using a roundabout path. A familiar sign, such as the sight of the entrance to the blocked path, has previously been the signal to enter the path. In this case, it acts as a signal for a new behavior, the taking of the roundabout path. The shortcuts taken as evidence for place learning were used by Tolman as evidence for inferential behavior.

Insight The sudden solution of a problem in the apparent absence of much prior practice with very similar problems.

Latent learning Learning without reward. Rats wandering through a maze still learn something about it, even when there is no reward in the goal box. When the goal box is then baited, the learning previously gained is shown in more rapid learning to reach the goal box than is true for inexperienced rats.

Means-ends capacities "The innate (and acquired) capacities whereby a given organism or species is capable of having commerce-with and expecting means-end-relations." The latter refers to the sensitivity of the organism to direction, distance, similarity, multiple trackness, and so on.

Means-ends readiness Abbreviated to MER. "It is equivalent to a judgement that commerce-with such and such a type of means object should lead on by such and such direction-distance relations to some instance of the given demanded type of goal-object."

Mnemonization One of the three moods of the sign-Gestalt expectation. In this case, a sign is present, but the goal object is not.

Molar behavior "Any organic activity the occurrence of which can be characterized as docile relative to its consequences." To be docile, a behavior must be teachable, or improvable as a function of its consequences. This amounts to goal-directed behavior.

Molecular behavior "A conception of behavior which stresses its underlying physical and physiological character." Molecular views tend to emphasize natural elementary units of behavior, such as the reflex, conditioned reflex, or habit.

Motor patterns A principle of learning added to Tolman's theory in 1949. This consists of Guthrie's theory of contiguity learning as an account for the way in which cognitions are translated to action.

Perception One of the three moods of the sign-Gestalt expectation. In this case, all relevant stimuli (sign, goal object, and so on) are present.

Place learning The learning of the location of objects in space. This is the chief form of learning in a theory that posits cognitive maps.

Primacy A perhaps enduring effect produced by the first exposure to a new situation. For example, one's first experiences with a maze (or a bicycle or a spider) may have lasting effects even though more recent experiences were quite different.

Purpose "A demand to get to or from a given type of goal-object." Such behavior shows persistence, as in continued trial and error, and docility, or improvement with practice.

Recency The influence on behavior produced by the most recent experience in a situation. For example, the repetition of an error made on the preceding trial could be viewed as a recency effect.

Regression The often pathological return to earlier modes of behavior. A middle-aged person may begin to think and act as he or she did as an adolescent or even as an infant.

Response equivalence The successful performing of learned behavior even when the specific actions required to carry out the act are greatly altered. A rat which had learned to run a maze may as skillfully swim to the goal box when the maze is flooded.

Sign One of the three parts of the sign-gestalt, along with the means-ends readiness and the signified (goal) object. A colored panel that signals impending food is a sign.

Signified sign-Gestalt expectation A sign-Gestalt expectation that consists of sign, means-end readiness, and goal object, when a demand for the goal object is present and the goal object is thus signified.

Spatial orientation Another term for place learning.

RECOMMENDED READINGS

MacCorquodale, K., & Meehl, P. E. (1954). Edward C. Tolman. In W. K. Estes, *et al.* (Eds.), *Modern learning theory*, 177–266.

This is a classic analysis of Tolman's theory, from the point of view of writers more sympathetic to Hull's theory.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review, 56*, 144–155.

This is the best single source if you are interested in the general principles of Tolman's theory and in the research that he presented in support of it.

Chapter 8

BURRHUS F. SKINNER: RADICAL BEHAVIORISM

BIOGRAPHY

Burrhus F. (Fred) Skinner was born in Susquehanna, Pennsylvania, in 1904 and received his undergraduate degree in English from Hamilton College (see Figures 8.1 and 8.2). His early ambition was to be a writer, and he spent some time after graduation in a workshop in his parents' home, attempting to write the "great American novel." He finally realized that an author needs more than a knowledge of grammar and a good writing style; an author must also have something to say! Since he had nothing to say, he considered writing to be a waste of time. To remedy this situation, he applied to graduate school at Harvard, where he received a doctorate in psychology in 1931. During the years at Harvard he found something to say and he has been saying it in various ways during the past half century. After graduation, he remained for several years as a research fellow at Harvard and then joined the faculty at Minnesota, where he remained until 1945. He was then chairman at Indiana for three years, after which he returned to Harvard, where he remained until his death in August 1990.

Skinner was never an S-R theorist, as were Thorndike, Guthrie, and Hull, although he is often portrayed as one. The only other major theorist with whom he felt any kinship was Tolman, since he shared Tolman's concern for molar, rather than molecular, behavior. His best work was probably the strategy he proposed in the 1930s for the discovery of order in behavior. This was an insightful solution of the problem of the appropriate units into which behavior should be divided and thus provided an alternative to the molecular/molar controversy represented in the views of Hull and Tolman.

Skinner's view of operant conditioning and its application, presented in popular books, has made him familiar to the general public. *Operant conditioning* is the process whereby a particular class of response is shown to be more frequent as a function of the consequences it produces. In addition, the *Journal of the Experimental Analysis of Behavior* was founded in 1958 as an outlet for operant research. With the loss in popularity of Hull's theory, Skinner became the representative of behaviorism, in the opinions of both academicians and laypersons. His views are often confused with those of Hull. However, these two theorists are as different as night and day, as a reading of Skinner's major publications will show. His most influential publications include:

The Behavior of Organisms, 1938
Science and Human Behavior, 1953
Schedules of Reinforcement, 1957 (with Ferster)
"Behaviorism at Fifty," 1963

The following are two of his best-known popular books:

Walden Two, 1948
Beyond Freedom and Dignity, 1971

INTRODUCTION

Section 1 examines Skinner's significance at the time his theory was introduced, as well as his significance today. Unlike most psychologists, he believed that *theory*, in the usual sense, is not only unnecessary, but harmful. His analysis of the reflex led to a conception of the behavioral unit that was far more useful and sophisticated than was the S-R habit of Hull and Guthrie and the molar definition of behavior proposed by Tolman. His emphasis on contingencies promised a fruitful alternative to the intervening variables used by Hull and Tolman, and it was clear very early that his views were far more useful in application than were the views of most of his rivals and critics. Unlike Hull and Tolman, Skinner held (since 1945) that private experience is an important part of the phenomena our theories must explain.

Section 2 presents the basics of Skinner's theory, beginning with his arguments *against* the usefulness of theories! We will then discuss his early analysis of the reflex and his strategy for discovering order in behavior and experience. This will

FIGURE 8.1 Burrhus F. Skinner, 1904–1990. *Photo courtesy of Random House*

be followed by a discussion of his rationale for the operant/respondent distinction, or the basic distinction between classical conditioning and instrumental learning. This leads to a consideration of the empirical law of effect, which we will see differs greatly from the conceptions of Thorndike and Hull. We will then discuss Skinner's views on stimulus control, secondary reinforcement, punishment, drive, and emotion. Next, we will examine schedules of reinforcement, considering their applications and reasons for studying them. Finally, we will examine Skinner's views on the relation between heredity and environment and his treatment of private experience.

In Section 3, we will consider criticisms of Skinner's position. The most important issues (and the most relevant to his theory) are the role

FIGURE 8.2 Skinner in the laboratory at Harvard, 1930. *Photo courtesy of Bettman Archive*

of species-specific behavior (instinct) and the place of cognition, particularly in behavior therapy. Section 4 deals with other aspects of Skinner's theory. In this section we will consider behavior therapy, performing animals, and pigeon simulations of cognitive activity.

1. THE SIGNIFICANCE OF SKINNER'S WORK

When John B. Watson "founded" behaviorism in 1913, he stressed that activity is the basic subject matter of psychology and that our goal should be to discover the factors that influence that activity. As we saw in Chapter 4, Watson's "theory" was almost incidental; what was important was the answer to the question: Given a stimulus, what may we expect as a response; given a response, what may we discover to be the stimulus (Watson, 1930)? Recall that by *stimulus* Watson sometimes meant such things as a flash of light and at other times such things as an education to become an architect. By *response* he meant anything from an eyeblink to the building of a bridge.

But as behaviorism grew more popular, attention turned from the observation of the world and its effects on us to the postulating of internal mechanisms causing our behavior. The dozens of intervening variables proposed by Hull and the cognitive maps and expectancies of Tolman were substitutes for the world and for behavior. Skinner advocated a return to the basic strategy of Watson, which meant abandoning intervening variables and the theories that proposed them.

Instead, he proposed the intensive study of behavior itself. If we can discover the principles that influence the occurrence of a simple behavior, such as bar pressing by a rat, we may find that we can apply them to behavior in general, including human behavior and experience.

We find that simple behaviors are influenced by their consequences and that a schedule of reinforcement that keeps a rat steadily pressing a bar will also keep a gambler operating a slot machine. In the 1950s it seemed that the study of reinforcement schedules would answer many important questions concerning human behavior, and the success of Skinner's methods in commercial animal training, education, and mental health seemed to lend strong support for this possibility.

Skinner felt that an analysis of *contingencies*, or the relations between stimuli, behaviors, and consequences, would eventually account for attending, remembering, learning, perception, dispositions, traits, and any other activity we might consider. To date, some progress has been made and many hold Skinner wholly responsible for advances in programmed instruction and behavior therapy. Although this is an exaggeration, his influence in these areas is undeniable. Let us examine the basics of Skinner's "theory" and discover just what he means by contingencies and see what their study can do for us.

2. BASICS OF SKINNER'S THEORY

Why Theories Are Unnecessary

When Skinner argued against the use of theories (1950), he puzzled a great many of his colleagues and students, and this puzzlement has continued. In the view of many people, science *is* the development of theories; do we not praise Newton's, Einstein's, and Darwin's theories? What can we do if we abandon theories? Even if we simply make observations and classify organisms we seem to require a theory that specifies what goes in what class!

But Skinner was really arguing against a specific definition of *theory*, which was cast in terms other than those referring to the phenomena the theory was intended to explain. For example, when we train a child to say "bunny" when we present a rabbit, we might interpret the learning involved as the formation of an association between the sight of the animal and the sound of the word. We might represent the association as a habit ($_SH_R$) and conclude that the connection between sight and naming is strengthened by our praise when the behavior occurs. The child remembers the word and the rabbit, and expects to be praised for properly associating the two. Perhaps the new knowledge is represented as a change in the child's cognitive structure, and the child processes the image of the rabbit, retrieves a memory copy, compares the image and the copy, retrieves the corresponding name, and so on.

One could go on endlessly explaining this simple behavior, attributing it to actualizing tendencies, need for achievement, an innate tendency to imitate or to name everything. But what we actually see is quite simple: In the presence of a specific stimulus (the rabbit), a specific behavior (saying "bunny") is followed by a reinforcer (our praise). Thus, it is an example of a simple and familiar rule popularized by Thorndike and used by Hull. That rule explains the behavior and no further theory is necessary or helpful.

Consider some other so-called explanations that Skinner criticized.

Behavior	*Explanation*
argumentativeness, testiness	aggressiveness, frustration
recalling a list of names	memory
not eating a candy bar	willpower
scoring well on tests	intelligence
working hard for long hours	industriousness, achievement need
salivating to the sound of a bell	associating bell and food
having a cup of coffee	thirst

Such explanatory terms not only are used in theories, they are used by all of us every day—and for good reason. Why would Skinner or anyone else object to their use? The reason is simple; although we often treat such terms as explanations, they are more often only names for the behavior in question. Further, in treating them as explanations, we are creating what Skinner called *unfinished causal sequences* (1963). This is naming-as-explaining; a few examples will illustrate the point.

When someone behaves rudely or testily toward you, you may seek the cause for such behavior by asking someone who knows the person better than you do. You may be told that he or she is a nasty, aggressive person and that is why he or she acted in that way. For practical purposes, that may be all you want to know—that the person typically behaves in this way and that you can take care in the future to avoid the person. But all you learned is that the names you already gave to the behavior—rudeness and aggressiveness—are appropriate names for the person's behavior much of the time. You by no means have an explanation for that behavior.

The real explanation for that behavior lies in the history of the individual. Perhaps the manner in which he or she was raised, a history of failure, recent disappointments, a chronic toothache, or a thousand other factors are responsible. It may be the case that aggressive behavior in the past has been rewarded in the form of deference or attention from others. You may not care what the real cause of the rudeness is, but merely naming it as aggressiveness or nastiness is no explanation.

Similarly, when we see someone scan a list of 30 words and recite it back in perfect order, we attribute it to a good memory. Is that an explanation, or is it another case of mere naming? Is memory the cause or do we have another unfinished causal sequence? We call such behavior a feat of memory; but how do we explain memory? Is it possible that it is caused by past practice with such lists? Was the individual using a special mnemonic device?

We call someone who works long and hard industrious; what past influences have led to hard work or industriousness? What past training led to the high test scores that determine what we call intelligence? What past experiences put a person in a position where he or she can turn down a candy bar? Does willpower add anything but a name for such behavior? When I drink a cup of coffee, do I do so because I am thirsty? Do I want to wake up? Or do I drink it because the president has invited me to coffee at the White House? Motives are never explanations; at best, they are labels.

By now, you should get the point; explanations should refer to the activity we wish to explain and to the factors that influence the activity. As Skinner put it, when we talk about "increasing the self esteem of the poor and reducing their frustration," we are saying nothing about how we will do it and how this will change their behavior (1953). On the other hand, suppose we say that we will improve their living conditions, get them jobs that have high social status and promise for advancement, and ensure their safety and health. If we do these things for the poor and they tell us they are happier, then we are going in the right direction.

To sum up, if we discover that practice learning lists of words leads to improvement in the ability to read and then recite lists, do we aid matters to attribute the change to an improved memory? Skinner felt that our knowledge of the conditions that produced the improvement is the only explanation necessary. Is it helpful to say that practice influences memory, which is then responsible for improved performance? By the same token, the many intervening variables of Hull and Tolman are really not necessary.

The Behavioral Unit

Skinner published two papers (1931, 1935) outlining his strategy for the discovery of order in behavior. Together, they constitute a brilliant piece of work. They show that the unit of behavior need not be molecular, as Hull proposed, and it need not be molar, as Tolman proposed. Skinner's rejection of the traditional molecular and molar alternatives and his proposal of a different kind of molar unit has often been misunderstood. In fact, half a century after they were written, Skinner himself seemed to misunderstand them! Malone (1987) and Staddon (1967) pointed out that he tended to emphasize molecular interpretations, despite his earlier arguments against them. His inconsistency concerns his followers and provides fuel for his critics. The essential points of these early and classic papers by Skinner are presented here.

The Concept of the Reflex Skinner's 1931 paper, "The Concept of the Reflex in the Description of Behavior," traces the history of the *reflex* from Descartes's time to the present, showing that some superfluous properties associated with the reflex have been dropped over the centuries. For example, in his original treatment of the reflex, Descartes proposed that animal spirits were responsible for reflex action. The spirits filled muscles, causing them to swell and thus decrease their length. This, in turn, produced the movement of a limb.

Once it was discovered that muscles do not "expand when they contract," theorists dispensed with the concept of animal spirits, but other superfluous properties were left. Even as late as the nineteenth century, there was talk of a vis nervosa, or "life force of the nerves." And the great physiologist Friedrich Goltz spoke of a "soul of the spinal cord" that governed reflex action! During this century, the acceptance of the neuron theory and evidence for the synapse left the reflex arc of neurons as the only essential basis for reflex action. And this seems as it should be; is it not true that the reflex arc *is* the reflex?

Skinner argued that even the arc itself was superfluous. It is not an essential part of the definition of the reflex and, after all, reflexes had been discovered many years before the first reflex arc was traced. When we get down to it, the reflex is really a conceptual expression, referring to a certain correlation between stimuli and responses. There are rules for the discovery of reflexes that tell us when we have isolated one and the discovery of a reflex is an explanation for part of our behavior.

For example, light causes the pupils of the eyes to constrict; when this (pupillary) reflex was discovered, that small part of behavior was explained. It was not explained because a reflex arc of specific neurons was shown to be responsible and Skinner never suggested that behavior was a collection of simple reflexes. Consideration of simple physiological reflexes is simply the starting point for the discovery of higher forms of order.

Discovering Reflexes How do we isolate even a simple physiological reflex? Ordinarily, the procedure is to stimulate organisms in various ways and to look for effects on behavior. I tickle a dog on its side and it scratches with a hind foot. Have I discovered the scratch reflex or is the behavior a combination of reflexes? I know that it is a true reflex if it obeys the laws that govern reflex behavior in general; these laws may be arranged in two groups, which Skinner called the *static laws* and the *dynamic laws* of the reflex. For example:

Static Laws	*Dynamic Laws*
threshold	fatigue
afterdischarge	summation of subliminals
magnitude	refractory phase

Thus, if I find that a certain magnitude of stimulation is required before scratching occurs, the behavior obeys the law of the threshold. If scratching continues for a while after stimulation ends, and if stronger stimulation produces more vigorous scratching, the behavior follows the laws of afterdischarge and the relation of magnitude of response to magnitude of stimulation. I probably have found a reflex.

If I stimulate repeatedly, I find evidence for behavior conforming to the dynamic laws. Repeated stimulation leads to decreased responding (that is, fatigue) and, if I softly but repeatedly tickle the dog, the individual weak stimulations seem to add up and finally produce a response (summation of subliminals). Finally, a refractory phase is evidenced in the fact that after I elicit a response, I will be unable to evoke another for a brief time thereafter. My reflex *is* a reflex, since it obeys both the static and the dynamic laws governing reflexes. If I add up all of the reflexes that I can evoke from a dog or from a human, I find that I can account for some of its behavior. I can explain the action of the digestive system and the rest of the viscera, as well as the parts of behavior that contribute to walking and maintaining balance. But I am really interested in other things; how do I deal with more interesting behavior?

Larger Units I continue the same strategy for larger units but change the overall level of my analysis. Reflexes are correlations between stimulation and responding that obey static and dynamic laws, and I can apply the same method to the discovery of Pavlovian conditioned reflexes. I will find the same static and dynamic laws that I found to be true of the real reflex, and I can add two important dynamic laws—the law of conditioning and the law of extinction of Type S (Skinner's term for classical conditioning). That is, when I have found a behavior that I believe to be controlled by a CS, I can strengthen the reaction by pairing the CS with a UCS, and I can weaken the reaction by repeatedly presenting the CS alone (extinguishing the CR).

This allows me to explain much more of the behavior of my subject; I can list the reflexes that appear at birth or shortly thereafter and then add the acquired reflexes that are elicited by one or another CS. For example, if I am attempting to account for an infant's crying, I may search for stimuli that I have discovered elicit reflexive crying. Is the infant reacting to pain produced by a diaper pin? Is the infant cold or wet? If I can find no stimulus that might elicit reflexive crying, I seek a CS that may be eliciting conditioned reflex crying. Perhaps a dog's bark has evoked reflexive crying in the past and the sight of a dog has thereby become a CS, producing conditioned reflex crying.

Earlier theorists, such as Watson, had left matters there. The analysis of all behavior is restricted to the discovery of stimuli that elicit reflexes and those that have become CSs. (Skinner called behaviors that are elicited by specific stimuli *respondents*. Classical conditioning is therefore respondent conditioning.) But Skinner felt that such a strategy was hopeless; even if there are stimuli responsible for every aspect of behavior and experience, they are too numerous and too difficult to discover to make such a strategy plausible. What we should do is use the same method we used to discover reflexes and conditioned reflexes and move our level of analysis higher, to the behavior that seems to occur spontaneously and for which no obvious eliciting stimuli can be found. If we can show that such behavior, properly defined, is also reflex (in the sense of being orderly), then we will have taken a great stride forward.

