psychology, 6*, 257–269. [ch 29]

Thomas, J. R. (1979). Matching-to-sample accuracy on fixed-ratio schedules. *Journal of the Experimental Analysis of Behavior, 32*, 183–189. [ch 16]

Thomas, L. (1992). *The fragile species.* NY: Touchstone. [ch 3]

Thompson, C. R., & Church, R. M. (1980). An explanation of the language of a chimpanzee. *Science, 208*, 313–314. [ch 24]

Thompson, R. H., Iwata, B. A., Hanley, G. P., Dozier, C. L., & Samaha, A. L. (2003). The effects of extinction, noncontingent reinforcement, and differential reinforcement of other behavior as control procedures. *Journal of Applied Behavior Analysis, 36*, 221–238. [ch 5]

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review Monograph Supplements,* 2(4). [ch 2, 5, 29, 30]

Thorndike, E. L. (1911). *Animal intelligence.* New York: Macmillan. [ch 29]

Thorndike, E. L. (1921). *Educational psychology. Volume II. The psychology of learning.* New York: Teachers College. [ch 4, 30]

Thorndike, E. L., & Woodworth, R. S. (1901). The influence of improvement in one mental function upon the efficiency of other functions. *Psychological Review, 8*, 247–261. [ch 26, 29, 30]

Timberlake, W. (1980). A molar equilibrium theory of learned performance. In G. H. Bower (Ed.), *The psychology of learning and motivation. Volume 14* (pp. 1–58). New York: Academic Press. [ch 6]

Tinbergen, N. (1960). *The Herring Gull's world (Revised ed.).* New York: Basic Books. [ch 2]

Tinbergen, N. (1972). *The animal in its world. Volume 1. Field studies.* Cambridge, MA: Harvard University Press. [ch 4, 12]

Tinbergen, N., & Perdeck, A. C. (1950). On the stimulus situation releasing the begging response in the newly hatched Herring Gull chick (Larus a. argentatus Pontopp). *Behavior Analyst,* 3, 1–38. [ch 2]

Titchener, E. B. (1898). The postulates of a structural psychology. *Philosophical Review, 7*, 449–465. [ch 30]

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review, 55*, 189–208. [ch 6, 12, 30]

Tolman, E. C., & Honzik, C. H. (1930). Introduction and removal of reward, and maze performance in rats. *University of California Publications in Psychology. 4*, 257–275. [ch 6]

Tomasello, M., & Farrar, M. J. (1986). Joint attention and early language. *Child Development, 57*, 1454–1463. [ch 11, 22]

Touchette, P. E. (1969). Tilted lines as complex stimuli. *Journal of the Experimental Analysis of Behavior, 12*, 211–214. [ch 11]

Townsend, J. T. (1971). A note on the identifiability of parallel and serial processes. *Perception and Psychophysics, 10*, 161–163. [ch 28]

Truax, C. B. (1966). Reinforcement and nonreinforcement in Rogerian therapy. *Journal of Abnormal Psychology, 71*, 1–9. [ch 23]

Tulving, E. (1962). Subjective organization of free recall in "unrelated" words. *Psychological Review, 69*, 344–354. [ch 21]

Tulving, E. (1969). Retrograde amnesia in free recall. *Science, 164*, 88–90. [ch 21]

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 381–403). New York: Academic Press. [ch 27]

Tulving, E. (1974). Cue-dependent forgetting. *American Scientist, 62*, 74–82. [ch 21, 26]

Tulving, E. (1985). How many memory systems are there? *American Psychologist, 40*(4), 385–398. [ch 27]

Tulving, E., & Madigan, S. A. (1970). Memory and verbal learning. *Annual Review of Psychology, 21*, 437–484. [ch 26, 27]

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior, 5*, 381–391. [ch 26]

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90*, 293–315. [ch 29]

Twitmyer, E. B. (1974). A study of the knee jerk (1902). *Journal of Experimental Psychology, 103*, 1047–1966. [ch 17]

Twyman, J. S. (1998). The Fred S. Keller School. *Journal of Applied Behavior Analysis, 31*, 695–701. [ch 13]

Underwood, B. J. (1957). Interference and forgetting. *Psychological Review, 64*, 49–60. [ch 27]

Underwood, B. J. (1961). Ten years of massed practice on distributed practice. *Psychological Review, 68*, 229–247. [ch 21]