However, if the reflex is defined in terms of obeying static and dynamic laws, we have an immediate problem. Static laws are those characteristics of the reflex that are observable with single elicitations of the reflex (for example, the threshold), and dynamic laws refer to characteristics of the reflex when it is elicited repeatedly (for example, fatigue). Spontaneous behavior, by definition, is not elicited; it appears in the absence of any eliciting stimulus, as far as we can tell. Of what use are the static and dynamic laws in identifying the reflex in this case?

As a matter of fact, the static laws *are* useless here; we cannot elicit spontaneous behavior, so there is no way to examine it when it is elicited once. But we may yet appeal to the dynamic laws of the reflex, since such behavior occurs repeatedly on its own, even though we cannot identify an eliciting stimulus. To make a long story short, we find that such behavior is influenced by the consequences it produces; we can influence the rate of occurrence of such behavior by manipulating such consequences.

Thus, the law of conditioning of Type R (that is, operant behavior and the law of effect) and the law of extinction of Type R may be used as dynamic laws to help us identify the unit of spontaneous behavior, or the operant. As long as orderly changes occur as we vary the consequences of such behavior (including the use of various schedules of reinforcement), we may be certain that we have extended the concept of the reflex to the level of freely occurring behavior.

Let us return to our example of the crying infant. It may well be that we can find no eliciting stimulus; we must then assume that the crying is neither reflex nor conditioned reflex. But, when we vary the consequences produced by crying, the frequency of crying changes. Perhaps it is our concern and attention that acts as a reinforcer for operant crying. If we ignore bouts of crying and find that it decreases in frequency, then we know that it is operant crying.

Practical Application The same strategy may be applied to the analysis of any behavior, whether it be aggression, crying, reading, anxiety attacks,

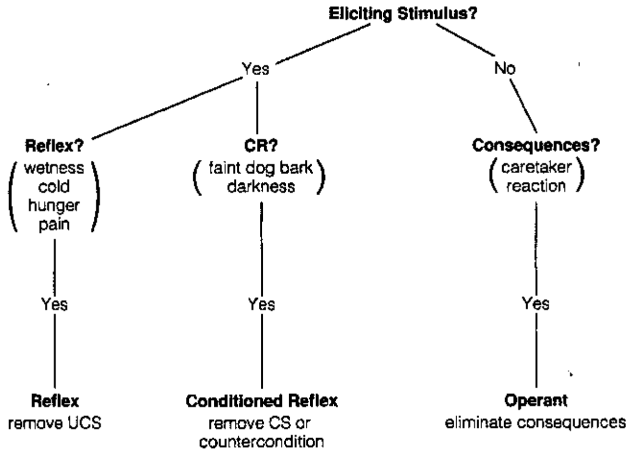

FIGURE 8.3 Identifying behavioral units: causes for an infant's crying.

hard work, and so on. We identify reflexes (show order) in the following manner:

1. We look for an eliciting stimulus. If it is a stimulus that elicits traditional reflex behavior (for example, painful withdrawal of a limb), we influence the behavior by removing or altering the eliciting stimulus. If it is a CR, we find the CS and extinguish responding to it or use counterconditioning.
2. If we can find no eliciting stimulus, we look for consequences that may be maintaining the behavior. If we find them, and if we alter the behavior when we alter the consequences, then we know that we are dealing with an operant (or operants).

Figure 8.3 presents a diagram of this method of identifying units of behavior.

For example, if we wish to alter the speed of handwriting in a slow writer, we find no eliciting stimulus. We may ask the writer to copy written prose and find that altering the consequences of such behavior alters its rate. But we find that when we reinforce this behavior according to different schedules of reinforcement, its rate does not change as does the rate of operants thus reinforced. Either handwriting is not a true operant or our definition of handwriting is inappropriate; we have not defined it in the proper units. Suppose that we reinforced the handwriting of so many pages per day of original, good-quality prose, as occurs in many business situations. Perhaps then we will find that writing speed behaves as do other operants when we reinforce it according to different schedules.

The discovery of a reflex, conditioned reflex, or operant does not guarantee that what we are specifying as the stimulus (in the first two cases) or the response is appropriate. Showing that a behavior is elicited by stimuli or influenced by consequences is only the beginning.

The Generic Nature of Stimulus and Response

The discovery of a reflex is not complete with the isolation of a unit that obeys the static and dynamic laws of the reflex. Whether we speak of a flexion reflex, as did Skinner (1935), or a more complex phenomenon, such as aggression, it is necessary to refine our definition of stimulus and response if we are to have a useful unit of behavior. The unit that we discover initially almost certainly includes irrelevant aspects of stimulation and responding. It also probably fails to emphasize aspects that are essential.

Let us consider an example. Rather than the flexion reflex, which Skinner used (1935), we will use aggression in our example. Skinner would not have approved of such an example in 1935, since aggressive acts must include instances of all three varieties of reflex behavior—the biological reflex, the conditioned reflex, and the operant. But the use of such an example will make it easier to illustrate Skinner's strategy for discovering useful behavioral units and will serve as an example of the way in which the strategy may be extended. Let us begin as the popular press often does and treat all behavior that could be viewed as aggression as a single entity. Thus, we must include aggressive play by children, the killing of prey by predators, angry words produced by a headache sufferer, random killings done by a murderer, the aggressiveness of a businessman or a football player, international conflict, sibling rivalry, and whatever else could be called aggression.

In searching out eliciting stimuli, we find that some of these forms of aggression are produced by obvious stimuli, as in the cases of the headache and predation. In other cases, we find that the consequences of the aggressive activity influence it, as in the case of children's aggressive play and in the aggressiveness of the businessman and football player, as well as in international conflict and sibling rivalry. Random killings could fall into either case.

We find that, although we have applied a common class name—aggression—to all of these cases, we really have a number of fundamentally different cases, which is why I chose the example. Let us narrow down our definition of aggression so that we can be more certain that one case of aggression is equivalent to another case. For example, we might consider the aggressiveness of the football player, the businessman, and the children. We soon find that all of them are dependent upon the consequences they produce and therefore qualify as operants. Our only problem is to identify instances of what we are calling aggression in the three cases. We could begin by including all actions in which damage or the threat of damage to people or property is present and note the effect of the consequences produced; this is what Skinner called the extreme generic position.

But we find that this doesn't serve us very well. Some aggressive acts, defined in this way, are clearly atypical: The businessman inadvertently ruins a friend, the football player breaks a teammate's leg, and the child accidentally pushes a playmate into a lake. These consequences do not increase the occurrence of aggression in general. Our definition, then, is too general. Let us go to the other extreme and call each individual act that involves damage to others the unit of aggression. The child pushing an individual playmate is a different form of aggression from the kicking of the playmate; the football player aggressively tackling a particular opponent is fundamentally different from the same player blocking an opponent. But clearly this is not a useful strategy to take; it limits us to predicting and influencing only specifically defined instances, even though we feel that many instances have something in common. So how do we find what it is they have in common?

The answer lies somewhere between the first and the second alternative, of course. What we must do is observe instances of aggression and refine our definition, beginning with the extreme generic case. We will find that at some point order will appear, as the frequency of the behavior we are counting changes in a lawful way with changes in the consequences it produces. What will this definition include? It depends upon the behavior we are concerned with.

No one can safely predict in advance exactly what behaviors will hang together to form our definition of one or another kind of aggressive behavior. We may find commonalities among some behaviors of the child, the businessman, and the football player, but we may not. If we were to find that the three kinds of aggressive behavior, when defined in a certain way (for example, producing injury to other people belonging to the same social group and not resulting in retaliation) vary in frequency as we vary the consequences (in money, praise, and so on), then we could say that we are dealing with essentially the same unit of behavior.

More likely, we will be unable to do this and we will have to be satisfied with concentrating on one of them. For example, we observe aggression among children, specifying as clearly as we can what we mean by aggression and noting the effects of praise and other payoffs for such behavior. At some point our definition will be refined so that we can predict quite precisely what will happen to the whole class of behaviors when we vary their consequences. When this occurs, we may be sure that we have adequately described what we mean by aggression of this sort; we have discovered a reflex, in Skinner's sense of the term. We have a unit of behavior that varies in an orderly way as we vary its consequences or when other variables (for example, lack of sleep) produce changes in the overall level of behavior.

Our unit, therefore, is not determined in advance. We do not treat behavior and experience as a set of habits, or as goal-directed behavior in general, as was done by Hull and Tolman, respectively. We seek out eliciting stimuli or consequences that influence the behavior in question and then narrow our definition of that behavior until we discover the class of behaviors that vary together in an orderly way.

In practice, the discovery of an operant class may reveal some behaviors as members we could not have foreseen. Suppose we alter the consequences of specific aggressive behaviors on the playground by removing the consequences that normally maintain such behavior. We pay no attention to the aggressive behaviors and we arrange that other children do likewise. As the pushes and hitting that we have identified as aggression on the playground decrease in frequency, we may find that other behaviors decrease in frequency as well, such as surliness toward adults and giggling in class. The class of behaviors we termed *aggression* may belong to a larger class, including disruptive behavior in class; whether it does or not depends upon what behaviors covary, since that is the definition of the operant class and the criterion for the demonstration of order in behavior.

The Stimulus and Response Classes

In defining the reflex, the conditioned reflex, and the operant, we always refer to correlations between aspects of the world and behavior and to the things that are correlated as classes. The reflex and the conditioned reflex are correlated classes of stimuli and responses, whereas the *operant* is a class of behaviors that depends upon similar consequences. As we saw earlier, the membership of all of these classes may not be obvious; we discover those behaviors that constitute the class.

In the simplest case, the membership of the *response class* may be more or less obvious. We may train a rat to press a bar, with reinforcement presented occasionally in the form of a food pellet. Once the rat is reliably pressing, we find (as Skinner did in 1938) that the operant class we have strengthened is defined as any behavior that depresses the bar enough for us to count it as a press. It matters not whether the rat used its left or right paw, a hind foot, or even its snout. Any act that depresses the bar counts and, should we increase or decrease the frequency of reinforcement, we find that this whole class of behaviors changes in frequency.

We may refine the operant class if we wish by altering the requirements for reinforcement. If we require that the bar be pressed only with the right forepaw, we may find that the other members of the operant class decrease in frequency and that the new class consists of the various ways in which the bar may be pressed with the right forepaw. We have created a new operant class, using a method often used by animal trainers, as we will see later.

The concept of stimulus and response classes was always central in Skinner's thinking, and we should note that this is very different from the ordinary S-R association theories of others, most notably Hull. For such theories, reinforcement involves the strengthening of specific responses, and if the strengthening appears to affect other behaviors it must be due to the generalization of the specific reinforcing effect. Such response generalization is not the same as the concept of a response class, where the response is the class of behaviors that covary.

The confusion between Skinner's definition of the response class and traditional S-R interpretations may be evidenced often in attacks on Skinner, considered to be the foremost behaviorist and often assumed to hold the same S-R position as Hull. Critics such as Noam Chomsky (1959) argue that the learning of language is not the learning of specific associations between objects and words; rather, it involves an innate knowledge of the deep structure of a language. But consider the way in which Skinner's notion of classes applies to the learning of language.

A child learns that the sight of his or her mother is an appropriate time to say "mommy," because when the child says this there is general rejoicing among other family members, who pay loving attention to the child. Saying "mommy" does not happen all at once, of course; we may assume that it is the product of refinement, similar to the rat's pressing of the bar with its right forepaw. Even Chomsky would not question the fact that words are learned, so there is no need to belabor this process. What is it that the child has learned? Has the child learned to say "mommy" in a very specific situation (at home, in its room, with mother present) in a specific tone and volume?

If we watch the child we find that what has been learned is a correlation between a class of stimuli and a class of behaviors. In a variety of

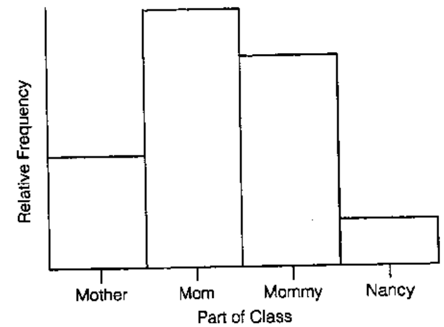

FIGURE 8.4 Stimulus and response classes. The names the child uses to refer to his or her mother form the response class. The names the child hears his or her mother called form the stimulus class.

situations (but not all) we find that the child says "mommy" in a variety of tones and volumes. This set of situations is the *stimulus class* and the child's names for his mother form the response class. Other women may initially be called mommy, and other terms may be added to the class of inflections and volumes in which *mommy* is uttered. With time, the child learns that *mother*, *mom*, and other words produce the same effects that *mommy* produces, so these words become part of the same class.

When the child is older, he or she learns that mommy is called another name, such as *Nancy*, by other people, and this name may become a less frequently used member of the "mommy" class (see Figure 8.4). In a recent newspaper story, a child trying to attract his mother's attention ran through stronger members of the "mommy" class, such as *mother*, *mom*, and so on, but these words produced no effect. At last, the child called *Nancy* and finally caught the mother's attention. In the future, one might expect that this will occur more frequently, especially if it continues to produce effects that *mommy* and *mom* fail to produce.

Since we do not learn specific responses, there is no reason why we cannot learn rules, such as those referring to noun-verb order and other rules of grammar. We will see later that animals can easily learn far more abstract concepts, so that there is no reason to doubt that children can learn more than discrete reactions to concrete and specific stimuli. We will discuss this later; suffice it to say that arguments such as Chomsky's are valid only insofar as they criticize simple S-R associationism, a view that never characterized Skinner.

Summary

In summary, the reflex, conditioned reflex, and operant are all defined in terms of classes of stimuli and responses or responses and consequences. Only through observation may we know the membership of these classes, and we may often be surprised at the identity of some of the members. This interpretation, which is unique to Skinner, is far different from the old associationist behaviorisms, which assumed that a stimulus was a stimulus and a response was a response. It is toward such views that many current criticisms of Skinner are properly directed.

We may note that Skinner's definition of the reflex and identification of stimulus and response classes are actually a refinement of Watson's cruder treatment of stimuli and responses. For Watson, *stimulus* and *response* are molar terms if the phenomena of interest are molar things, such as the kind of education (stimulus) that produces efficient reading (response). In other cases, *S* and *R* refer to molecular things, as when one discovers a cue that produces a facial tic.

The Empirical Law of Effect

A great many psychologists have concerned themselves with the problem of reinforcement: How is it that reinforcers work as they do? We have seen that Thorndike struggled with this and concluded that satisfiers have their effect (or lack of it) depending on the activity of the organism prior to the receipt of a satisfier. Thus, he appealed to a law of readiness.

It is often difficult to determine for what consequences an organism is ready until after the fact, so Hull proposed a more specific (and popular) explanation: All reinforcers work as they do because they either reduce a drive (such as hunger) or they are stimuli that have been paired with the reduction of a drive (such as the sight of food). Both of these explanations are attempts to specify the *theoretical law of effect*, which means they attempt to explain why reinforcers work as they do. Other possible explanations hold that reinforcers work because they provide pleasure, activate reward centers in the brain, increase general sensory stimulation, act to further life, or provide information. We are still far from a clear understanding of why all reinforcers work under all of the conditions they do. Is it really important that we know?

Skinner long argued that it is not important and that efforts to explain the action of all reinforcers are misdirected. For him, the law of effect is an empirical fact; we find that many events under many conditions act to increase the frequencies of many operants, and that is enough. Food, water, attention, and praise reliably reinforce many behaviors. Depending on circumstances and our purposes it is no great difficulty to discover an effective reinforcer and, once we have done that, we may discover what operants are influenced by it. According to this *empirical law of effect*, a reinforcer is defined as that which influences the strength of an operant, and that is the end of it (see Table 8.1). The law of effect is thus a useful rule that allows us to show order in behavior and experience. This, in turn, allows us to predict and control behavior and experience, which is our main aim.

Treated in this way, the law of effect is circular, a fact that has concerned and dismayed many supporters and critics of Skinner. An operant is a class of behaviors that is affected by reinforcers. Reinforcers are things that affect operants. We are left with the operant as a class of behaviors affected by things that affect such classes of behavior! A reinforcer is an event that affects behaviors affected by reinforcers! The law of effect is

TABLE 8.1 Law of Effect

Theoretical law:	Specifies an independent criterion that tells us when a consequence will act as a reinforcer.
	For example: readiness (Thorndike) drive reduction (Hull) brain centers (James Olds)
Empirical law:	Defines reinforcers by their effects.
	operants (behaviors that are sensitive to their consequences) — followed by — reinforcers (consequences that affect behaviors sensitive to such consequences)
	Note the circularity in the definition of operants and reinforcers. In physics, terms such as *force* are difficult to define independently.

circular because we cannot define a reinforcer or an operant independent of one another. Can this be true? And if it is true, can we leave it at that?

It is true and we can leave it at that, as far as Skinner is concerned. It is enough to say that we find specific kinds of correlations between classes of behavior (lever pressing by rats, crying by infants, working by factory workers) and specific consequences (food, attention, money). We seek out such correlations and, having found one, show that much more order in the world. The empirical law of effect may be circular, but we find that there are many other circular laws and that these laws are useful. For example, we may use the equation force (F) = mass (M) × acceleration (A) without being able to define independently *mass*, *force*, or *acceleration*. If $F = MA$, then $M = F \div A$, right? It is true that we can independently define *acceleration* (as change in velocity per unit time), but *mass* and *force* are still as difficult to define as they were for Newton. The great physicist was apologetic about the use of such terms at the beginning of his *Principia* but urged the reader to bear with him, since the use of such terms as *force* would be seen to be justified by its usefulness in his overall system. This is essentially Skinner's position concerning the law of effect.

Positive and Negative Reinforcers We normally think of a reinforcer as increasing the frequency of a behavior because the behavior produces the reinforcer; the receipt of food for bar pressing increases the frequency of pressing, and praise or a high test score increases the studying behavior that produced such consequences. Reinforcers that produce an increase in the frequency of the behavior class that produce them are termed *positive reinforcers*. Other consequences may act as reinforcers because an operant removes or ends them. For example, our withdrawal from a hot stove top is reinforced by the cessation of the discomfort the stove produces, just as a rat's bar press may be reinforced because such action turns off a powerful electric shock. Reinforcers that strengthen the behavior that removes them are termed *negative reinforcers*.

Note that all reinforcers, by definition, produce an increase in the frequency of a behavior class. The distinction between positive and negative reinforcers refers to whether the reinforcer is produced or removed by the behavior in question (see Figure 8.5). There has been confusion over the years about the distinction between negative reinforcement and punishment, although the difference is simple enough. A punisher is a consequence that decreases the frequency of the

Responding	Consequence: Add	Consequence: Remove
Increase	positive reinforcement	negative reinforcement
Decrease	punishment	omission (see Chapter 3)

FIGURE 8.5 Classification of effects of consequences. Note that reinforcers always increase the frequency of behavior. Omission could also be called negative punishment; in this case the occurrence of a behavior prevents a consequence.

behavior that produces it; thus, an angry word from a parent or the delivery of electric shock may decrease the frequency of the misbehavior or the bar press that produced it.

The problem is that negative reinforcement seems to imply the involvement of an aversive event (which is true), and we think of aversive events as punishers. To remember the proper usage of negative reinforcement and punishment, remember that all reinforcers increase the frequency of a behavior class, either because they are produced or terminated by members of that class.

Punishment The decrease in frequency of an operant behavior as a function of the consequences it produces. Because of the results of a few informal experiments, Skinner (e.g., 1938) felt that punishers do not directly affect behavior in the way that reinforcers do. For example, when a rat's bar pressing, which previously had been reinforced with food, produced only a slap on the paws by the lever, bar pressing was temporarily suppressed. However, when the slaps were discontinued and food was again provided for pressing, the rate of bar pressing quickly increased and even exceeded the level present before punishment was begun. This led Skinner to conclude that the effects of punishment were temporary and that any behavior that was suppressed during punishment was "saved up," only to reappear once punishment ended. The suppression of responding during punishment was due, in his opinion, to emotional side effects, rather than to a real decrease in the strength of the punished behavior. In Chapter 9, we will discuss Skinner's evidence for this view, along with later and more convincing evidence provided by Estes (1944). Azrin and his colleagues have shown that punishment does work; Skinner's and Estes's results were not representative.

The Operant/Respondent Distinction

In his early work, Skinner argued that elicited behavior and operant behavior form two fundamental classes and that the processes involved are different in a number of ways. In 1938, Skinner proposed the following classification:

Type S (respondent conditioning)	*Type R (operant conditioning)*
Involves stimuli (CS and UCS)	Involves responses and stimuli (R and reinforcers)
Stimuli elicit responses	Stimuli set the occasion for responding
May begin at zero strength	Must occur with some frequency
Involves a new reflex	Involves intensifying an existing behavior
May have a negative reserve	Never has a negative reserve
Involves the autonomic nervous system	Involves the skeletal nervous system
Deals with involuntary behavior	Deals with voluntary behavior

The first of these distinctions is valid, since it refers to the methods used in classical (respondent) and operant conditioning. In the first case, stimuli (CS and UCS) are presented independent of the behavior of the subject, whereas in the second case reinforcers are presented dependent on the occurrence of a behavior. The remaining distinctions are less certain. We will discuss them in a later section.

Is Conditioning Gradual or Rapid?

Operant conditioning is usually thought to occur gradually, as is the case in respondent conditioning. Each successive reinforcement adds an increment to the strength of the operant, and it may take many reinforcements before maximum strength is reached. Even Guthrie, who emphasized one-trial learning, felt that operant strength increases gradually, since an operant is an act; it is each specific response that contributes to the act that is attached on one trial.

Yet, Skinner (1938) devoted several pages to presenting evidence showing one-trial learning of operants (1938, pp. 69–73). He suggested that a single reinforcement may suffice to raise an operant to its maximum strength and that such an instantaneous change is common, if not actually typical. Skinner suggested that, when maximum effects do not occur on one trial, it may depend on a number of interfering factors. Thus, a sluggish feeder may actually be interposing a delay between response and reinforcer, the form of the response required may be unusual and thus difficult for the subject, other behavior may be reinforced and competing with the behavior in question, the subject may be in an emotional state, and so on.

Generalization and Discrimination

As noted above, operants commonly are influenced by environmental stimuli, which set the occasion for responding. Such stimuli are termed *discriminative stimuli*, or S^D; an operant under such *stimulus control* is a discriminated operant. Thus, a rat may press a lever only when a tone signals that reinforcement is available, and we learn as children to react to discriminative stimuli such as the sight of a policeman or the verbal signals "Watch out" or "Look what I have for you." In all cases, the S^D tells the rat and us that some operant behavior (pressing, watching out, or coming to see what is there for us) will be reinforced.

We generalize to other stimuli, an effect that Skinner always called *induction*, after Sherrington's simultaneous induction. After the rat learns to press when the tone is on, it will press when other tones are on. Similarly, we learn to react to uniformed elevator operators and to boy scouts and girl scouts as we react to policemen. Early in life we may react also to a loud shout "You are correct," as we react to a loud shout that warns us of danger.

Following Thorndike and Hull, Skinner suggested that generalization occurs to the extent that other stimuli are similar to the S^Ds that have signaled consequences for operant behavior. But, unlike Thorndike, who defined similarity as "common elements," and unlike Hull, who defined it in terms of discriminability, Skinner treats similarity as a relation defined by responding.

When we speak of a black square and a gray square being more similar than a black square and a white square, we may refer to the possession of shared elements or differences in discriminability. In other cases, it is harder to specify what we mean by similarity. Is a horse more similar to a cow or to a camel? Is an *X* more similar to a *Y* or a *V*? In such cases, shared elements and discriminability lead to less clear-cut conclusions, so we base our assessment of similarity on the behavior of our subject. If we are dealing with a child, we may tell him or her to say "horse" when we show a picture of a horse, but not otherwise. If the child then says "horse" more often when shown a picture of a cow than when shown a picture of a camel, we may conclude that the cow is judged to be more similar to the horse than is the camel. On the other hand, we may note the similarity of functions of the three animals and judge the camel to be more similar to the horse.

According to Skinner, the reinforcing of an operant in the presence of an S^D leads to generalized or induced responding to similar stimuli and it is only through *differential reinforcement* that discrimination occurs and responding is restricted to the S^D. This simply means that the rat receives food only for presses in the presence of the S^D tone and that we receive verbal confirmation when we say "policeman" only when a policeman is present. When a different tone is on or when we call an elevator operator a policeman, food and verbal confirmation are not forthcoming and the behavior is extinguished in the presence of such stimuli. These stimuli, for which responding is not reinforced, are termed S^{Δ}s. Hence, we learn to respond when an S^D is present and not when an S^{Δ} is present.

That is as far as Skinner goes in explaining generalization and discrimination. Operant conditioning leads to induced responding to stimuli similar to the S^D, and differential reinforcement (reinforcement when S^D is present and not when S^{Δ} is present) leads to the extinction of responding to S^{Δ}. Since the response that is extinguished is an operant, its strength may approach zero but does not go below it, as Pavlov and Hull claimed it could.

Shaping and Chaining

Skinner always viewed operant behavior as a constantly changing thing; we act and the consequences of our action modify our behavior. The child who mumbles requests will find them often unanswered and the scholar whose study habits do not produce the desired consequences will alter them if possible. Since our environment shapes our behavior, the complexity and originality of our actions, both overt and private, depend in large part on the complexity of the environment with which we interact. Our behavior is constantly being shaped, much as a sculptor shapes a piece of clay.

Shaping by Successive Approximations When we purposely shape an operant class, we call the technique used the method of successive approximations or *shaping*. This technique enables us to modify greatly an existing behavior class or create an entirely new behavior. We may apply this method to teach a dog to stand up, to teach a new word to a child, or to teach pigeons to play Ping-Pong, one of the first uses of the method. If we have higher aspirations, we may shape creative behavior in human or beast and even shape patriotic American prisoners of war to readily and frequently inform on their buddies!

The method used in all of these cases is the same. We begin by monitoring the behavior we are interested in (the dog's posture, the child's utterances, the pigeon's behavior toward a Ping-Pong ball) and find something we can use as a reinforcer, be it food, praise, or whatever. In the case of the dog, we may give it a bit of food with our arm raised in front of us and reinforce the dog's moving upward for the food until the sight of us makes the dog raise its head upward. If we do not want the dog following us around with upward-gazing eyes, we may reinforce head raising only when we give the command "stand up." In this way we place the behavior under the control of a discriminative stimulus (S^D).

Once the dog is reliably raising its head, we narrow the range of head raises that we reinforce, thus extinguishing lesser head raisings. The process of extinction typically leads to increased variability in behavior and we find new and higher head raisings occurring, with the dog's forepaws leaving the ground. We now have a new class of behaviors, with the average head raise far higher than was the case a few minutes ago and we again restrict the range that we reinforce. This leads to extinction of the lower raises in this behavior class, which leads to increased variability of behavior and new and higher leaps that we can selectively reinforce.

What we could have in the end (if we are tall enough or the dog is small enough) is a dog that leaps high in the air on command. This may be a behavior never before shown by the dog; we have created it through the method of successive approximations. The same method will work with the child, the pigeon, or whomever. In the case

of creativity, we need only specify what we mean by *creative* and shape behavior in that direction. In the famous case of the creative porpoise (Pryor, Haag, & O'Reilly, 1969), two porpoises' novel behaviors were reinforced. This was an unusual demonstration, since it was not specific responses that were shaped, but the class of behavior that could best be called doing something new. (That *is* what we mean by creativity, is it not?) After a period in which only new behaviors were reinforced, the porpoises exhausted the behaviors commonly shown by their species (for example, porpoising, inverted swimming, and tail slap) and even behaviors infrequently seen (for example, beaching, back porpoise, spitting, and tail walk). Then they began showing heretofore unheard of behaviors (for example, back flip, spinning, inverted leap, and tail wave). As the members of the class of common behaviors were extinguished, overall behavior became more variable and new behaviors occurred that could then be reinforced. If the shaping of creative behavior is possible in such animals, should we be surprised when we find that we can do the same with children?

In the last example, psychiatrists who interviewed American POWs after the Korean War found that the Chinese had used ingenious shaping methods in their prison camps, which probably contributed to the large number of our soldiers who remained in China after the war. Their simple method consisted of reinforcing soldiers' "confessing" to minor misdeeds of their fellows, with the assurance that no punishment would follow. For example, if a soldier told the guards that his buddy took more than his share of potatoes at dinner, then both the soldier and his buddy would be rewarded with cigarettes and privileges. Prisoners soon arranged to regularly inform on one another, thinking the whole thing a joke and evidence of the stupidity of their captors, at least at first. As informing became a way of life, the possibility of escape decreased immensely. Who would confide his plans to his fellows, whose help he would need, after he had informed on them and they on he dozens and dozens of times (Whaley & Malott, 1971)?

Establishing a Chain Once a behavior is shaped, it may be made a member of a chain of behaviors, as is the case with many of the behaviors shown by trained animals at amusement parks and in circuses. This procedure is called *chaining*. For example, suppose we wish to duplicate Skinner's "tour de force," as he called it (1937), and train a rat to pull a string that releases a marble from a rack. The rat then picks up the marble with its forepaws and carries it to a tube that sticks up two inches from the floor. The rat raises the marble to the top of the tube and drops it in, whereupon the rat receives a food pellet. How would you go about training a rat to do such a thing? A rat's normal lifetime would probably be too short a time for you to accomplish such a feat, unless you know what you are doing.

Before considering how we would do it, consider a simpler task. Suppose we wish to train a rat to press a bar four times when a green light goes on, which turns the light to red, and four further presses to turn the light to blue. After the light turns to blue, four more presses are followed by receipt of a food pellet. How would we train a rat to do this?

The trick is to begin with the final member of the chain, which is more obvious here; the first thing to do is to train the rat to press the bar while the blue light is on. We would first shape this behavior, reinforcing movements of the rat as it orients toward the bar, and then only when it is progressively closer to the bar. If we can keep it close enough to the bar by skillfully delivering pellets, the rat is bound to touch the bar sooner or later and we can require touches and then presses before food is delivered. Without too much difficulty we can thus train the rat to bar press.

Once we have established pressing, it is easy to increase to four the number of presses that bring food; we now have a rat busily pressing four times (or many more if we wish) for each pellet of food. Note that the rat is pressing while the blue light is on. What if we replace the blue light with red? Chances are the rat will still press, and, if it does not, we may shape its pressing as we did before; but this time we find that if we momentarily change the light from red to blue, the shaping effect is similar to that produced by food! The blue light, which was on in the past when pressing brought food, now acts as a reinforcer itself; it has become an S^D, signaling food, and a *conditioned reinforcer*.

When the rat presses the bar and the light is red, we change the light to blue and the rat goes through its four presses, which bring it food. Then we require two, three, and finally four responses when the light is red to change it to blue. Each time we allow the rat to go through its four-press sequence, which ends in food. Now we have a rat that presses four times in red, followed by four presses in blue, which brings food.

Finally, we begin with a green light, and when the rat presses, as it surely will, we change the light to red; after four presses, we change it to blue and deliver food for four presses in blue. We raise the response requirement to four presses in green and our chain is established. We have a rat that presses four times in green for the (conditioned) reinforcement provided by the change to red. Four presses in red are reinforced by the change to blue and the presses in blue are reinforced with food.

A *homogeneous chain* such as this, in which the same response is required in each member of the chain, is vastly easier to establish than is the *heterogeneous chain* described earlier, in which different responses are required in each member of the chain. But the principles involved are the same in the two cases. We begin with the final member of the chain, reinforcing the behavior required there with food, or some other potent reinforcer. Then we go to the penultimate member of the chain, shaping that behavior with reinforcement in the form of the opportunity to perform the final behavior of the chain. In the case above, changing the light from red to blue meant that the rat could then perform the behavior required in the final member of the chain. Then we work backwards, shaping each subsequent behavior with the conditioned reinforcement provided by access to the next member of the chain.

Training Pliny In the case of Skinner's demonstration (1937), the performance seems far more difficult to train, which is why the original account appeared in *Life* magazine. What was the final behavior in that chain? A large part of the skill required in establishing such a chain lies in the identification of the successive behaviors that we define as the chain of behaviors. What we want is a rat that walks over to a string, pulls it, picks up a marble, carries it to the tube, lifts the marble, and drops it in the tube. What is the last member of the chain? Is it dropping the raised marble into the tube? How would we establish such a behavior?

The method described in the following paragraphs is probably the best strategy for establishing a chain, but it is not clear that this is the method that young Skinner followed with his famous rat, Pliny. The *Life* account is merely a picture story, and Skinner does not spell out the method he used in his 1938 reference to Pliny. He did view it as quite a feat, however, and it appears to have taken him a long time, so he may have used a less enlightened method.

If we were to place a rat next to the tube and give it a marble at a height sufficient to allow the rat to drop the marble into the tube, we probably would not be wise. True, we might be able to place it in the rat's paws and the rat would surely drop it. But I think you will agree that the likelihood is not very high that the marble would enter the tube, and both we and the rat would tire of the handing-dropping sequence long before we and the rat were successful. How else could we do it?

Suppose we simply place the marble at the edge of a hole, flush with the floor. When the rat in its travels happens to push the marble into the hole, we immediately give the rat a food pellet.

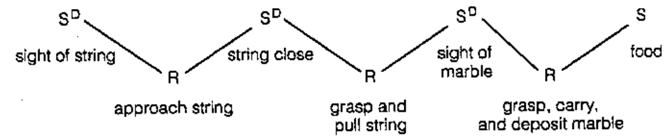

FIGURE 8.6 Establishing a chain of behavior with Pliny, Skinner's famous rat. This diagram shows one way of dividing the sequence of steps involved.

Soon, this simple behavior will be established and the rat will routinely walk over to the marble and push it into the hole. Then we place the marble an inch away from the hole and feed the rat when it pushes the marble over to the hole and drops it in.

Then what? Suppose that we then insert the tube in the hole so that the marble could not just roll in and the rat has to lift it a bit to push it in the hole. We progressively raise the tube, say a quarter of an inch at a time, until the opening is a full two inches above the floor. If we are patient and careful, it should be no great trick to shape this final behavior in the chain. The rat will lift the marble to the top of the tube; the dropping of the marble will be anticlimactic. One could define the dropping in as the final member of the chain; the lifting up of the marble is reinforced by the opportunity to drop it.

Next, we move the marble farther and farther away from the opening, requiring the rat to retrieve it, carry it to the tube, and drop it in. That becomes a third and easily established member of the chain. Our setup here will likely do much of the shaping for us; the sight of the marble should be enough to set the occasion for carrying it to the tube, as long as we do not position it too far away at the outset. Now we are ready for the final (actually, the initial) member of the chain—pulling the string to release the marble. How do we establish this step?

We could place marbles in a rack so they are visible to the rat but out of its reach and held in place by a gate that is opened when a string is pulled. We place the rat in the box, with the tube and the rack of marbles visible. Both are paired with the receipt of food and we may therefore expect our rat to be active. As it moves around the area of the marble rack, we shape its approach to the string by releasing marbles one at a time, each time allowing the rat to go through the rest of the chain, thus more securely establishing the final members. In time, we require touching the string and finally pulling the string before a marble is released. In the end, we have a rat that readily pulls the string, takes the released marble to the tube, raises it, and drops it in the hole. We have established a sequence of behaviors, each of which occurs in the presence of specific stimuli, and each of which is reinforced with the stimuli that allow the next step of the chain to be entered. The string pulling occurs in the presence of the string and the rack of marbles and is reinforced by the release of the marble. Carrying the marble occurs in the presence of the marble (as an S^D) and is reinforced by the stimuli of the marble over the tube opening. The dropping of the marble occurs in the presence of the marble over the opening and is reinforced with the delivery of food. Figure 8.6 illustrates this sequence.

The Significance of Shaping, Chaining, and Conditioned Reinforcement

These principles account for a great deal of our behavior and experience. They are effective whether they are used intentionally, as in the examples above, or as used by society or nature, where the intention is less clear or absent. Let us consider some examples.

Obviously, we learn to pronounce words when our attempts are shaped by the listeners in our family. Are the sentences we learn simply chains of such words that are joined together? Think of the way in which we establish behavioral chains and you will conclude that this is not the way in which we learn sentences! Shaping is important in the learning of words, in acquiring manual skills, and in every activity in which we engage. If we work for a supervisor who compliments us for suggesting new ways of doing things, then our behavior is shaped toward making suggestions; we become more creative.

Might we view the sequences of activities in which we engage on a daily basis as a chain or even as a series of chains of behaviors? Early behaviorists suggested that we could, but they were certainly mistaken. Let us consider why.

To remain simple and concrete, consider the words in a sentence such as "I plan to go to the beach." If we naively think of this as a chained sequence of spoken or read words, we miss an important point concerning the proper identification of behavior classes. What sort of education is required to enable us to use such sentences? Is it practice in the linking together of words, or is it practice in the rules of sentence construction, that leads to the ability to use and respond to such sentences? We learn that sentences must have at least a subject and a verb and that modifiers, objects, and subordinate clauses may be used only in prescribed ways. What is shaped is the proper use of syntax, not a chain of spoken words.

In other cases chaining seems more probable. In a boxing match a fighter's behavior is shaped and maintained by the reactions of his opponent, and signs of fatigue or damage to the opponent may act as signals that the desired outcome is approaching; this reinforces the behavior that produced these signals and sets the occasion for new behavior that has been useful in the past, when other opponents have shown similar signs. But even here, we are not talking about a rigid chain of behaviors "rattled off" in a simple and fixed way. The class of behaviors comprising each member of the chain is far more complex and seems clearly lawful only to a boxing expert or to someone who is familiar with the techniques of the fighter involved.

The point is that shaping, chaining, and conditioned reinforcement are surely important in regulating our daily activities, but we must remember that the behaviors and conditioned reinforcers involved are not as easily analyzed as is the case with the trained rat.

In addition, we may return to Skinner's early definition of the reflex and ask whether an established chain may act as a unit of behavior in its own right, rather than as a sequence of activities. Is the very experienced rat progressing through a series of string pulling, marble lifting, carrying, and dropping, or can the whole series be viewed as a single behavior? In the example of the boxer, thousands of hours of training may establish a chain consisting of a left jab followed by an overhand right, but does this remain a two-element chain in the champion?

According to Skinner's criteria for the identification of behavioral units, we may vary factors such as food deprivation, sleep, amount of reinforcement, and other things and note whether orderly changes occur in the behavior under consideration. If we find that the rat goes through its chain more or less quickly or that the boxer uses his jab–overhand-right combination more or less frequently, we are justified in referring to these chains as unitary behaviors. But Skinner continued to treat chains largely as sequences of individual behaviors, although he occasionally refers to a chain as a response.