Underwood, B. J. (1964). Degree of learning and measurement of forgetting. *Journal of Verbal Learning and Verbal Behavior, 3*, 112–129. [ch 27]

Underwood, B. J., & Postman, L. (1960). Extraexperimental sources of interference in forgetting. *Psychological Review, 67*, 73–95. [ch 27]

Urcuioli, P. J. (1985). On the role of differential sample behaviors in matching-to-sample. *Journal of Experimental Psychology: Animal Behavior Processes, 11*, 502–519. [ch 12, 14]

Urcuioli, P. J. (1991). Retardation and facilitation of matching acquisition by differential outcomes. *Animal Learning and Behavior, 19*, 29–36. [ch 14]

Urcuioli, P. J. (2005). Behavioral and associative effects of differential outcomes in discrimination learning. *Learning and Behavior, 33*, 1–21. [glossR, ch 14]

Urcuioli, P. J., & DeMarse, T. B. (1994). On the relationship between differential outcomes and differential sample responding in matching-to-sample. *Journal of Experimental Psychology: Animal Behavior Processes, 20*, 249–263. [ch 14]

Urcuioli, P. J., & Zentall, T. R. (1986). Retrospective coding in pigeons' delayed matching-to-sample. *Journal of Experimental Psychology: Animal Behavior Processes, 12*, 69–77. [ch 27]

Vaughan, W., Jr. (1988). Formation of equivalence sets in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 14*, 36–42. [ch 12]

Vellend, D. H. (2006). The consequences of genetic diversity in competitive communities. *Ecology, 87*, 304–311. [ch 9, 13]

Verhave, T. (1967). Contributions to the history of psychology: IV. Joseph Buchanan (1785–1829) and the "law of exercise" (1812). *Psychological Reports, 20*, 127–133. [ch 4]

Verplanck, W. S. (1957). A glossary of some terms used in the objective science of behavior. *Psychological Review, 64 (Supplement, Part 2)*, i-42. [gloss]

Verplanck, W. S. (1962/1992). Verbal concept "mediators" as simple operants (reprinted from: Unaware of where's awareness, pp. 130–158 in C. W. Eriksen, ed., Behavior and Awareness, 1962). *Analysis of Verbal Behavior, 10*, 45–68. [ch 13]

Verplanck, W. S. (2000). *Glossary/thesaurus of behavioral terms (CD-ROM).* Columbia, MD: CMS Software. [gloss]

Vesonder, G. T., & Voss, J. F. (1985). On the ability to predict one's own responses while learning. *Journal of Memory and Language, 24*, 363–376. [ch 19]

Vihman, M. M. (1996). *Phonological development: The origins of language in the child.* Oxford: Blackwell. [ch 19, 20]

Vollmer, T. R., & Athens, E. (2011). Developing function-based extinction procedures for problem behavior. In W. W. Fisher, C. C. Piazza & H. S. Roane (Eds.), *Handbook of applied behavior analysis* (pp. 317–334). New York: Guikford. [ch 7, 8]

Vollmer, T. R., Borrero, J. C., & Wright, C. S. (2001). Identifying possible contingencies during descriptive analyses of severe behavior disorders. *Journal of Applied Behavior Analysis, 34*, 269–287. [ch 14]

Vollmer, T. R., & Iwata, B. A. (1991). Establishing operations and reinforcement effects. *Journal of Applied Behavior Analysis, 24*, 279–291. [ch 6]

von Békésy, G. (1967). *Sensory inhibition.* Princeton, NJ: Princeton University Press. [ch 16]

Von Holst, E. (1973). *The behavioural physiology of animals and man. Selected papers.* Coral Gables, FL: University of Miami Press. [ch 4]

Von Restorff, H. (1933). Über die Wirkung von Bereichsbildungen im Spurenfeld. *Psychologische Forschung, 18*, 299–342. [ch 21]

Vonnegut, K. (1963). *Cat's cradle.* New York: Dell. [ch 3]

Waddington, C. H. (1957). *The strategy of the genes.* London: Allen & Unwin. [ch 3]

Waddington, C. H. (1959). Canalization of development and genetic assimilation of acquired characteristics. *Nature, 183*, 1654–1655–. [ch 3]