How may chaining apply to language? In considering language, it is apparent that, although sentences may in some cases amount to the chaining together of words, typical language is surely more than that. We learn classes of behavior that are more accurately called *using sentences that conform to the rules of syntax*, and the individual word is no more an elementary unit than is a particular muscle contraction that occurs in pressing a bar or throwing a ball. Skinner's critics, such as Chomsky (1959) should not be overly criticized for assuming that he, like some earlier psychologists, viewed language as no more than the chaining of words. We will see later that this was by no means the case.

Schedules of Reinforcement

Until the 1950s, psychologists were very unsuccessful in accounting for behaviors that occur over any appreciable span of time; it was risky enough to predict the choice a rat might make in a maze or the amount of forgetting a human subject would show after practicing a list of words. Ferster and Skinner (1957) were the first to show that the scheduling of reinforcers could produce extremely reliable patterns of behavior that could be maintained as long as was desired, often extending over a considerable fraction of the life of an experimental subject such as the rat. The rule by which reinforcers are delivered is called a *schedule of reinforcement*.

Skinner (e.g., 1974) identified parallels to simple reinforcement schedules in the everyday life of humans and, in a number of cases, explained some puzzling aspects of human behavior. If we show that compulsive gambling is maintained because the schedule of reinforcement involved is one that has been shown to maintain persistent behavior in rats or pigeons, we have explained gambling in a real sense. Unfortunately, the number of such examples has remained small and in most cases there is no obvious counterpart in our lives for the great number of reinforcement schedules that have been studied. Nonetheless, the study of schedules of reinforcement has told us some extremely interesting things, many of which have not yet had a real impact on our views of the control of behavior and experience. Before we consider these recent findings, we will discuss the more common reinforcement schedules and describe the patterns of behavior they produce.

Fixed-Interval Schedules The simplest schedules provide reinforcement based either on the passage of time or on the completion of a response requirement. A *fixed-interval* (FI) *schedule* reinforces the first response that occurs after a set period of time has passed. Thus, a fixed-interval one-minute schedule provides reinforcement for the first response that occurs after one minute has elapsed since the last reinforcement was delivered. A reasonable organism on such a schedule would wait for a minute or more, respond, wait for a minute or more, respond, and so on; only one response is required for each reinforcer delivered.

Such is not the case, however. Pigeons, rats, and other animals (including humans under many conditions) respond many times more than is necessary, even after months and months of daily experience with the same FI schedule. The typical pattern of responding appears as a scallop, which consists of a pause at the onset of the interval, responding beginning after a third to a half of the interval has elapsed, and an increase in the rate of responding up to the time that food is delivered. This pattern is illustrated in Figure 8.7. One might think that subjects behave in such a way because they simply fail to time properly; they underestimate the time remaining until reinforcement is available and thus begin responding too early. But this does not appear to be the case. One would think that some improvement in time estimation would occur after weeks or months of experience with the same schedule value, but this does not occur. Instead, the scallop simply becomes smoother, with little change in the number of responses involved. In addition, the one-third to one-half interval pause occurs over a fairly wide range of schedule values. If time estimation were involved, there would be some intervals (for example, ten seconds) that a reasonable subject could accurately estimate.

Note that the scallop pattern of responding is exactly what occurs in Pavlovian delayed conditioning, where a CS is presented and a UCS appears after a delay period. The Pavlovian procedure does not require a response, of course; the UCS comes no matter what the subject does. The analogous schedule of reinforcement is a *fixed-time schedule*, which is a fixed-interval schedule without the response requirement. Many

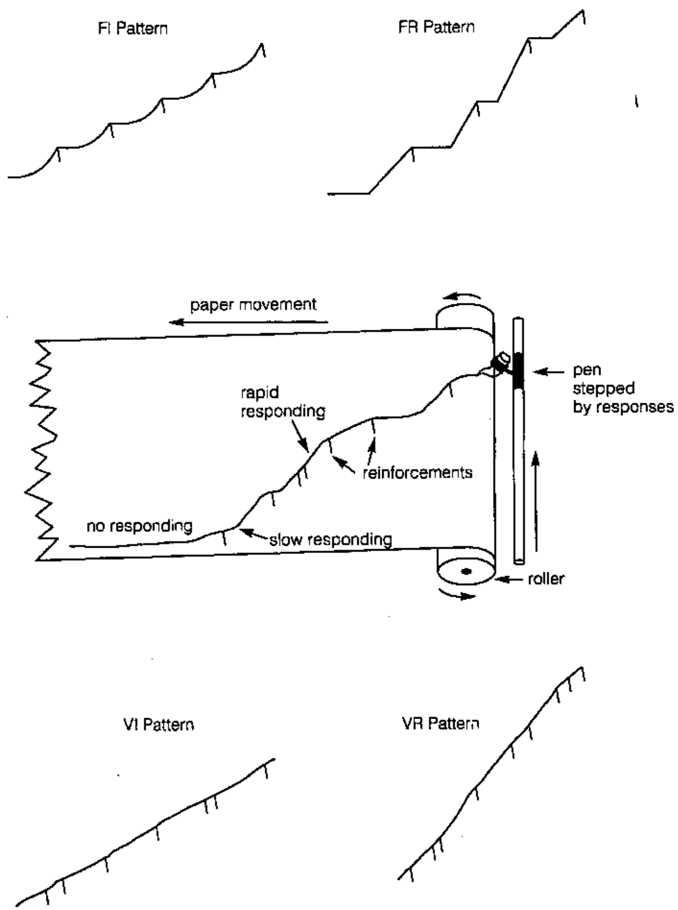

FIGURE 8.7 Patterns for basic schedules of reinforcement. *After Mazur, 1986*

popular examples of FI schedules in daily life actually refer to fixed-time (FT) schedules, as we will see later.

Fixed-Interval Schedules in Daily Life We may train a rat, a pigeon, or a child on a fixed-interval schedule and be assured that the appropriate pattern of responding will endure day after day for as long as we want to keep the schedule in effect. Are such schedules in effect in daily life? If we could show that they were and that the appropriate pattern of behavior were maintained, we could thereby explain such behavior. Suppose we find that students study little immediately after an examination and that they increase their rate of studying as the next scheduled exam approaches. Similarly, we may find that workers slack off just after payday and increase their rate of working as the next payday approaches. Is this an instance of the influence of FI schedules in daily life?

This may seem to be the case until one considers what behaviors are occurring and what behaviors are being reinforced. In the case of students' studying, the behavior of writing the examination answers and turning them in is actually what is reinforced. True, studying may be a prerequisite for writing proper answers, but it is not the studying itself that is reinforced with a good grade, in the sense that the first key peck after a fixed interval has elapsed is reinforced with food. Unlike the studying that precedes the exam, the key pecks preceding the reinforced one are not in any sense required. The same holds with the relation between rate of working and time since payday. Pay is actually given according to a fixed-time schedule, with the proviso that the worker be an employee. To be an employee usually requires that work be done between paydays, but the schedule clearly is not a simple FI arrangement.

Examples of fixed-interval reinforcement in daily life must refer to simple cases. Although a watched pot never boils, we may repeatedly check it to see if it has; eventually, one of these checks will be reinforced. If we live in an area where buses run on reliable schedules, we may see fixed-interval behavior in the form of looking down the street in the direction from which the bus will come; as the scheduled time nears, the frequency of looking may be seen to increase.

Variable-Interval Schedules A *variable-interval* (VI) *schedule* provides reinforcement for the first response occurring after some average period of time. Thus, on a VI one-minute schedule, reinforcement may come for the first response after ten seconds has passed, then after five minutes, and then after thirty seconds. But the average value of the interval is one minute. Subjects on such schedules respond at moderate and fairly constant rates, even when the frequency of reinforcement varies quite widely. For this reason, VI schedules have been used often as a baseline against which to assess the effects of other factors, such as punishers or drugs.

Variable-Interval Schedules in Daily Life In daily life we have to look long and hard for simple examples of VI schedules. In fact, it is so difficult to find examples that we will suggest one that is only close to a valid example (Schwartz, 1978). If we are trying to call New York on Christmas Day, we will frequently get busy signals or recorded messages telling us that all lines are busy. At some time, we know not when, the lines will be free and our dialing will be reinforced; our call will get through. This is a variable-interval schedule even if we try our call only until the effort pays off; whether we continue dialing after that time, lines are free after variable periods of time pass and those times indeed do have some average value. But on a true VI schedule, once an interval has elapsed and a reinforcer is made available, it remains available, unlike that in our example. A line may be free for only a matter of seconds, and if we do not call quickly enough it will no longer be free. This is thus a VI schedule with a *limited hold*, which in this case is also of variable duration. This means that after the interval elapses, reinforcement is available for only a fixed time and if a response is not made within that time the reinforcer is lost and we must wait for the next interval to pass. If you live in an area where the time of mail delivery is variable, checking your mailbox may provide an example of behavior maintained by a VI schedule.

Fixed-Ratio Schedules A *fixed-ratio* (or FR) *schedule* requires that a given number of responses be made before reinforcement is delivered. For example, an FR 20 schedule requires twenty responses for each reinforcement. Animals on FR schedules behave similarly—each reinforcement is followed by a pause that is then followed by a high steady rate of responding. These high response rates remain more or less the same when the ratio requirement is changed; it is the pause after reinforcement that increases when the ratio requirement is increased and that decreases when fewer responses are required.

Fixed-Ratio Schedules in Daily Life We can easily recognize instances of FR schedules in daily life. One common example is paying workers on a piecework basis; fixed amounts of pay are given for so many units of work performed. It may seem, as we watch a subject on an FR schedule responding at a high rate, that he, she, or it must like responding under such conditions. But this apparently is not the case. When we ask workers, they often express dissatisfaction with piecework pay. When we ask the pigeon on an FR schedule, we get the same answer. How do we "ask" the pigeon? Needless to say, we cannot ask it directly, but we can do so indirectly, as was shown by Azrin (1961).

If a pigeon is pecking a lighted key and an FR schedule with a large response requirement is in effect, we may add a second key that has no effect other than to turn off the FR key and stop the FR schedule. Azrin found that pigeons occasionally pecked such a key, typically during the pause after FR reinforcement. This is an interesting finding, since no one is forcing the bird to peck the FR key in the first place. If it wants a respite from pecking, it need only stop, ignore the key, and rest! Why not simply do that, rather than pecking a second key to turn off the FR key?

Perhaps this is why workers complain about piecework pay, and perhaps it tells us something important about reinforcement schedules. An FR schedule prescribes what behavior is required for reinforcement and any organism under the control of an FR schedule behaves in the same way; it pauses after reinforcement and responds at constant high rates. Whether it likes the schedule is not relevant.

Variable-Ratio Schedules On a *variable-ratio* (VR) *schedule*, reinforcement is delivered after a varying number of responses, depending upon the average requirement of the schedule. That is, a VR 25 schedule requires an average of 25 responses per reinforcement, although some reinforcements may be produced by two responses and others by a hundred; the average number is 25. Animals on VR schedules respond at fairly high constant rates, even through periods in which the response requirement is in the hundreds. Once such a large requirement is met, the next reinforcement may come after a requirement of only two or three responses; these easily obtained reinforcers more than make up for the sections of the schedule that require hundreds of responses.

This conclusion is indicated by the finding that animals choose a VR schedule over an FR schedule of the same value; that is, given a choice, a rat or pigeon will spend its time on a VR 25, rather than an FR 25 schedule, even though the average number of responses required for each reinforcer is the same on both schedules (Fantino, 1967; Herrnstein, 1964). Evidently, the animals weight the easily gotten reinforcers more than those less easily obtained.

Variable-Ratio Schedules in Daily Life Here we have no trouble finding parallels in the daily life of many humans. Slot machines pay off according to VR schedules, and it is a well-known fact that a brief "lucky streak," which occurs when in a part of the schedule requiring a few low response requirements, is enough to keep the gambler responding through the large response requirement portions of the schedule. This is when the house wins, having found long ago that humans, like rats and pigeons, are more affected by the low ratios than by the high ones. Perhaps pigeons and rats believe in lucky streaks, though they probably deserve more credit than that.

Another common instance of VR-like scheduling appears in the video arcades, which have replaced the pinball machines and the pool halls of the past. The parallel is not exact: A video game does not require so many repetitions of the same response for payoff, but if we consider such behavior as a class of responses, the parallel is close. Thus, a class of responses all have the common effect of destroying an invader from space, or saving a small bulemic person from ghosts. Some skill may be involved, but a pigeon's skill in pecking at high rates and in hitting the key also is required on a simple VR schedule. And, unlike the video-game "addict," the pigeon at least receives food for its labors. Those who feel that an unfathomable gulf exists between human and beast are well advised to watch a pigeon pecking for food delivered according to a VR schedule and then go to a video arcade.

Other Schedules of Reinforcement The number of ways of scheduling reinforcements is infinite, of course, and researchers have studied a large variety of schedules, both with animal and human subjects. Time is a defining property of FI and VI schedules, as well as of DRL and DRH schedules. Differential reinforcement of low rates of response schedules, or *DRL schedules*, require that responses be spaced. For example, a DRL ten-second schedule requires that a response occur and that the next response occur no sooner than ten seconds later. If a response occurs earlier (for example, at seven or nine seconds), a timer resets and no response must occur for ten seconds from that time. A response that occurs once every nine seconds would therefore never be reinforced.

A schedule such as this is rare in daily life, although it may be imposed to alter certain kinds of behavior. For example, an obese patient may receive pieces of food for pressing a button, with the food delivered according to a DRL schedule. This would produce low rates of responding if food were to be obtained regularly. Such a method has proven effective in changing the eating habits of obese patients.

One may well ask how this could be; how can such training change eating habits? It happens that there is evidence that the obese do eat more rapidly than do nonobese people (Schachter, 1971), independent of their state of "hunger" (time since last feeding). In addition, they eat fewer meals, so it follows that they eat more per meal. Unless a person is eating alone, the duration of a meal is more or less fixed by the person's eating mates. So if one could induce an obese patient to eat more slowly, he or she would eat less overall.

Differential reinforcement of high rates of response schedules, or *DRH schedules*, provide reinforcement only if a specified number of responses occur within a set time. This resembles the procedure on some TV game shows, in which answers must be given within a set time. Whether we are talking about a pigeon pecking twenty times in five seconds or a person naming all of the states beginning with the letter *N*, the result is the same: The procedure produces high rates of response.

Some other schedules combine time and response requirements in other ways. A *conjunctive schedule* requires that both a response requirement and a temporal requirement be met: One might have to respond at least twenty times over 30 seconds for reinforcement. On an *alternative schedule*, the meeting of one requirement is enough; the subject may respond either twenty times *or* once after 30 seconds. An *adjusting schedule* may change the response requirement as time passes. For example, a response requirement of twenty may change to 25 after ten seconds if the required twenty responses do not occur. At ten seconds the requirement could jump to 30, and so on. Or the change could go in the opposite direction; an initial requirement of 100 responses could decrease to 80 after ten seconds, to 60 after twenty seconds, and so on. A lazy subject could just wait for a minute and respond once for reinforcement.

Schedules may be combined in many ways. A *concurrent schedule* provides two or more available schedules of reinforcement simultaneously; a subject may respond to the lever or key of only one of these schedules at a time, but he, she, or it may switch among schedules as often as desired. Such schedules are widely used in the study of the *matching law*, which holds that relative response rates match relative reinforcement frequencies. The matching law will be discussed in greater detail in Chapter 9.

Tandem schedules, like tandem tractor-trailers, are linked together in sequence. A TAND FI 2 FR 10 DRL 2 requires that two minutes pass, whereupon a response changes the schedule to FR 10. After ten responses, the schedule becomes DRL 2 (usually seconds), so that one response, followed by at least a two-second pause, is reinforced. All three schedule requirements must be met, and the fulfilling of one schedule requirement puts the next schedule in effect. In a tandem schedule, no stimuli signal which schedule is in effect. When there *is* a signal, we have a *chained schedule*. Hence, a red light may mean that FI 2 is in effect; when a response occurs after two minutes, the light may change to green, signaling the FR 10 member, followed by orange, signaling DRL 2.

Schedules may also be presented successively, independent of the subject's behavior. A *mixed schedule* provides fixed periods of different schedules in succession with no distinctive S^D signaling each. Thus, a MIX FR 20 VI 5 may provide a two-minute period during which FR 20 is in effect, followed by two minutes of VI 5, followed by two minutes of FR 20, and so on. If a distinctive stimulus signals the schedule in effect, thus identifying the *component* of the schedule in effect, we have a *multiple schedule*. Multiple schedules typically are used in studies of discrimination learning, as discussed earlier. We might alternate periods of red and green accompanied by VI reinforcement and extinction (EXT) respectively, to assess the ability of our subject to distinguish colors.

Human Performance on Reinforcement Schedules

You may wonder if people really show these patterns of responding when they are used as subjects in laboratory experiments studying reinforcement schedules. It seems clear that piecework and gambling do influence human behavior just as fixed- and variable-ratio schedules influence animal responding. But would a normal human on a fixed-interval schedule show "scallop-patterned" responding, and, if the schedule were changed to fixed-ratio, for example, would the pattern of responding change appropriately? The answer to this question depends on whether the human subjects receive the same "instructions" as animal subjects. That is, if human subjects are told to press a button to make a counter count, they press and their pattern of pressing does not change much as the reinforcement schedule changes.

This illustrates the powerful effect of rules on human behavior, a topic to which we will return. Subjects who follow the verbal rule "press the button" do just that; if the counter occasionally counts, that maintains their pressing, but the patterning of counts does not strongly influence the patterning of their responses.

To avoid such effects, Matthews, Shimoff, Catania, and Sagvolden (1977) ingeniously shaped humans to press a lever, using a counter and a light and a button corresponding to "food magazine operated" and "eat," respectively. Such labels were not actually used, but readers familiar with the shaping of a pigeon's response to a key would view the procedure as identical. The experimenters were not visible to the subjects and shaped the subjects' pressing and eating without the use of verbal instructions. After this was done, the human subjects' responding was appropriate for the different reinforcement schedules that were used.

The Significance of Schedules of Reinforcement

It is difficult to imagine the excitement of researchers of the 1950s, when the study of reinforcement schedules was probably at its peak. For the first time, we could completely control the behavior of a vast number of living things. Not only could we precisely predict the behavior of an individual over days, weeks, or months, but we could also predict that one individual rat would behave like another rat, or like a pigeon, or like a cat, dog, rabbit, child, or Nobel prize winner. Surely, we were on to something!

In addition, if we look to the world around us, we can identify schedules operating in daily life; after all, we have the pieceworkers and the gamblers responding on FR and VR schedules. To change behavior in these and other cases, we need only change the schedule in force.

Besides controlling behavior and identifying schedules already in force in the world around us, many researchers felt that the study of schedule behavior was important in another way. Through their study, we can find out more about the mechanisms underlying behavior. Performance on interval and DRL schedules can tell us something about a subject's ability to estimate time, for example. FR performance may provide clues regarding the ability to count. Through the analysis of schedules we can assess the importance of conditioned reinforcement, chaining, temporal discrimination, and even of such faculties as attention and short-term memory.

Such hopes were dashed during the 1960s and 1970s. As Jenkins (1970) and Mackintosh (1974) concluded, on the basis of masses of evidence, schedule behavior is not explainable in terms of any simpler, more complex, or more vague principles, such as conditioned reinforcement, expectancies, or memory. Subjects that can easily time a twenty-second interval nonetheless waste thousands of responses on a fixed-interval schedule, even after months of daily sessions.

Reinforcement schedules are just *there;* they are basic laws governing behavior, as Morse and Kelleher (1977) and others have argued. They have made a good case for this, showing that the effect of a schedule per se is enough to keep a cat, rat, or monkey responding for weeks when the only reinforcer is electric shock. These data have implications for human masochism. The study of schedules is justified when it leads to important discoveries such as this.

Treatment of Human Psychotics

Many patients in mental hospitals spend decades as so-called hopeless cases, unresponsive to psychotherapy or to drugs. Such patients often refuse to eat, to dress themselves, and to practice the most basic personal hygiene. Ayllon and Haughton (1962) proposed to apply Skinner's principles to treat a population of such patients, diagnosed as chronic schizophrenics.

But how may that be done? What could act as a reinforcer to alter the behavior of patients who appear to want nothing and whose files show them to be "out of reality contact," subject to psychotic intrusions," and suffering from faulty "ego identification" (Ayllon & Haughton, 1962)? Using food as a reinforcer seemed out of the question, since half of the 45 patients refused to eat, even after tube and intravenous feeding and electroshock treatment. When they ate, they were spoon-fed. How does one induce such patients to eat without assistance?

In a situation such as this, the therapist considers the reinforcement contingencies that currently maintain the patients' behavior and those that might alter the behavior. These patients had refused to eat unless spoon-fed; in one case, this had gone on for seventeen years. Could it be that the attention accompanying spoon-feeding acted as a reinforcer, maintaining the refusing-to-eat behavior? How could that situation be changed so that the health of the patients would not be endangered?

Ayllon and Haughton began a simple and effective procedure involving 32 patients. Each meal was available at a fixed time and access to the dining room was available for 30 minutes; then the door was locked and those patients who did not come to the dining room on their own missed the meal. Meals were missed at first, but patients who had not done so in years soon entered the dining room and ate without coaxing. Two patients refused to cooperate and went hungry for seven and fifteen days, during which time their health was closely monitored by the staff. But, eventually, they came to the dining room as well. The 30-minute access was reduced to twenty minutes, then to fifteen minutes, and finally to five minutes. The change in the patients' behavior was a "revelation to the nurses" (Ayllon & Haughton, 1962), who had spoon-fed the patients for years, but it should not have occasioned surprise. Food acts as a reinforcer under many circumstances, but so does spoon-feeding, which provides personal attention, as well as food.

The therapists later required that a penny be dropped in a collection can for access to the dining room and they even established some social cooperation behavior, by requiring that two buttons be pressed for access to the room. The buttons were far enough apart that two patients had to press simultaneously. All but one of 43 patients learned to do this. All of this is not much improvement, you might say. But imagine what the staff charged with the care of these patients thought of it! It was only weeks before that many of these patients had to be fed, clothed, and groomed by others.

In a later study, Ayllon and Azrin (1964) used a simple strategy to alter the eating behavior of eighteen patients who refused to pick up their cutlery and consequently did much of their eating with their hands. It took three months to achieve almost 100 percent success, but all it took was the addition of contingencies for a specific behavior. Patients who picked up all three pieces of cutlery immediately went to the serving line, while those who did not were sent to the end of the line awaiting trays and cutlery. This meant a delay of only five minutes or so, but it acted as a sufficient consequence to establish the new behavior.

The most elaborate treatment using Skinner's principles is the famous token economy established by Ayllon and Azrin in the early 1960s. This was described in 1965 and presented in more detail in 1968. A *token economy* arranges reinforcement for socially-desirable behaviors. The patients in their study were female psychotics, many of whom had spent twenty to 30 years in the chronic schizophrenic ward. Instead of using food as the sole reinforcer, the therapists chose things that seemed valued by specific patients. Some of these were:

Choice of one of five sleeping rooms	A twenty-minute walk
Use of a personal chair	A trip to town
Choice of eating group	A room divider
A personal coat rack	Candy, books, and other goods

All of these things were purchased with metal tokens that were contingent on work done by the patients. A choice of room cost from one to 30 tokens per day, depending on the room, a trip to town cost 100 tokens, a twenty-minute walk required two tokens, and so on. The television set was coin operated (fifteen minutes per token) and access to the dining room was through a token-operated turnstile.

In one situation, eight psychotic subjects worked outside the ward; they served food, typed and answered the phone, cleaned the laboratory, or worked in the laundry. Each job paid 70 tokens per day. Patients volunteered for the jobs and shifted to nonpreferred jobs when payment was no longer given for preferred work. When tokens were given free, work ceased, so their behavior was not solely dependent on the opportunity to leave the ward or to do work that was intrinsically satisfying.

The program was expanded to include 44 patients, most of whom were diagnosed as chronic schizophrenic and who had been hospitalized for an average of sixteen years. All suffered from "severe psychosis" (Ayllon & Azrin, 1968), many had extremely low IQ test scores, and some were very deficient in verbal abilities. Yet all were affected by the token system; when appropriate behavior produces tokens exchangeable for little things that are valued by patients, patients respond. Aside from eight patients who did little but sleep and eat, all worked steadily as waitresses, clerks, janitors, and launderers. The least responsive still showed improvement in grooming and other self-care behaviors.

Bear in mind that these patients had been diagnosed as hopeless, had been hospitalized for many years, and initially were incapable of even the most rudimentary self-care. Yet, the alteration of consequences for their behavior transformed them. Where passivity had been reinforced with attention (for example, in spoon-feeding), work that could become satisfying in itself was reinforced with tokens, a currency that could purchase things of value to the patient. And it seemed at the outset that there were no things of value to them.

Concept Formation

The class of stimuli controlling an operant class need not be composed of simple and concrete elements. Often the effective stimulus class takes the form of a concept; rather than being controlled by lights, sounds, and tastes, behavior is controlled by concepts such as truth, beauty, people, water surface, edible object, and noun-verb order. There is nothing mysterious about such concepts and their influence has even been studied with animal subjects. Simply put, a concept refers to the fact that groups of particulars are treated as a class; they are exemplars of that concept. For example, the concept of truth includes all that I count as an aspect of my experience that has not been negated by other experience. My concept of people is the set of common characteristics according to which I identify individuals as people. This includes bodily characteristics such as arms and legs, as well as behavioral characteristics, such as the ability to use language. How did I come to have such a concept? Following Plato, we could assume that I was born with a concept of people; when I encounter an instance of this concept it awakens whatever innate archetype I possess and I recognize an individual as a person. On the other hand, my concept could arise from my training and education. Since I was an infant, I have been taught that such and such a being was a person and that other beings, though they have arms and legs, were not. Which alternative is more likely correct?

During the Second World War Skinner and his colleagues, in their attempt to aid the war effort, showed that pigeons were capable of learning fairly abstract concepts. The military had developed missiles for use against the Japanese, but a reliable guidance system was lacking. The missile was in the form of a glider that could be launched from a large bomber, but where it went then was very uncertain. The Japanese had solved the problem with a living guidance system and the kamikaze was essentially a missile directed by its human occupant. Skinner proposed a living guidance system in the form of pigeons.

In the laboratories of General Mills in Minnesota, pigeons were trained to peck a semiconductor screen, through which was visible a target projected on the wall. A target could be a film of a Japanese ship, although a photo of an intersection in northern New Jersey was also a frequently used target. The pigeon was strapped in a harness and its beak tipped with gold; pecks on the screen controlled the direction of movement of a "missile," consisting of a table on wheels driven by small motors. As the pigeon pecked the screen for occasional food, the mock missile moved across the room to the target.

Eventually, several pigeons were incorporated into the system to increase reliability. Suppose that two Japanese ships appeared on the screen and that the pigeon pecked alternately at both; the missile's course could be disrupted so that both were missed. This possibility was avoided through the use of a panel of three pigeons, which "voted" on the target of choice. As many as eight pigeons were sometimes placed on such panels (personal communication from Norman Guttman, a worker on the project).

The device worked beautifully. The pigeons acquired the concept of "Japanese ship" of such and such class and developed a peculiar interest in an aerial view of a location in New Jersey. (The system was demonstrated to the military in Washington, which saw its potential usefulness, but which declined to fund its further development. Work going on at the time in New Mexico was receiving higher priority, and it promised to eliminate the need for precise bombing!) See Skinner (1983b) for further details concerning the project.

Another example of *concept learning* in pigeons was reported by Thom Verhave (1966), then a scientist working for Lilly pharmaceutical company. In the manufacture of capsules for their drugs, Lilly relied on skillful women, who sorted capsules as they passed along a conveyor belt. They picked out faulty capsules and their pay was docked if later inspection revealed that they had made many mistakes. This was not an enviable job! Yet, it had its glamour; accurately spotting faulty capsules required a lot of experience and could not be done by novices. The concepts "good capsule" and "bad capsule" took time to learn.

Verhave quickly trained pigeons to do the job as well as and even better than the human sorters. By placing a pigeon in a box by the conveyor, the pigeon could see the capsules as they came and it was occasionally fed for picking a clear key (or window) when a good capsule went by. One could also train it to peck a second key when a bad capsule was present. After food reinforcement for identifying clearly good and faulty capsules, samples more difficult to classify were used until the pigeons were correctly sorting capsules that previously required a skilled human inspector to classify. The birds had learned the concept of "good capsule."

Like Skinner's Project Pigeon, this method remained an example of the learning of concepts by pigeons; it was never used by Lilly. Apparently, their public relations people advised against pigeon pill inspectors, since undesirable images were bound to be conjured up by the public. Nonetheless, Verhave, supported by Lilly, applied for a patent, which was ultimately denied. A suit was filed against the Federal Patent Office, challenging the denial; Verhave believed that the denial could not have been due to the fact that living things were involved, since many other devices (such as horsecollars and typewriters) require living things for their operation.

The answer finally came and revealed the government position on animal mentality. The patent was denied because "one cannot patent a mental process," thereby indicating that the simple performances of these birds was evidence for mental activity. There is little doubt that the women capsule inspectors were relieved that the project was not used by Lilly; many workers have been replaced by machines, but being replaced

TABLE 8.2 Concept Learning by Pigeons. General method used by Verhave (1966), Herrnstein and Loveland (1966), Herrnstein, Loveland, & Cable (1976), and Skinner (1983b).

Training	
1. Samples of clear "positive instances" (of ships, pill capsules, trees, people) —mixed with—	Pecking produces occasional food
2. Samples of clear "negative instances" (of friendly ships, bad capsules, scenes with no trees, people) —followed by—	Pecking produces no food
3. Samples of less clearly different positive and negative instances —followed by—	Pecking positive produces occasional food
Testing	
4. Presentations of positive and negative instances not included in the training sets	Pecking has no consequence

by an animal, especially a pigeon, would be especially unpleasant.

In a similar study, Herrnstein and Loveland (1966) showed that pigeons readily learn the concept of *people*. They trained pigeons to peck a key when a slide included people or parts of people and not to peck when people were absent. Then they showed the subjects some new slides. The pigeons were amazingly good at distinguishing new slides based on the presence or absence of people. If pigeons can easily learn such a concept, it is likely that humans learn such concepts as well, which is contrary to Plato's theory. See Table 8.2.

However, a few years later the same authors reached different conclusions. Herrnstein, Loveland, and Cable (1976) showed that pigeons could learn the concepts of *trees*, *water surface*, and a *particular young woman* extremely quickly, too quickly in the authors' opinion. They suggested that pigeons must be born with knowledge of

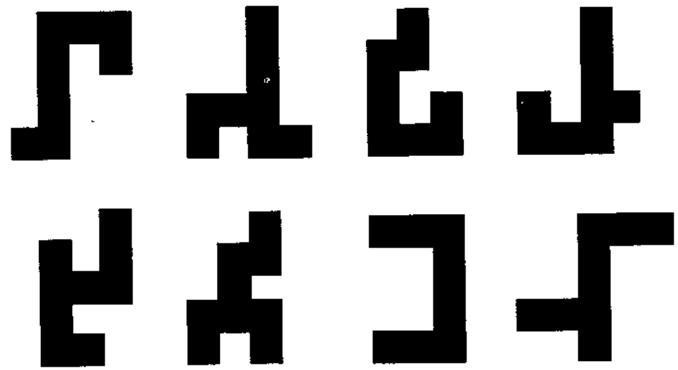

FIGURE 8.8 Some of the forms presented to the subjects (pigeons and humans) in Hollard and Delius's mental rotation experiments.

many things, in the form of templates or schemata that assist them in rapidly learning concepts that correspond to common features of the world. Since there could be no precise schema for the young woman the pigeons came to know, their foreknowledge must be in the form of a readiness of some sort to learn particular concepts.

The authors referred to these concepts that the pigeons seemingly came prepared to learn as "natural concepts." Later, Herrnstein's group showed that pigeons also seem prepared to learn the concept of *fish* extremely quickly (Herrnstein & deVilliers, 1980). Why pigeons should come prepared to learn such a concept is anyone's guess; perhaps such a finding suggests that Herrnstein and Loveland's earlier explanation was accurate after all.

Mental Rotation

In Skinner's opinion, mental activity is actually no different in kind from other behavior (e.g., 1963). Further, if conditions have been arranged to demonstrate an aspect of mental activity in humans, they may be arranged similarly to demonstrate the same activity in animals. This was demonstrated in the case of mental rotation.

One of the most publicized pieces of research during the 1970s was the demonstration of mental rotation by Shepard and his colleagues. For example, Cooper and Shepard (1973) presented human subjects with letters or digits as they normally appear or as mirror images. Then they presented these numbers and digits rotated to various extents and asked the subjects to identify them as normal or mirror image. An interesting effect appeared. Subjects' reaction time (RT) depended upon the degree of rotation; it increased linearly as a function of angular rotation. Thus, it seemed that subjects were "mentally rotating" the display in order to determine whether it was normal or mirror image, and the farther it had to be rotated, the longer it took. Is the ability to manipulate images mentally an example of the higher mental functioning of which we humans are capable? Or can pigeons duplicate the feat?

Hollard and Delius (1982) cleverly investigated the capacity for mental rotation in pigeons (see Figure 8.8). Pigeons were trained in a matching to sample (MTS) task in which one of the forms was presented on a response key. After fifteen pecks to that key, a display appeared on each of two keys on the left and right of the center key. On one comparison key was the form presented on the center sample key; on the other comparison key was the mirror image of the sample form. For half of the birds, pecks to the matching key were reinforced with food; for the others, pecks to the mirror image key led to food.

After training on a variety of forms, tests were conducted with various of the forms presented with comparison forms (normal and mirror image) rotated in one of five steps, ranging from zero degrees (normal) to 180 degrees. The reaction times of the birds, which were performing essentially the task required of Cooper and Shepard's subjects, was recorded, as were the errors made with various degrees of rotation.

After the data were collected, the pigeon chamber was disassembled, except for the end of the box holding the three keys and the feeder. Twenty-two college students were brought in and trained to perform the pigeons' task; correct responses were reinforced with the equivalent of half a cent each (pigeons had been reinforced with grain). The students were tested much as the pigeons were, and the results were analyzed and compared with the pigeon data.

The human subjects performed much as did Cooper and Shepard's subjects. Their reaction time and accuracy depended on the degree of rotation of the forms, with the greatest accuracy and fastest RT when the comparison forms were not rotated. Accuracy decreased and RT increased as the degree of rotation increased up to 180 degrees. How did the pigeons compare?

Amazingly (perhaps), the authors found that the birds were "more efficient than humans"! Human RT varied from about 1.5 to 2.75 seconds as the degree of rotation increased; the pigeons' RT was always less than a second. And the degree of rotation had no effect on the pigeons' RT. The percentage of errors was approximately the same for the birds and the humans, with the salient difference being that the birds made more errors at rotations of 90 degrees, while the human errors were greatest at rotations of 180 degrees.

The authors suggested that the superior performance of the pigeons may be due to the fact that human estimates of orientation are made with reference to gravity, what we perceive as "down." Birds, on the other hand, use no such reference; for them, orientation is arbitrary and relative to the position of the observer. No one teaches them the artificial coordinates in which we frame the world.

Problem Solving

Skinner's specific views on problem solving appeared in 1966 and are generally consonant with the views of other learning theorists on the topic. He believes, as did they, that there is nothing special about problems set for solution. As Skinner wrote:

> [T]here is probably no behavioral process which is not relevant to the solving of some problem. . . . [A]n exhaustive analysis of techniques would coincide with an analysis of behavior as a whole. (1966a, pp. 225–226)

Unlike his predecessors, Skinner emphasizes the distinction between contingency-shaped and rule-following behaviors. An outfielder's response to the problem of the approaching fly ball has been shaped by past contingencies; what has been the consequence of various responses in the past? Like the wisdom of the old country doctor and the answers that intuition provides, the contingencies involved have not been formulated in verbal rules. (In what situation does a given behavior produce such and such a consequence?)

Contrast that with the commander in charge of the retrieval of an incoming satellite. His behavior has been shaped by an education in the rules of ballistics and the calculating of trajectories. So-called rational, or rule-governed, behavior is shaped without need for exposure to the actual contingencies involved. Such rules are passed on by cultures in the form of maxims and proverbs and in all of the verbal education we receive.

But thinking and problem solving are by no means purely verbal; Skinner believes that operant behavior in all its forms is involved. The contingencies involved may be subtle, and they may

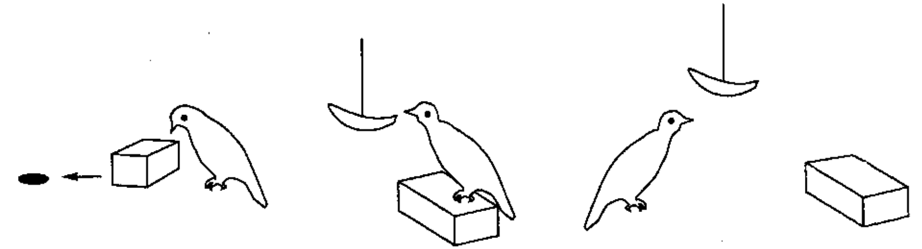

FIGURE 8.9 Simulation of insight in problem-solving experiments by Epstein, Kirshnit, Lanza, and Rubin (1984).

shape behavior even though they are experienced only in the form of learned rules. Needless to say, Skinner opposes models such as those of information processing. The fundamental error committed is the assumption that we are input/output mechanisms; such an assumption then requires cognitive processes to relate input and output and research done only for the purpose of testing predictions of specific cognitive models. Are we interested in how we solve problems or in how contrived cognitive processes may work?

As an illustration of a possible Skinnerian approach to problem solving, consider Epstein, Kirshnit, Lanza, and Rubin's study of insight in pigeons (1984). They proposed that the insight shown by Köhler's chimpanzees, evidenced in the subjects' use of a box as a jumping stand to retrieve a suspended banana, was possible only because of specific past learning experiences. Köhler's subjects must have learned to push objects toward targets and to climb on objects to reach other objects. Köhler's chimps were adults and had had plenty of time to acquire such learning.

The two learned skills are irrelevant to the lives of most pigeons. But could pigeons that were trained to push and to climb then insightfully put the objects together to retrieve an object hanging out of reach? Epstein, et al. trained pigeons to push a small cardboard box toward a green spot located at various positions over a period of one to eight weeks. The pigeons also were shaped to climb onto a box placed under a suspended facsimile of a banana. They were then tested with the banana suspended out of reach and the box positioned away from it (see Figure 8.9).