Wagner, A. R., Thomas, E., & Norton, T. (1967). Conditioning with electrical stimulation of motor cortex: Evidence of a possible source of motivation. *Journal of Comparative and Physiological Psychology, 64*, 191–199. [ch 17]

Wagner, K. R. (1985). How much do children say in a day? *Journal of Child Language, 12*, 475–487. [ch 22]

Wahler, R. G. (1975). Some structural aspects of deviant child behavior. *Journal of Applied Behavior Analysis, 8*, 27–42. [ch 9]

Walcott, C., Gould, J. L., & Kirschvink, J. L. (1979). Pigeons have magnets. *Science, 205*, 1027–1029. [ch 3, 12]

Wales, R. (1986). Deixis. In P. Fletcher & M. Garman (Eds.), *Language acquisition* (2nd ed., pp. 401–428). New York: Cambridge University Press. [ch 25]

Walters, G. C., & Glazer, R. D. (1971). Punishment of instinctive behavior in the Mongolian gerbil. *Journal of Comparative and Physiological Psychology, 75*, 331–340. [ch 7]

Wanchisen, B. A., Tatham, T. A., & Hineline, P. N. (1988). Pigeons' choices in situations of diminishing returns: Fixed- versus progressive-ratio schedules. *Journal of the Experimental Analysis of Behavior, 50*, 375–394. [ch 16]

Ward, T. B. (1980). Separable and integral responding by children and adults to the dimensions of length and density. *Child Development, 51*, 676–684. [ch 28]

Washburn, D. A., Hopkins, W. D., & Rumbaugh, D. M. (1991). Perceived control in rhesus monkeys (macaca mulatta): Enhanced video-task performance. *Journal of Experimental Psychology: Animal Behavior Processes, 17*, 123–129. [ch 12]

Wasik, B. H. (1970). The application of Premack's generalization on reinforcement to the management of classroom behavior. *Journal of Experimental Child Psychology, 10*, 33–43. [ch 6]

Wason, P. C., & Johnson-Laird, P. N. (1970). A conflict between selecting and evaluating information in an inferential task. *British Journal of Psychology Monograph Supplements, 61*, 509–515. [ch 29]

Wasserman, E., Franklin, S., & Hearst, E. (1974). Pavlovian appetitive contingencies and approach vs. withdrawal to conditioned stimuli in pigeons. *Journal of Comparative and Physiological Psychology, 86*, 616–627. [ch 17]

Wasserman, E. A., Kiedinger, R. E., & Bhatt, R. S. (1988). Conceptual behavior in pigeons: Categories, subcategories, and pseudocategories. *Journal of Experimental Psychology: Animal Behavior Processes, 14*, 235–246. [ch 12]

Watkins, C. (2011). Indo-European and the Indo-Europeans. In *The American Heritage Dictionary of the English Language* (pp. 2023–2031). Boston: Houghton Mifflin Harcourt. [ch 19]

Watkins, M. J. (1981). Human memory and the information-processing metaphor. *Cognition, 10*, 331–336. [ch 28, 30]

Watkins, M. J. (1989). Willful and nonwillful determinants of memory. In I. H. L. Roediger & F. I. M. Craik (Eds.), *Varieties of memory and consciousness* (pp. 59–71). Hillsdale, NJ: Erlbaum. [ch 27]

Watkins, M. J. (1990). Mediationism and the obfuscation of memory. *American Psychologist, 45*, 328–335. [ch 26]

Watkins, M. J. (1996). Encore: Mediationism and the obfuscation of memory. *The Behavior Analyst, 19*, 89–110. [ch 26]

Watkins, M. J., & Tulving, E. (1975). Episodic memory: When recognition fails. *Journal of Experimental Psychology: General, 104*, 5–29. [ch 21]

Watson, J. B. (1913). Psychology as the behaviorist views it. *Psychological Review, 20*, 158–177. [ch 30]

Watson, J. B. (1919). *Psychology from the standpoint of a behaviorist.* Philadelphia: Lippincott. [ch 4, 30]

Watson, J. B., & Rayner, R. (1920). Conditioned emotional reactions. *Journal of Experimental Psychology, 3*, 1–14. [ch 17]

Watson, J. D., & Crick, F. H. C. (1953). Molecular structure of nucleic acids. *Nature, 171*, 737–738. [ch 3, 30]

Waugh, N. C., & Norman, D. A. (1965). Primary memory. *Psychological Review, 72*, 89–104. [ch 27]

Weiner, H. (1983). Some thoughts on discrepant human–animal performances under schedules of reinforcement. *Psychological Record, 33*, 521–532. [ch 23]

Weiner, J. (1994). *The beak of the finch.* New York: Knopf. [ch 3]

Weiss, B., & Laties, V. G. (1961). Behavioral thermoregulation. *Science, 133*, 1338–1344. [ch 8]

Weiss, B., & Laties, V. G. (1969). Behavioral pharmacology and toxicology. *Annual Review of Pharmacology, 9*, 297–326. [ch 16]

Wellman, H. M. (1990). *The child's theory of mind.* Cambridge, MA: MIT Press. [ch 25]

Wells, H. G. (1920). *The outline of history.* Garden City, NY: Garden City Publishing. [ch 23]

Werker, J. F. (1989). Becoming a native listener. *American Scientist, 77*, 54–59. [ch 20]

Wertheimer, M. (1959). *Productive thinking.* New York: Harper and Row. [ch 29]

West, M. J., & King, A. P. (1980). Enriching cowbird song by social deprivation. *Journal of Comparative and Physiological Psychology, 94*, 263–270. [ch 24]

West-Eberhart, M. J. (2003). *Developmental plasticity and evolution.* New York: Oxford University Press. [ch 3, 13, 30]

Wetherington, C. L. (1982). Is adjunctive behavior a third class of behavior? *Neuroscience and Biobehavioral Reviews, 6*, 329–350. [ch 4]

Whalen, C., & Schreibman, L. (1997). Joint attention training for children with autism using behavior modification procedures. *Journal of Child Psychology and Psychiatry, 44*, 456–468. [ch 19]

Whitehurst, G. J., & Valdez-Menchaca, M. C. (1988). What is the role of reinforcement in early language acquisition? *Child Development, 59*, 430–440. [ch 25]

White, K. G. (2001). Forgetting functions. *Animal Learning and Behavior, 29*, 193–207. [ch 27]

White, K. G. (2002). Psychophysics of remembering: The discrimination hypothesis. *Current Directions in Psychological Science, 11*, 141–145. [ch 27]

White, K. G., & Wixted, J. T. (1999). Psychophysics of remembering. *Journal of the Experimental Analysis of Behavior, 71*, 91–113. [ch 27]

Whitley, D. S. (2009). *Cave paintings and the human spirit: The origin of creativity and belief.* Buffalo, NY: Prometheus. [ch 19]

Wildemann, D. G., & Holland, J. G. (1972). Control of a continuous response dimension by a continuous stimulus dimension. *Journal of the Experimental Analysis of Behavior, 18*, 419–434. [ch 19]

Wilkes, G. (1994). *A behavior sample.* North Bend, WA: Sunshine Books. [ch 7]

Wilkes, G. (2009). *Punish or perish: The desperate need for punishment in behavior modification.* Paper presented at the Association for Behavior Analysis International, Phoenix, AZ. [ch 7]

Williams, D. R., & Williams, H. (1969). Auto-maintenance in the pigeon: Sustained pecking despite contingent non-reinforcement. *Journal of the Experimental Analysis of Behavior, 12*, 511–520. [ch 17]

Wilson, D. M. (1959). Long term facilitation in a swimming sea anemone. *Journal of Experimental Biology, 36*, 526–531. [ch 4]

Wilson, D. S., & Sober, E. (1994). Reintroducing group selection in the human behavioral sciemces. *Behavioral and Brain Sciences, 17*, 585–608. [ch 19]

Wilson, E. O. (2012). *The social conquest of earth.* New York: Liveright. [ch 3, 19]

Wilson, K. G., & Hayes, S. C. (2000). Why it is crucial to understand thinking and feeling: An analysis and application to drug abuse. *The Behavior Analyst, 23*, 25–43. [ch 13]

Wilson, T. D., & Lassiter, G. D. (1982). Increasing intrinsic interest with superfluous extrinsic constraints. *Journal of Personality and Social Psychology, 42*, 811–819. [ch 14]

Winett, R. A., & Winkler, R. C. (1972). Current behavior modification in the classroom: Be still, be quiet, be docile. *Journal of Applied Behavior Analysis, 5*, 499–504. [ch 7]

Winner, E. (1979). New names for old things: The emergence of metaphoric language. *Journal of Child Language, 6*, 469–491. [ch 22]