Like Köhler's chimps, the birds first appeared "confused" (Epstein, et al., 1984) and looked back and forth from the banana to the box. Then "suddenly" they pushed the box to a position under the banana, "sighting" the banana as they did so. They placed the box appropriately, stepped onto it, and pecked the banana. Is that insight, in the sense that Köhler defined it? The birds never were trained to push the box toward the banana, but their experience in pushing and climbing allowed their insight. Birds that were trained to peck the banana but not to climb, or to climb and peck but not to push, or to climb and peck and push randomly showed no insight when tested.

This is one way to study problem solving. It is of little help to argue that the occurrence of insight requires past learning experience unless the nature of that experience is specified. Similarly, it is of little help when Köhler (1925) suggested that insight is not "blind trial and error" but requires that attention be paid to "structural properties" of the situation. In the box-and-banana problem, insight depends on having learned to push objects toward targets and to climb on objects to reach other objects. Needless to say, other cases of insight require other previously learned skills.

Skinner and Language

Skinner was always interested in literature, as any reader of his popular works can see. After decades of work, Skinner published *Verbal Behavior* (1957), which presents his views on language. In this book, he attempted to provide a *behavioral* account of language; it consisted of only an outline of an operant analysis of language use. No data were included and no great interest was shown initially by Skinner's followers. *Verbal Behavior* did not seem to be a barn burner.

Noam Chomsky, the noted linguist, noticed the book and evidently felt that, as a behavioral treatment of language, it epitomized what he disliked in older linguistic theories. In 1959, Chomsky published a scathing review of it, which was far more widely read than Skinner's book. No one responded to Chomsky until 1970, when MacCorquodale reviewed Chomsky's article. The reason for the long delay was that Chomsky's attack was so unrelated to Skinner's views that no one felt the appropriate object of the attack.

Chomsky's criticism was directed toward the behaviorisms of the 1940s and 1950s that derived from Hull's theory. It attacked the theory of reinforcement based on drive reduction and the extreme molecular view that might explain language as the stringing together of letters and words. In fact, Skinner held no such views. Skinner's view of language stresses its function, and letters, words, sentences, paragraphs, or *whatever* may be appropriate units (1957).

Skinner's point, which is not original, is simply that verbal behavior *is* behavior and, as such, it is learned as is other behavior. Does this mean that language derives from the deliberate reinforcing of spoken words that are then hooked together in chains? Since other behavior is not so characterized by Skinner, there is no reason that language should be.

The functional unit of any behavior is that class of behaviors that covaries as a function of consequences (that is, as a class). Reinforcers are consequences that increase the frequency of occurrence of a class of behavior. Behavior also comes under the influence of a stimulus class that includes whatever stimuli or situations set the occasion for the behavior. Does all of this hold for verbal behavior?

The unit of verbal behavior is not a letter, word, or phoneme (basic sound); as is the case with other behavior, there is no fixed unit. As we learn to speak, our verbalizations are shaped by the reactions of those around us. The function of speech is to influence the behavior of others, and, therefore, the form of our speech is molded so it is effective in doing that. I speak loudly and clearly because that has been effective in the community in which I live. The unit of speech is sometimes a word, and at other times a phrase, idiom, or passage (Skinner, 1953, 1957). I use proper grammar because I am more readily understood when I do so. All of this should be no revelation; speaking is behaving and it is shaped by consequences as are other behaviors.

Not everyone shares such a functional view of language. For example, the philosopher Bertrand Russell marveled that the whole of human knowledge and thought was the consequence of 26 shapes! Letters and their combinations were all that seemed essential in Russell's view of language. Skinner (1957) criticized him, suggesting that one could use Morse code and marvel that all knowledge and thought is expressable with *two* symbols. In fact, of course, the bare symbols we use are a small part of language in general; meaning comes not from combinations of symbols, but from our experience behaving in a complex environment. There is no context-independent pure language. This functional view of language is similar to the pragmatic view of James and Dewey (Chapter 1).

Whether we speak and what we say at a given time is influenced by a host of factors, including the current situation, what we have just said, and so on; verbal behavior typically is controlled by multiple factors. Our current repertoire of verbal behavior depends upon the past consequences of various kinds of verbalization; some reinforcers of verbal behavior are attention, approval, affection, and submission (Skinner, 1953). Skinner (1936) invented an "auditory Rorschach" in the form of the verbal summator as a means of examining a person's current verbal stock.

The summator was a recording that was played in a nearby room; it was designed to resemble difficult-to-understand speech. Subjects that were told to interpret the conversation in the next room actually heard sounds like "eye-uh-ah-uh" and "oo-ee-uh-uh," presented either faintly or against a noisy background. Their reports, like responses to the ink blots of the Rorschach, reflected their current repertoire of verbal behavior.

Private Experience

Recall that early behaviorists, such as Thorndike and Watson, believed that thinking, remembering, seeing, and other mental activities could be studied as other behavior is studied. But Hull and his associates cared little about private events, and Tolman was clear in throwing "raw feels" out of science. Due to their great influence, behaviorism came to be seen as a mindless position and, to this day, a good many behaviorists and critics of behaviorism feel that behaviorism and mental activity are incompatible.

Radical Behaviorism

Skinner's first statement on private experience and behaviorism appeared in a 1945 paper: "The Operational Analysis of Psychological Terms." In it he criticized contemporary operational analyses, citing what he called *"methodological behaviorism."* This is the view that there is a distinction between public and private events and that psychology (to remain scientific) can deal only with public events. Skinner noted that this is the position that accepts the arid philosophy of truth by agreement; something is real if at least two observers agree. Methodological behaviorism leaves the mind to philosophers.

In the same paper, Skinner referred to his own view, which he called *radical behaviorism,* and which does not distinguish between private and public events. It treats a toothache as a stimulus as physical as a typewriter and "seeing" as an activity no different in kind from walking. This view follows Watson in denying the "mental" as different in kind from the physical. It does not mean that all is physical, but that all is of a kind.

Skinner certainly does not deny the existence of private experience, nor does he feel that its study is beyond us. What he does deny (as did Thorndike and Watson) is the mind-body dualism of the mentalists and the methodological behaviorists. There is no separate realm of nonphysical images, thoughts, and feelings: private and public, inner and outer are the same in kind. Thinking is something that we do, just as walking is something that we do. And we think mental thoughts no more than we walk mental steps. Our names for feelings are words that refer to real stimuli just as words refer to objects we see. To feel pain is a reaction to pain-producing stimulation, just as recognizing a friend is a reaction to other stimulation. It is important, however, to realize that feelings do not make us act. When we touch a hot stove it hurts and it makes us pull the hand away. However, withdrawing the hand is a reflex response that would occur whether we felt pain or not. We do not withdraw because it hurts (Skinner, 1974).

Awareness Our families and educational systems obviously teach us to name the objects of the world we seem to share. Society allows the very young child to refer to all four-footed animals as "bow-wows," but it later expects that proper names be used corresponding to different species and later requires that names for subspecies be used correctly. We learn to name things because society (the verbal community) reinforces correct naming. "That's right; she *is* a cocker spaniel," or "That *is* a fir tree." The learning of language is in large part the learning of the verbal community's names for things.

The verbal community also teaches us to name aspects of private experience. "That is a bad bruise; it must hurt." "You have eaten too much, you must have a stomachache." "After all of that work, you must feel tired." "You tried so hard, but failed; you must be disappointed." "After what he did, you must be angry." Based on what circumstances it finds us in and what it can see us doing, the verbal community teaches us the labels for our private experiences. We learn the names for pain, fatigue, disappointment, anger, impatience, sadness, and the other aspects of our experience.

But Skinner suggests that we ultimately know our subjective experience less well than we know the public world around us. This is because of the limitations faced by society as it teaches us to name the varieties of our private experience. Those around us can often tell when we are in pain or tired or disappointed or angry by seeing what has happened to us and/or what we are doing. But in many cases they cannot.

Society may or may not be able to tell when we are diffident, embarrassed, hungry, determined, or in pain. On occasions when it cannot tell, it cannot teach us to label our experience appropriately. Much of the function of great literature, art, and music is to produce private experiences and, in some cases, to provide labels for them. A literary piece that evokes wonder or sadness in the reader helps to teach us to properly label such private experiences when they occur in the future. Great literature is highly valued, in part because it is so scarce; it is no mean feat to evoke strong private reactions in readers.

Once society has taught us to name whatever private experiences it can, it urges us to report such experiences later. "Are you hungry?" "What do you see?" "How does your stomach feel?" We are repeatedly drilled in the reporting of internal events and we become more or less adept at doing it. But reporting that we see a car coming is no different in kind from reporting a stomachache; the main difference is that we can lie more easily about the stomachache, since the questioner cannot see this experience, as he or she can see the oncoming car. There is no reason why a behaviorist cannot study the factors that determine that we see cars or that we feel stomachaches.

However, since society finds it harder to train us to label private experience, we will always be less able to know our private experience. One way in which we do discern such experience uses the same method that society uses to infer it in us so that the proper labels can be taught. That is, just as the verbal community judges our private experience in part by observing our actions, we often judge our own experience by observing our actions.

What does it mean to say that I was "hungrier than I thought?" Such a statement might be made by someone who had ordered food, eaten it, and still wanted more. The individual had estimated his or her feelings of hunger and ordered food that was judged sufficient to quell that hunger. The food is now gone but the hunger remains; the earlier estimate of hunger was mistaken, as is revealed by the fact that the individual still feels like eating.

In another common circumstance, we may believe that we have forgiven someone for a wrong done to us; yet, when we next meet the person we find that we scowl and clench our fists. We didn't realize that we were still angry, but our reaction shows that we are. Similarly, tears and a feeling of sadness when we hear the name of a lost loved one tells us that we still grieve, although we were certain that our grief was past.

Such cases are common and show that we are often poor judges of our feelings. In a great many cases we judge our private experience on the basis of what we know to be our actions, just as society does when it teaches us to name those experiences. What are our feelings toward brown bread, nuclear war, the plight of the poor, or the existence of God? When we consider our opinions on

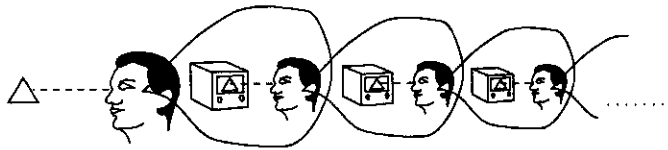

FIGURE 8.10 The copy theory of perception and thought.

such matters, we often think of our behavior with respect to the issues. If an acquaintance eats nothing but brown bread, writes pamphlets opposing nuclear war, and gives freely to the poor, we assess his opinions on such matters accordingly. The next time that you are asked for an opinion on some matter and you must consider it, see whether part of what you consider is what you have recently done with respect to the matter at hand. If it is, then you are judging yourself in the same way that you judge others, based on your own past actions (cf. Bem, 1967).

Conscious Content The British philosophers of the eighteenth and nineteenth centuries believed that the real world was within us, and early psychologists saw their task as the analysis of this "content of consciousness." We do have dreams and reveries, imagine distant places, and recall events from the far distant past. Wilder Penfield (1958) reported that electrical stimulation of the brain apparently produced detailed memories in his patients, evidently roused from their storage chambers in the cortex. How does radical behaviorism deal with such conscious contents?

Skinner illustrates the ancient point of view regarding the content of consciousness by referring to a popular film sponsored by Bell Telephone, called *The Gateways to the Mind.* According to this film, the nervous system resembles (not surprisingly) a telephone system. Stimulation of the body surface produces messages, resembling little lightning flashes, that travel to the brain. In the brain is a television screen that shows a display, which is seen by a "little man" in our heads. Depending on the message, the little man presses various buttons and pulls levers that lead to an outgoing series of little lightning flashes, producing appropriate movement of our body. Thus is explained sensation and perception, thinking and judgment, and action: The little man does it all. What we see, hear, and feel is represented by copies produced on the TV screen and (presumably) by the audio equipment available to the little man.

Needless to say, the little man is intended to represent the wonders of the functioning of the brain and the explanation is fine unless one wonders how the little man operates. Does he have a little TV screen in his head and a yet smaller little man to see it? Such explanations provide only an infinite regress; we never get to any more than copies within us and copies of *us* within us. What have we gained?

It is true that copies of stimulating objects do exist in the nervous system. Visual stimulation is copied on the retina, in the optic nerve, in the lateral geniculate of the thalamus, and in the cortex. But do such copies tell us anything about seeing? A Leica camera makes far better copies than does the retina, but does it see?

The copy theory (Figure 8.10) was invented and remains popular because it seemed to explain how experience changes us; we react to the world by copying it, as a camera or a computer copies and saves things. Skinner (1987) proposed a far more apt metaphor for the way in which experience changes us. He suggests that, instead of being like a computer saving its input, we are more like storage batteries. We put electricity in the battery and we take electricity out of the battery. But there is never any electricity *in* the battery.

When we put it in and take it out, we are changing the battery, much as experience changes us. In both cases nothing is copied.

Mental Way Stations In Skinner's view, the "little man," or "inner man," as he calls it, has been pressed into service to do much more than just our perceiving and reacting. He is actually a summary term for all of the unfinished causal sequences that abound in explanations of our behavior. It is he who has memory, ambition, aggressiveness, self-control, and all of the other mental way stations that have been proposed. He has moods, personality disorders, cognitions, and expectancies; in short, he is the embodiment of all of the intervening variables that Skinner believes act to obstruct a real understanding of our workings.

Skinner (1980; 1989) examined the etymology of some common English words that refer to inner states, or mental way stations. Interestingly, their original referent was to external things; they were descriptors of activity. Presumably, their meanings changed over time in such a way that they came to refer to internal states. Here are some examples:

Term	*Original Meaning*
wish	to strive for
yearn	to ask for
envy	to emulate
scorn	to humiliate (as by removing horns)
sad	to be sated

Other examples are easy to find. To the Greeks, *idea* meant "an image or likeness," and *glad* originally meant "shiny or bright." How many other terms originated as referents to observable situations, objects, or actions and eventually came to refer to inner states?

Skinner believes such mental way stations to be harmful, not because the things to which they refer do not exist, but because they represent unfinished causal accounts. When we describe someone's actions as due to a bad mood, we have only applied a name to certain kinds of behavior that we classify as representative of a bad mood. Why the behavior occurs is left unexplained. Similarly, when we recall or fail to recall a name, referring to a good or bad memory adds nothing to the observation. As discussed earlier, such naming is often treated as explaining, yet the observed behavior has merely been referred to something imagined to be inside the individual. The factors leading to our ability to recall or not (that is, to have a good or bad memory) lie in the past experience of the individual, as well as in aspects of the present situation.

Secondly, assuming such way stations often requires that people can accurately report subjective experience and that they be aware of factors influencing them. Masses of research, as well as daily experience, shows that this is not true, as we have noted already. "Inner man" remains a mystery to most of us, as well "he" should.

Phylogeny and Ontogeny

Skinner presents his position on the influences of heredity and environment in "The Phylogeny and Ontogeny of Behavior" (1966b), which begins with a quote from Blaise Pascal:

> "Habit is a second nature that destroys the first. But what is the nature? Why is habit not natural? I am very much afraid that nature is itself only first habit as habit is second nature."

The French physicist was suggesting that the distinction between effects of environmental experience (habit) and heredity (nature) is misplaced. Rather than assume that habit replaces instinct, it seems more natural to view learning and heredity as depending on similar processes; that is, habit is natural and nature is first habit.

Skinner agrees with this interpretation. He points out that behaviorists realize that much of our behavior is inherited; for example, we do not have to learn to pump our blood or digest our food. If internal behavior is thus genetically determined, so might be much of our other behavior. But what does it mean to say that behavior is inherited?

Inherited behavior is commonly believed to have been selected in the course of evolution. It is chosen for its survival value from some pool of behaviors and it is defined in terms of some sort of unit, like the fixed action pattern. It may take intermediate forms in successive generations, as different gradations prove more useful in promoting survival. We should expect variation in behaviors (and structures) as the selection pressures change.

All these characteristics of genetically selected behaviors apply as well to operant behaviors "selected" during an individual's lifetime (that is, in ontogeny). One need only substitute the phrase *contingencies of reinforcement* for *natural selection*! Thus, operants are chosen either by nature or human design. They are defined as a unit by the class of responses producing a common consequence and they come under the control of discriminative stimuli. And when reinforcement contingencies are altered, we find increased variation in operant behavior until the new contingencies select the final form of behavior.

Along these lines, recent critics of behaviorism have pointed to exceptions to the laws of learning. Among these exceptions is the discovery by the Brelands that the performances of animals trained for commercial purposes (for example, television and movies) were often disrupted by the intrusion of species-specific behaviors. The influence of phylogenic behavior on learned behavior is by no means exceptional, and there are cases in which the reverse occurs. Skinner (1977) noted that one bird in Project Pigeon came to peck at targets on the screen very rapidly and it also came to peck too rapidly at its food to allow it to eat. If this had not been noticed the bird could have starved. In this case, the so-called misbehavior involves the influence of learning on innate behavior.

In these and other cases there is no point in attempting to clearly separate learned and innate behavior. In many cases, such as aggression, territoriality, and courtship, there are varying contributions by innate and learned factors. In the case of aggression it is obvious that selection pressures in phylogeny could favor aggressive individuals, and it is obvious that aggressive behavior is often reinforced as an operant. A single case of aggression could depend on either or both factors.

3. CRITICISMS OF SKINNER

Skinner is extremely well known, largely because of his popular writings and autobiography. He has come to represent behaviorism in the minds of many, and thus be has been criticized as if he held the views of all other behaviorists, both living and dead. Most frequently, he is charged with holding the mechanistic S-R position of Hull. This is not to say that Skinner is not criticizable—far from it. Let us consider briefly the criticisms that are not legitimate, followed by those that are.

A critic who argues against the following is not criticizing Skinner; the real target is Hull and simple-minded S-R associationism:

1. Learning involves the establishment of specific S-R associations. Our goal is to analyze behavior into such units.
2. Reinforcers reduce biological drives.
3. Private experience is irrelevant.

This is *not* Skinner's view; *learning* for him is a poor term that refers to the formation of correlations between classes of behaviors and classes of stimuli. These are in no sense S-R associations; orderly relations between behavior and the world are something to be discovered. By the same token, reinforcers do not share any fundamental property aside from the fact that they may be observed to act as reinforcers under given conditions. Private experience is important and may be dealt with in the same manner as other behavior; experiencing *is* behaving (see Morris, 1988).

The following criticisms of Skinner are more legitimate. The first criticism is the assumed lack of concern for cognitive processes, especially in behavior therapy. The second is his relative neglect of instinctive behavior.

Cognitive Behavior Therapy

Skinner's critics have included behavior therapists; this may seem paradoxical, since the methods used by behavior therapists derive in large part from Skinner's writings. Behavior therapy is a general label that applies to a number of methods for changing behavior. An early example of behavior therapy is Watson and Raynor's (and Mary Jones's) demonstration that phobic reactions could be both instilled and removed (Chapter 4). Behavior therapy was not really popular until the 1960s, following the success of systematic desensitization (e.g., Wolpe, 1958), the application of operant methods to children's problem behavior (e.g., Bijou & Baer, 1961), and the treatment of mental patients (e.g., Ayllon & Azrin, 1968).