Winograd, T. (1980). What does it mean to understand language? *Cognitive Science, 4*, 209–241. [ch 25, 29]

Wixted, J. T. (1989). The vocabulary of remembering: A review of Kendrick, Rilling, and Denny's Theories of animal memory. *Journal of the Experimental Analysis of Behavior, 52*, 441–450. [ch 27]

Wixted, J. T., & Gaitan, S. C. (2002). Cognitive theories as reinforcement history surrogates: The case of likelihood ratio models of human recognition memory. *Animal Learning and Behavior, 30*, 289–305. [ch 28]

Wolf, M. M., Risley, T. R., & Mees, H. (1964). Application of operant conditioning procedures to the behavior problems of an autistic child. *Behavior Research and Therapy, 1*, 306–312. [ch 7]

Wollen, K. A., Weber, A., & Lowry, D. (1972). Bizarreness versus interaction of mental images as determinants of learning. *Cognitive Psychology, 3*, 518–523. [ch 26]

Wolpe, J. (1958). *Psychotherapy by reciprocal inhibition.* Stanford, CA: Stanford University Press. [ch 17]

Wolpe, J. (1990). *The practice of behavior therapy* (4th ed.). New York: Pergamon. [ch 17]

Woodworth, R. S. (1929). *Psychology* (revised ed.). New York: Holt. [ch 26, 28]

Woodworth, R. S. (1938). *Experimental Psychology.* New York: Holt. [ch 26, 29]

Wright, A. A. (2012). Memory processing. In T. R. Zentall & E. A. Wasserman (Eds.), *The Oxford handbook of comparative cognition* (pp. 239–260). Oxford: Oxford University Press. [ch 12, 13, 26]

Wright, A. A., Cook, R. G., Rivera, J. J., Shyan, M. R., Neiworth, J. J., & Jitsumori, M. (1990). Naming, rehearsal, and interstimulus interval effects in memory processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 1043–1059. [ch 26]

Wright, A. A., Katz, J. S., & Ma, W. J. (2012). How to be proactive about interference: Lessons from animal memory. *Psychological Science, online first*, doi:10.1177/0956797611430096. [ch 12, 27]

Wright, A. A., & Roediger, H. L., III. (2003). Interference processes in monkey auditory list memory. *Psychonomic Bulletin and Review, 10*, 696–702. [ch 12]

Wundt, W. (1900). *Die Sprache.* Leipzig: Engelmann. [ch 25]

Yates, F. A. (1966). *The art of memory.* Chicago: University of Chicago Press. [ch 26]

Yates, F. E. (1986). *Self-organizing systems.* New York: Plenum. [ch 30]

Yerkes, R. M. (1907). *The dancing mouse.* New York: Macmillan. [ch 30]

Yerkes, R. M., & Watson, J. B. (1911). Methods of studying vision in animals. *Behavior Monographs,* 1(2). [ch 2]

Young, F. A. (1958). Studies of pupillary conditioning. *Journal of Experimental Psychology, 55*, 97–110. [ch 17]

Zangwill, O. L. (1972). Remembering revisited. *Quarterly Journal of Experimental Psychology, 24*, 123–138. [ch 26]

Zeigarnik, B. (1927). Das Behalten erledigter und unerledigter Handlungen. *Psychologische Forschung, 9*, 1–85. [ch 26]

Zelazo, P. R., Zelazo, N. A., & Kolb, S. (1972). "Walking" in the newborn. *Science, 176*, 314–315. [ch 4]

Zener, K., & McCurdy, H. G. (1939). Analysis of motivation factors in conditioned behavior. I. Differential effect of change in hunger upon conditioned, unconditioned and spontaneous salivary secretion. *Journal of Psychology, 8*, 321–350. [ch 4, 10]

Zentall, T. R., & Galef, B. G., Jr. (1988). *Social learning.* Hillsdale, NJ: Erlbaum. [ch 19]

Zentall, T. R., & Levine, J. M. (1972). Observational learning and social facilitation in the rat. *Science, 178*, 1220–1221. [ch 19]

Zentall, T. R., & Sherburne, L. M. (1994). Role of differential sample responding in the differential outcomes effect involving delayed matching by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 20*, 390–401. [ch 14]

Zentall, T. R., & Urcuioli, P. J. (1993). Emergent relations in the formation of stimulus classes by pigeons. *Psychological Record, 43*, 795–810. [ch 12]