The procedures used in behavior therapy derive in large part from the basic methods used in classical and operant conditioning. Over the past few years therapists have found that reinforcement, extinction, shaping, conditioned reinforcement, and other so-called behavioral methods long used in animal research are very effective means for treating human problem behaviors (e.g., Martin & Pear, 1983). Judicious use of reinforcement and basic learning principles have been used to alter behavior in problem children, to foster beneficial new behaviors in mental patients, to help the obese lose weight, and to deal with a host of other problems. Behavior therapy works because the major principles of the classic learning theories are sound.

The majority of behavior therapists were trained at universities that stressed what Seligman (1970) called General Process Learning Theory (GPLT), an extreme and simplified view of conditioning and learning theory, because it served as a clear alternative to the more traditional therapies, which rely on explaining psychopathology through reference to intrapsychic forces and similar mechanisms. So, many behavior therapists learned their trade in such university training programs, which Mahoney (1977) called "behavioral boot camp." *Behavioral* here refers to this simple view, and in the institutions that make up that "boot camp" it is evidently taught that behavior therapists deal only with behavior and that behavior is movement in space. Good behaviorists do not take note of or deal with patients' private experience, imaginings, or thoughts! We have already discussed this view; it is what Skinner (1945) called methodological behaviorism.

What a state of affairs, given the position on private experience held by Skinner, as well as by Thorndike, Guthrie, and (even) Watson! But much of psychology had become GPLT by 1960 and it is not surprising that behavior therapists should accept it wholeheartedly. Part of the reason for this acceptance was that even GPLT worked better in the treatment of problem behavior than did competing therapies. But GPLT has its drawbacks, and these have not gone unnoticed by behavior therapists.

Many therapists became dissatisfied with the strictures of GPLT, which seemed to require that human patients be treated essentially as are the animals in Skinner boxes (a term actually coined by Clark Hull). But humans can do a bit more than other animals. Humans can be taught to monitor their own behavior, to imagine various things, to set goals, to think positively, to notice what they are thinking when a headache comes on, and so on. The precepts of GPLT taught in the training centers of only a few years ago did not allow therapists to take advantage of such capabilities, and thus a rebellion of sorts occurred.

One result of this rebellion was the founding in 1977 of the journal *Cognitive Behavior Therapy*, edited by Michael Mahoney. The title is somewhat misleading, since it does not mean "cognitive" in the sense that this term is normally used (see Chapter 1). Instead, cognitive behavior therapy simply allows mental activity back into conditioning and learning, as it was before the coming of GPLT.

If one views mental activity as behavior, as did Thorndike, Guthrie, and Watson, and as did Skinner, this new journal contains nothing that is radically new. For example, would earlier learning theorists object to a therapist teaching patients to dispel disturbing thoughts by saying "stop!" to themselves when the thoughts arise? Would they object to advising patients to note the thoughts they experience when a severe headache occurs so that such thoughts can be avoided? Is it not acceptable to imagine oneself as very obese when tempted by a sweet and then to picture oneself as slim after refusing it? Why shouldn't patients be taught to monitor their own behaviors and to set goals for changing it? Only a very restricted view of conditioning and learning would deny the possible usefulness of such methods. Yet, that is what learning theory had become. "Behavioral boot camp" accepted GPLT, just as many research-minded psychologists did since the 1950s.

Bandura has attacked GPLT, and his alternative, unlike most others of the cognitive behavior therapy movement, *is* radical (1977; 1986). During the 1960s, he and his colleagues demonstrated the importance of imitation in human affairs by showing the effects of modeling, or observational learning (e.g., Bandura, 1965). Children especially are apt to imitate models, particularly those who are older and of higher status in other ways (cf. Chapter 6, discussion of Miller & Dollard's earlier work). It seemed to Bandura that GPLT (which he called "connectionist S-R psychology" in 1977, 1986) could not account for imitation; he proposed that such behavior requires central cognitive processes to code experience and to retain it as symbols so that it can later guide behavior ("Am I imitating accurately?").

Bandura was also aware of some of the data described in Chapter 9, and he cited it as further evidence that central cognitive processes are necessary to explain even simple classical and operant conditioning. For example, Bandura often mentions Rescorla's (1967) demonstration of the importance of contingency in classical conditioning and Herrnstein's (e.g., 1970) evidence for the matching law. These findings tell us that organisms are sensitive to relations among events occurring over time; this is true and interesting. But Bandura argues that central processing mechanisms must be invoked to explain such things, and he believes they are necessary to explain modeling.

This is a classic case of using the complex and vague (central cognitive processors) to explain the simple and clear (operant and classical conditioning). Beidel and Turner (1986) criticized Bandura and cognitive behavior therapy in general for this mistake. They charge that cognitive behavior therapists reject a simple and outdated S-R theory in favor of a so-called cognitive model. They then propose that cognitions be changed so that behavior will be changed.

Beidel and Turner examined the methods used by a variety of cognitive behavior therapists, including Bandura, and concluded that there is really nothing new or nonbehavioral in any of them. And, when methods are successful, it is always the case that the change in observable behavior is responsible. If anything, the change in cognitive activity accompanies or is produced by the behavior change, rather than the other way around.

Neglect of Species-Specific Behavior: Misbehavior?

Keller Breland worked with Skinner at General Mills during the 1940s; Marion, who was to become his wife, was a research technician. They established Animal Behavior Enterprises in Hot Springs, Arkansas, and successfully trained animals for commercial purposes. By 1961 they had trained more than 6,000 animals belonging to 38 different species to perform a variety of tricks; in 1951 they published a paper describing their methods and applications.

But the Brelands are known today largely for a paper published in 1961, in which they describe difficulties they encountered over the years. Oddly enough, they strongly criticized the sufficiency of the theory that had guided them over the years, arguing that the instinctive behaviors of the animals they worked with caused problems for them. Such instinctive behaviors were paid little attention in the theory of conditioning and learning which Skinner had taught them, and they felt that this was a serious failing of behavioral theories in general.

Although Skinner's *Behavior of Organisms* (1938) included the major methods that the Brelands had employed successfully, it neglected the instinctive behavior of animals, which often appeared to provide difficulties. This amounted to "misbehavior" in their view, and their 1961 report was thus entitled "The Misbehavior of Organisms." Here are some of their examples of misbehavior.

The "dancing chicken" involved a hen that walked three feet to a loop of string, which she pulled, thereby starting a four-note tune playing. She then stepped onto an eighteen-inch disk, which rotated as she stood on it. As the disk turned, she vigorously scratched and after fifteen seconds she was fed. As you may imagine, observers might be led to believe that the hen was turning on the juke box and then dancing, a cute trick.

It is a cute trick, but it is not what the Brelands originally had in mind. Their initial purpose was to train the chicken simply to stand still on the disk for fifteen seconds, followed by food reinforcement at the end of that period. But the chicken developed the annoying habit of scratching with her feet on the disk, just as a chicken scratching for grain or grubs might do. The behavior that normally precedes eating intruded on the behavior the Brelands were trying to establish. To salvage the situation, they took advantage of the intruding instinctive scratching and came up with the "dancing" chicken.

In a related case, a chicken was to pull a loop that released a plastic capsule containing a toy or charm, which slid down a chute. The bird was then to peck out the toy or charm through an opening for a human observer who (presumably) had deposited money in the slot. (Such devices, using mechanical rather than real chickens, used to be fairly common in the entranceways of discount stores.) As training progressed, the Brelands encountered problems that they could not overcome and that may account for the mechanical chickens. Their real chickens frequently grasped the capsule in their beaks, shook it, and banged it on the floor, much as a chicken does when trying to break open a seed pod or kill a grub or larva! Here again, instinctive feeding behavior intruded, disrupting the behavior that the trainers were trying to establish.

Other forms of instinctive behavior were evident in the case of a raccoon. The animal was trained to deposit a wooden coin in a small box; although it became difficult to get the raccoon to drop the coin, it eventually learned to do so. The trouble really began when the raccoon was required to drop two coins in the box. It then did what raccoons do with food: It rubbed the coins together for minutes at a time and dipped them into the box, just as it would dip food into water to soften it. The coins, which had become associated with food, came to be treated as food and the project was scrapped.

A pig showed similar "misbehavior," despite the fact that the Brelands had long found pigs to be excellent and rapid learners. This pig was to deposit a large wooden coin in a piggy bank, a task it learned quickly; soon it was carrying four and five coins to the bank for each food reinforcement. Then, after weeks or months of training (depending on the specific pig), the pig began to slow down and, finally, to drop the coins and root at them. With time, this misbehavior became worse and worse. Instinctive eating behavior disrupted the performance which the Brelands were trying to establish.

Similar cases of misbehavior occurred when the Brelands tried to train cows to kick, cats to move away from the feeder, and whales to refrain from swallowing inner tubes that they were supposed to manipulate for food reinforcers. It was clear that animals come with species-specific behaviors and that these place limits on what animals may be easily trained to do.

The Significance of Instinctive Drift The Brelands (1961) wrote that these things were a shock to them and that they represented a "clear and utter failure of conditioning theory." Why should animals conditioned to perform a specific learned response show a gradual *instinctive drift* to natural food-getting behaviors (see Table 8.3)? The misbehaviors delayed or prevented the delivery of food and required a lot of energy, thus violating the "law of least effort" (a seldom-mentioned law, which was really stressed only by Tolman). Misbehavior showed that species differences are not insignificant and that all responses are not equally conditionable to all stimuli. Instinctive behaviors do exist and one must be at least somewhat mindful of them.

But didn't everyone already know this? Was there a learning theorist who actually believed that we may take any organism, place it in any situation, and reinforce any behavior, without considering the species with which we are dealing? Thorndike surely would not believe that, nor would Watson, Guthrie, or Tolman. It is hard to believe that even Hull and his followers would be so obtuse!

Yet, by the 1950s and 1960s the majority of experimental psychologists did believe that one could ignore species differences. For years the favored subjects of psychological research had been rats and pigeons and whatever laws applied to the pressing of levers and the pecking of keys were assumed to be generalizable to all behaviors of all organisms at all times. The Brelands were the first animal psychologists to point out that this was not the case, and the revolution they initiated continues to the present.

Remember, however, that the Brelands were terribly successful in training animals, even though they had firmly believed the simple view of conditioning which they attacked in their paper until the cases of instinctive drift mounted to such a level that they could not be ignored. They admitted that they did not mean to "disparage the use of conditioning techniques," since these methods had served them well for many years. But one must take species-specific behaviors into account. Just after the paper was accepted for publication, Keller Breland wrote the following in a letter to Skinner:

> "We have been concerned here for a number of years with the area of observations commonly called instinct (a rosebush by any other name is just as thorny) and an article discussing the type of problem one encounters in this area will soon appear in the *American Psychologist*. That is the occasion for this letter since, after looking at the galley proofs, it occurred to us that it might convey impressions not intended. Perhaps we did not state strongly enough our feeling as to the efficacy of operant conditioning in the control of organisms. This conviction is so 'old hat' with us that I am afraid that we sometimes forget that it is not shared by all American psychologists." (Skinner, 1977, pp. 1006–1007)

TABLE 8.3 Instinctive Drift among the Brelands's Animals

Training	*Outcome*	*Reason*
chicken fed after fifteen seconds on rotating disk	scratched disk vigorously	food-getting (scratching) behavior intruded
chicken fed after pecking capsule through opening	grasped and shook capsule	eating behavior intruded
raccoon fed after depositing two coins in a box	rubbed coins together and "dipped" them	food-softening behavior intruded
pig fed after depositing coins in a bank	dropped and rooted at coins	food-getting (rooting) behavior intruded

Skinner added the following remark:

> His letter concluded: "Viva la operant! (Such a pregnant notion as the operant must be female.)" It was, of course, necessary for the Brelands to use the phylogenic repertoires of the various species with which they worked (and they did so in a most ingenious way), but they continued to be dependent upon the shaping technique we discovered in Project Pigeon in developing that behavior for commercial purposes. (p. 1007)

In the same article, Skinner noted that misbehavior may occur in the reverse fashion; learned behavior may intrude upon and disrupt instinctive behavior. This occurred during the ORCON project, in which pigeons were trained to peck at targets visible on a large screen in front of them, thus guiding a missile. One bird pecked at the targets extremely rapidly and this high rate of pecking occurred whenever it pecked. This made it difficult for the bird to eat, and it lost weight. It had to be trained to peck more slowly, so that its learned high rates of pecking would not disrupt its instinctive eating behavior! Despite the discovery of misbehavior, the Brelands found Skinner's methods generally sound, as we will see below.

4. OTHER ASPECTS OF SKINNER'S THEORY

Examples of Skinner's theory in operation usually refer to instances of behavior therapy, since that work is best known to the public. Yet, behavior therapy really illustrates the influence of behavioral methods in general, not just Skinner's version of behaviorism. More appropriate examples lie in the training of animals for commercial purposes and in Skinnerian interpretations of self-control, avoidance learning, the effect of contingencies per se, and the Columban simulations.

Performing Animals

While a graduate student of Skinner's, the same Keller Breland now known for his contribution to the misbehavior paper saw the possibilities of using operant conditioning methods in the commercial training of animals. He and his wife Marion had many clients; for example, Quaker Oats, General Mills, the U.S. Army and Navy, the Veterans Administration, the Department of the Interior, Marriott's Great America Parks, Six Flags Over Mid-America, and Opryland. Part of their advertising reads as follows:

> Want a show or exhibit featuring trained common domestic animals? Since 1947, ABE has produced shows featuring barnyard animals.
> Want a dolphin or sea lion to perform in the open ocean? Since 1963, ABE has pioneered in open ocean training of marine mammals.
> Want a seagull, raven, pigeon or other bird to fly free and perform at your command? Since 1965, ABE has freeflown birds over distances of several miles.
> Want long-range invisible control of dogs, cats, or other species of land animals? ABE has developed and tested techniques for controlling certain land animals at distances of several hundred yards.

Want to know how they do this? Breland and Breland (1951) reveal the basics.

Animal Behavior Enterprises will send you an animal version of a video game for a price. Their exhibits include Skyline Hen (the tightrope walking chicken), the fortune-telling chicken, the capsule vending chicken, the baseball-playing chicken, the piano- or drum-playing chicken (or duck), the kissing bunny, the gambling rabbit, the tic-tac-toe chicken, the basketball-playing chicken, and much more. The Brelands are now better known for their 1961 paper emphasizing "misbehavior." But their commercial success testifies to the effectiveness of Skinner's training methods.

The Columban Simulations

Through the years, Skinner had argued that anything that we may think important in psychology may be explained in terms of contingencies. In works such as *Science and Human Behavior* (1953) and *Beyond Freedom and Dignity* (1971), the term *contingencies* is used almost as a literary device to refer to the fact that something is learned. Hence, he attributes thinking, selective perception, the concept of self, beliefs, and many other things to the effects of contingencies, without going very deeply into how these contingencies operate and how they produce these products. But how could he illustrate how such things happen? The so-called Columban simulations are one way of doing this. (The term *Columban* refers to the order *Columbae*, the group of birds that includes pigeons and doves.) The Columban simulations are attempts to show how contingencies can shape behavior indicative of higher mental processes. Two simulations show how the behavior of pigeons may seem to suggest that they have a genuine concept of self and that they are capable of symbolic communication. They even leave themselves notes, seemingly foreseeing that their memory may fail.

The nature of self-concept was examined by Epstein, Lanza, and Skinner (1981), using a method devised by Gallup (1970) for assessing the existence of a self-concept in monkeys. Gallup found that chimpanzees, marked with red dye on an eyebrow, showed that their mirrored image was recognized as an image of themselves. The chimps examined the red area on their own bodies, while seeing it only in the reflection from the mirror. Monkeys were unable to use the mirror image in this way; in one case, a rhesus failed to show any recognition that the monkey in the mirror was itself after as long as 2,400 hours of opportunity to do so! Gallup concluded that only humans and the very highest subhuman primates are capable of recognizing a mirror image of themselves, which means that only they have (or are capable of having) a self-concept. This was a rash conclusion.

Epstein, et al. felt that the ability to respond to a mirrored image of oneself could be demonstrated in organisms far lower than chimpanzees. To show this, they trained pigeons to peck at a blue spot placed in various locations in their box, with pecks reinforced with access to grain. They went on to present the blue spot on the end of a clear Plexiglas rod, so that the birds would be used to pecking at it in various locations in space. Finally, they placed press-on blue spots on the birds' bodies, in places where the animals could see them. The birds soon reliably pecked at blue spots, whether on the wall, in space, or on their legs.

Then came the interesting part: Would a pigeon recognize a mirrored image of itself and thus show evidence for a self-concept, according to Gallup's definition? The blue spot was placed on the pigeon's breast and a bib was placed around the bird's neck. Mirrors were set up outside the Plexiglas box in which the birds worked. The bird could see the blue spot reflected in the mirror image, but if it bent to look at its breast, the bib covered the spot. Thus, when the bird stood upright, the blue spot was visible in the mirror's reflection. If the bird looked down, the bib covered the spot. The only way in which the bird could see the spot was in the mirror's reflection (see Figure 8.11).

The photographs in Epstein, et al. clearly show the birds pecking downward at the bib, over the point at which the spot was hidden. When the mirror was covered, the birds did not peck at this spot, showing that the bird was not simply pecking where it felt the spot to be. Do pigeons recognize themselves in mirrored images? Yes! Do they therefore have a self-concept, as we think of it? Don't be silly! And neither did Gallup's chimps. The ability to recognize oneself in a mirror image is not a good criterion for a concept of self, unless we extend the ability to have a self-concept to pigeons. At least, this was the conclusion of Epstein, et al.

A second experiment was done in response to demonstrations of animal communication through symbols, which have appeared in the popular press. Several projects have claimed to show human-like use of language by chimps and gorillas, such as Washoe, Koko, and Sarah. The animals do not actually speak, of course; they communicate through the symbolic gestures of American Sign Language (ASL) or manipulate plastic symbols mounted on a magnetic board. In other cases

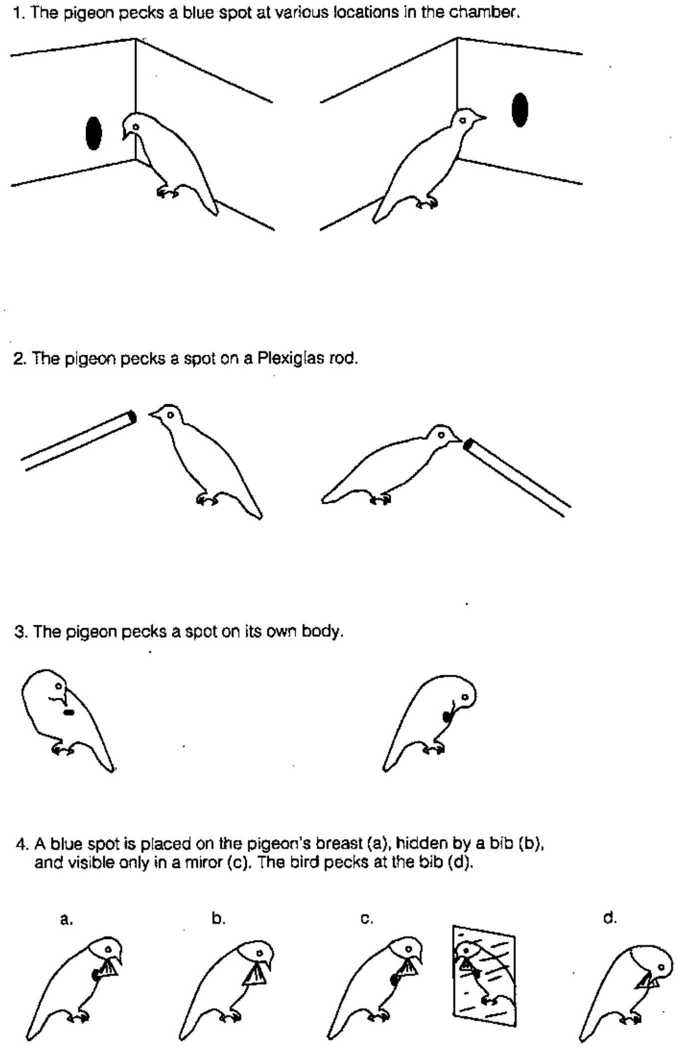

FIGURE 8.11 Self-concept experiment by Epstein, Lanza, and Skinner (1981) using pigeons.

the animals use symbols mounted on a computer keyboard. According to the investigators involved, these animals actually have a rudimentary use of language. They do communicate, just as we do. (Interestingly, they often seem to adopt other human behaviors, such as sleeping in beds, rather than in trees, and eating with plates and silverware.)