Zettle, R. D., & Hayes, S. C. (1982). Rule-governed behavior: A potential theoretical framework for cognitive-behavioral therapy. In P. C. Kendall (Ed.), *Advances in cognitive-behavioral research and therapy, Volume I* (pp. 73–118). New York: Academic Press. [ch 23]

Zimmerman, J., Hanford, P. H., & Brown, W. (1967). Effects of conditioned reinforcement frequency in an intermittent free-feeding situation. *Journal of the Experimental Analysis of Behavior, 10*, 331–340. [ch 16]

ACKNOWLEDGMENTS

We thank the following sources and the authors for granting permission to use material from copyrighted works. Full citations appear in the References section.

Copyright by and with permission of the Academic Press: Fig. 27-3, Melton (1963); Fig. 29-1, Thomas (1974).

Copyright by and with permission of the American Association for the Advancement of Science: Fig. 28-1, Shepard & Metzler (1971).

Copyright by the American Psychological Association: Fig. 7-3, Camp, Raymond & Church (1967); Fig. 18-3, Rescorla (1968).

Copyright by and adapted with permission of the Society for the Experimental Analysis of Behavior: Fig. 7-4, Holz & Azrin (1961); Fig. 11-3, from Reynolds (1961a); Fig. 15-3 Catania & Reynolds (1968); Fig. 15-12, Catania, Matthews, Silverman & Yohalem (1977); Fig. 15-7, Dews (1962); Fig. 18-1, Geller (1960); Fig. 18-2, Blackman (1968); Fig. 23-2, Catania, Shimoff, & Matthews (1989).

Index

[Page numbers with an asterisk (*) identify etymologies.]

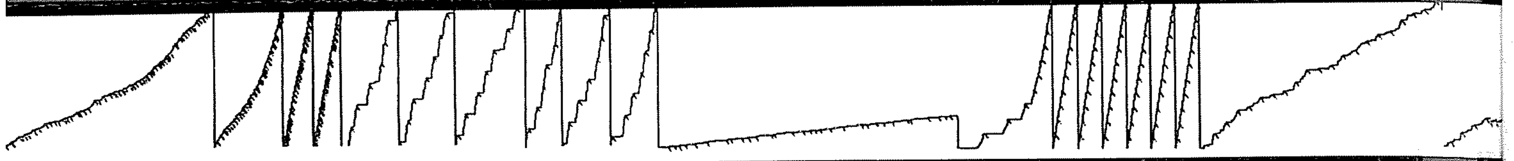

A. CHARLES CATANIA IS PROFESSOR EMERITUS IN PSYCHOLOGY AT UMBC. HE RECEIVED HIS UNDERGRADUATE DEGREE FROM COLUMBIA COLLEGE AND HIS PH.D. FROM HARVARD UNIVERSITY. HE SERVED AS EDITOR OF *The Journal of the Experimental Analysis of Behavior.* WITH VICTOR G. LATIES HE CO-EDITED B. F. SKINNER'S Cumulative Record (for the B. F. Skinner Foundation), WITH PHILIP N. HINELINE Variations and Selections (for the Society for the Experimental Analysis of Behavior), AND WITH STEVAN HARNAD The Selection of Behavior (*for Behavioral and Brain Sciences* and Cambridge University Press). HE IS PAST PRESIDENT OF THE ASSOCIATION FOR BEHAVIOR ANALYSIS AND DIVISION 25 OF THE AMERICAN PSYCHOLOGICAL ASSOCIATION.

THE SLOAN CENTURY SERIES IN BEHAVIOR ANALYSIS IS PUBLISHED IN COLLABORATOIN WITH THE CAMBRIDGE CENTER FOR BEHAVIORAL STUDIES. FOR INFORMATION ABOUT SLOAN PUBLISHING VISIT *sloanpublishing.com*. FOR INFORMATION ABOUT THE CAMBRIDGE CENTER VISIT *www.behavior.org*

CAMBRIDGE CENTER
FOR
BEHAVIORAL STUDIES

SLOAN
PUBLISHING

Sloan Publishing
220 Maple Road
Cornwall-on-Hudson, NY 12520
To order: (845) 534-4994
http://www.sloanpublishing.com

ISBN: 978-1-59738-023-2
9 781597 380232