Savage-Rumbaugh, Rumbaugh, and Boysen (1978) published an account of chimpanzee communication, in which one animal asked another where food was hidden and the other answered the question. This was described as a marvel in primate intelligence. What other animal had the innate capacity for language to allow it to accomplish such a feat? Chimpanzees do resemble humans to a degree, and most people are therefore willing to share such an intrinsically human faculty as language use with them.

But these chimps did not spontaneously begin to converse with symbols; they were taught to do so. Given that training is necessary, is it still the case that only the highest primates may benefit from training? Could the same sort of training lead to communication in pigeons? If such lowly animals could be trained to duplicate the chimps' behavior, we might then conclude that all of us learn to communicate in such a way but that no one bothers to teach animals, so most of them do not talk. Or we might conclude that the special training given the chimps produced only the appearance of true language use, since few would grant the lowly pigeon the ability to use language.

Savage-Rumbaugh, et al., had taught chimpanzees to name foods by pressing buttons marked with symbols corresponding to various foods. They then were taught to request foods. Finally, one chimpanzee watched as food was hidden; the experimenter then asked it to name the hidden food by pressing the appropriate button when another chimp was present. If the second chimp then correctly asked for that food, both were fed. Later, both chimpanzees were observed to spontaneously "ask" each other for food. Remember that in all of these cases the naming and asking was done by pressing buttons. Nonetheless, here was an example of language use in higher primates.

Epstein, Lanza, & Skinner (1981) showed that such behavior was not unique to chimpanzees. In a charming piece, which described pigeons' behavior in the same warm terms that the chimpanzee behavior was described by the Rumbaughs and others, Epstein's group showed that pigeons were capable of symbolic communication. Two pigeons named Jack and Jill worked in a chamber divided down the middle by a clear Plexiglas barrier. Jack began by pecking a key with *WHAT COLOR?* written on it. Jill then looked at a panel of three lights recessed in the upper right corner of her side of the box and covered by a curtain. She had to stick her head behind the curtain to see which of the three lights—red, green, or yellow—was lighted (see Figure 8.12). Jack could not see this three-key display from his side of the chamber.

Jill then moved to a vertical column of three keys on the left half of her chamber, which Jack could see. These were marked *R*, *G*, and *Y*, and Jill pecked the one corresponding to the hidden color that was lighted. When the key was pecked, it was illuminated and stayed illuminated. Jack looked at it and pecked a key on his side marked *THANK YOU*, at which time Jill was automatically fed. Jack then pecked the appropriately colored key on a horizontal display on his side, after which he was fed. Jack then requested another color name by pecking the *WHAT COLOR?* key and the cycle began again.

The final performance was remarkably like those reported by the Rumbaughs and other primate researchers. One animal "asked," another "answered," the first "thanked," and both were fed. After a description of the experiment, which was cast in the somewhat nonobjective terms used by the primate researchers, Epstein, et al., ended with:

> We have thus demonstrated that pigeons can learn to engage in a sustained and natural conversation without human intervention and that one pigeon can reliably transmit information to another entirely through the use of symbols. (p. 696)

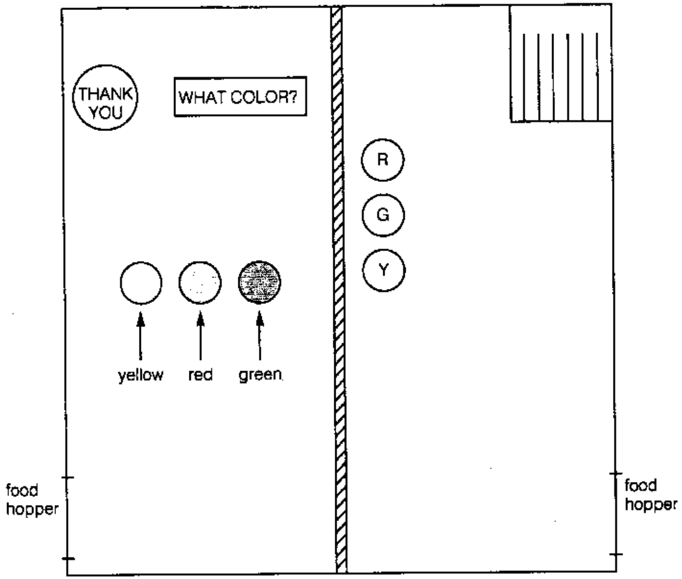

FIGURE 8.12 Panel of chamber for Jack and Jill, the pigeons in Epstein, Lanza, and Skinner's symbolic communication experiments.

They continued with the following:

> It has not escaped our notice that an alternative account of this exchange may be given in terms of the prevailing contingencies of reinforcement. (p. 696)

They go on to point out that they did *train* the birds to do this, using standard fading, shaping, chaining, and discrimination procedures. "A similar account may be given of the Rumbaugh procedure, as well as of comparable human language," they added.

After this experiment, Jack and Jill were given experience in both roles. Each learned to "ask," choose the communicated color, and "thank," as well as to look and communicate. The birds were then placed, one at a time, in the chamber, without the partition. To be fed, all the bird had to do was look at the hidden color (as Jill had done) and go to Jack's side and peck the same color. No one had to ask "what color?" or say "thank you," or peck the *R*, *G*, or *Y* keys.

But an odd thing happened. Jack, in going from the hidden colors recessed in the upper right wall of the chamber and covered by the curtain, sometimes forgot the color that was lighted. By the time he reached the left side of the chamber, he didn't know where he was to peck the color on the keys mounted there. How do humans solve such a problem? How do we guard against such a lapse in memory? We leave notes for ourselves!

Epstein and Skinner (1981) showed photographs of Jack "leaving a note," by pecking the *R*, *G*, or *Y* keys on the right half of the box, on his way to the keys on the left. The bird then went to the row of color keys on the left and clearly looked over his shoulder to see which of the labeled keys was lighted (that is, he looked at his

"note"). He then pecked the appropriate color key. Jill did the same thing, though less frequently; in both cases, the birds apparently were leaving a memorandum to which they referred when reaching the color keys on the left side of the chamber.

And they were doing just that, although they had not been trained to do so! When the color keys on the left were lighted only after a delay, which strained the birds' memories, both animals reliably pecked the "reminder keys" so that they could look back at them. When a disruptor, such as a loud buzzer, went off between the time that the hidden color was seen and the color keys were lighted, both birds more often "left a note."

Does this mean that pigeons, which have no neocortex and only a shred of cerebral cortex, are as intelligent as chimpanzees and humans? No, it only means that when animals are placed in situations in which human-like contingencies are in effect we may expect human-appearing behavior, insofar as the animals are capable of it. The same occurs when humans are placed in human and nonhuman situations. Thus, we need not be surprised to find contingencies (in the concentration camp or in the ghetto) producing behavior that we are reluctant to accept as human.

SUMMARY

Skinner carried on the spirit of Watson's rendition of behaviorism, although he disagreed with what Watson felt to be the fundamental principles of conditioning. Like Watson, he felt that theories were not only unnecessary, but were an impediment to progress. For Skinner, a theory was a way of explaining behavior and experience that translated what we wish to explain to other terms, such as the language of neurons and of mechanical devices, such as information processors. Most damaging were those theories that referred to the common explanations of popular psychology, such as the thoughts, traits, and faculties we all refer to when we speak of people with a good memory or an aggressive nature. Such so-called explanations, like references to higher cognitive processes, do no more than name what we wish to explain. They are better dispensed with, so that we can get to the real explanations, which lie in the past history of the individual (or the species).

Skinner's papers during the 1930s brilliantly outlined a strategy for the discovery of behavioral units. The unit is the reflex, although his definition of the reflex is not what most of us think of when we hear the term. *Reflex* means "order," and the discovery of the physiological reflex was possible long before it was possible to trace the anatomical path involved. Skinner proposed that the static and dynamic laws of the reflex are a more reliable means for identifying reflexes than is the anatomical path, and that the same strategy that allows us to identify reflexes may be applied to the identification of conditioned reflexes and operants, although the latter occur in the absence of any obvious stimulation.

Reflexes, conditioned reflexes, and operants are always defined as correlations between stimulation and responding (or responding and stimulation, in the case of the operant) defined at that level of specificity at which orderly changes occur as the result of manipulations of third variables, such as hunger, emotion, and fatigue. Stimuli and responses are never spoken of as such; we always speak of classes of stimuli and responses. This makes Skinner's behaviorism far different from the Š-R associationism of Hull and his associates. The result of the reinforcement of an operant may lead to changes in behaviors that seem very unlike the behaviors that were reinforced explicitly; the identity of the response class need not be intuitively obvious before conditioning occurs. Similarly, the stimulus class need not include only obvious and specific stimuli; there is no reason why abstract concepts may not be learned, just as reactions to more concrete stimuli are learned.

The law of effect is an empirical law; there is no property that all reinforcers have in common, such as the ability to produce pleasure or to reduce biological drives. The law of effect is a useful tool, and its usefulness is extended whenever we show yet another behavior that is influenced by its consequences. There are so many things that may act as reinforcers under one or another condition that the attempt to find a single basis for their effect (such as the activation of reward centers in the brain) is doomed to failure.

All reinforcers increase the frequency of behaviors that produce them. Positive reinforcers do so because they are produced by those behaviors, and negative reinforcers do so because a behavior removes them. Punishers decrease the frequency of behaviors that produce them. Although Skinner (1938), following Thorndike and Watson, downplayed the effectiveness of punishers, work by Azrin clearly shows that punishers are effective, although their application may be tricky.

Skinner is largely responsible for the distinction between operant and respondent learning; although there is a difference between the procedures used in the two types of conditioning, he is responsible for the belief that two different learning processes underlie the two methods. This belief, which is inconsistent with his position on theories of learning, has recently been decisively refuted.

Operants come under the control of discriminative stimuli, and the effects of reinforcement generalize to other stimuli, a process Skinner called induction. Differential reinforcement leads to extinguishing of responding to nonreinforced stimuli, which is the basis for discrimination learning. Operant behavior may be shaped, using the method of successive approximations, and a sequence of shaped behaviors may be joined to form a chain. In the course of shaping, we restrict reinforcement to only certain members of an operant response class, such as bar presses of a given force or greater. The other members of the class extinguish and behavior becomes more variable, increasing the number of instances that can be reinforced. In time, we may have a distribution of responses totally different from that with which we began. Once a behavior has been shaped, we may reinforce its occurrence only in the presence of a given S^D; we may then use that S^D as a conditioned reinforcer for a second behavior. When we pair a second S^D with that behavior, we have a two-member chain. In theory, we may extend the chain indefinitely, although in practice there are real limits, especially with animal subjects.

Schedules of reinforcement specify the rule by which reinforcement is delivered. The basic schedules specify that either a fixed or a variable number of responses occur or that a fixed or variable time period pass, after which responding is reinforced. In addition to the basic FR, VR, FI, and VI schedules, there are many variations specifying time and/or response requirements, including discriminative stimuli, and sometimes comprising a combination of simpler schedules.

Examples of reinforcement schedules in daily life are few, but obvious daily life applications are not important. Rather, the schedules tell us important things about the relations between behavior and its consequences. Perhaps the most important thing that psychologists have learned is that orderly relations between the world and behavior appear more frequently when we consider behavior over extended periods of time. The matching law and the work of Neuringer and Chung, which are described in Chapter 9, represent two salient examples of this.

Skinner's treatment of motivation, emotion, and other such terms follows from his opposition to the use of intervening variables. If we think of all of the things we mean when we describe ourselves or others as "hungry," or as "angry," we can appreciate his reasons for stressing the use of independent and dependent variables, rather than intervening variables.

Skinner's methods have been applied to the study of concept formation, beginning in his wartime efforts to build a pigeon-piloted missile. Further work with pigeons, porpoises, and humans has shown that our concepts are probably all learned rather than present at birth, as Plato had suggested. However, there is growing attention paid to the rapidity with which animals learn concepts, suggesting that they may come prepared to acquire some concepts.

Unlike Hull and Tolman, but like Thorndike and Watson, Skinner considers private experience to be part of what a behavioral science must explain. According to his radical behaviorism, mental activity may be treated in the same manner as other activity, and private stimuli, such as those that produce pain, may be treated as other stimuli. Skinner believes that society teaches us to identify items of personal experience, just as it teaches us to identify objects in the world. It is harder for our teachers to identify our private experiences, since they must rely on our behavior to do so; consequently, we can never know our own experience as well as we can know the objects of the world.

In 1966, Skinner set forth his views on innate behavior. He suggested that the factors that change behavior during the course of an individual's lifetime are similar to those that change the bodily and behavioral characteristics of a species during the course of evolution. Variation, selection, and heredity act in essentially the same way in both cases, and it is a mistake to try to separate the influence of nature and nurture in dealing with specific cases. Had Skinner published such a paper twenty years before he did, we might have seen less interest in the exceptions to the laws of learning. But he did not.

Criticism of Skinner is often misplaced. Frequently he is criticized as though he were the distillation of all views of all people who have called themselves behaviorists. Most often, though, he is criticized as though he were Hull, because Hull's influence was so great and because so many of his students have kept his basic ideas alive. There *are* serious criticisms that do apply to Skinner, however. The most serious is his failure to appreciate the significance of his early thoughts regarding the unit of behavior. Too often he reverts to molecular analyses that convince critics that his views really are the same as those of Hull.

The power of his approach is well illustrated in the success of professional animal trainers, such as the Brelands, and in the series of papers describing the Columban simulations. In the latter series, pigeons have been trained to show behaviors that indicate self-concept, communication, and note leaving. Other researchers have shown that pigeons can show self-control and that rats can react to fairly subtle differences in sequences of events acting over prolonged periods of time; these topics will be covered in Chapter 9. These are the sorts of demonstrations that show the applications of Skinner's theory, and we could have benefited from them long ago.

GLOSSARY

Adjusting schedule Schedule of reinforcement in which the value of the interreinforcement interval or the response requirement changes as time passes since the last reinforcement. For example, an adjusting schedule could begin with a value of FR 10 and increase that value by ten every twenty seconds. This should lead to high response rates, since, in effect, the schedule penalizes low response rates.

Alternative schedule Schedule of reinforcement in which reinforcement depends on the passage of time *or* the fulfilling of a response requirement.

Chaining (Chained schedule) The joining together of a sequence of behaviors by a series of discriminative stimuli that also act as conditioned reinforcers. The occurrence of the first behavior produces the S^D for the second member of the chain, and so on. The appearance of the S^D acts to reinforce the first behavior and set the occasion for the second member, which leads to the appearance of the next S^D.

Component Basic unit of a multiple schedule, consisting of a specific S^D, schedule of reinforcement, and duration. Components are presented successively. One example of a component would be a two-minute period in which a green light was lighted and in which a VR 25 schedule was in effect.

Concept learning Discrimination learning in which the class of stimuli involved do not consist of specific concrete things like lights and tones. The concept may be "all four-legged creatures" or all true statements. We train concepts by presenting numerous examples of instances of the concept and reinforcing responses by our subject and by presenting noninstances without reinforcement.

Concurrent schedule Schedule of reinforcement in which two or more independent schedules are simultaneously in effect, with a lever or key corresponding to each. For example, two different VI schedules may be available, with two response keys present, one for each schedule. A subject may get into the habit of switching keys, since the longer the time spent on one key, the more likely reinforcement becomes for responses on the second key. Switches between keys therefore may be reinforced. To prevent this, a changeover delay (COD) is usually used with concurrent schedules. A COD prevents the receipt of reinforcement for responding on a key for some fixed time (a few seconds) after a switch from the other key to that key.

Conditioned reinforcer According to some usages, such as that of Hull, all reinforcers that do not reduce drives arising from bodily needs are assumed to be conditioned reinforcers; their power derives from an association with primary reinforcers. Skinner's view defines conditioned reinforcers as things that act as reinforcers because of some pairing with already-effective reinforcers. The latter may act as such for a variety of reasons, and the difference between conditioned and other reinforcers lies only in the fact that conditioned reinforcers act as such only after association with effective reinforcers.

Conjunctive schedule Reinforcement schedule in which reinforcement is delivered only after the passage of time and the completion of a response requirement. For example, 47 responses may be followed by reinforcement if two minutes have passed since the last reinforcement.

Contingencies Another word for *schedule.* When Skinner speaks of the contingencies responsible for the development or maintenance of a behavior, he refers to the requirements that govern the delivery of reinforcement. These may depend upon the passage of time, the occurrence of specific responses, the presence of specific stimuli, or a combination of these things.

Differential reinforcement Another way of describing the selecting of behaviors, depending upon reinforcement of some responses in the presence of some stimuli and the nonreinforcement of other responses in the presence of other stimuli. Shaping, discrimination learning, and all of our learned behavior may be viewed as the result of differential reinforcement.

Discriminative stimulus Operants usually are reinforced only in the presence of some class of stimuli, the S^D. Skinner believes that such stimuli come to "set the occasion" for reinforcement, thus acting as discriminative stimuli. He also believes that stimuli that elicit reflex behavior or conditioned reflex behavior operate differently, by eliciting the responses which they control.

DRH schedule Schedule of reinforcement that reinforces high rates of responding. A differential reinforcement of high rates schedule requires that some number of responses occur within a fixed time period if reinforcement is to be received.

DRL schedule Schedule of reinforcement that reinforces low rates of responding. A differential reinforcement of low rates schedule requires that no responses occur during some fixed period since the last response if reinforcement is to be received.

Dynamic laws Factors that influence the strength of a reflex, conditioned reflex, or operant response, which show their influence as the behavior repeatedly occurs. For example, the law of fatigue and the laws of conditioning and extinction are such laws.

Empirical law of effect Skinner's version of the law of effect, which holds that reinforcers need not share any properties except their ability to act as reinforcers. We discover that something acts as a reinforcer under given conditions; we do not concern ourselves over why it does so. (For example, we do not ask whether it reduces drives or if it is pleasure producing.) The law of effect is thus an empirical generalization. It is a useful rule that we have found to apply in a large number of cases, and its value and usefulness depend upon how far we may extend its application. Why reinforcers act as they do is an unanswerable question.

Fixed-interval schedule Schedule of reinforcement that requires that a set period of time pass, after which the first response produces reinforcement. Fixed interval is symbolized as FI.

Fixed-ratio schedule Schedule of reinforcement that requires that a set number of responses occur to produce each reinforcement. Fixed ratio is symbolized as FR.

Fixed-time schedule Schedule of reinforcement that provides reinforcement after fixed periods of time since the last reinforcement, independent of the subject's behavior. This schedule is indistinguishable from classical delayed conditioning. Fixed time is symbolized as FT.

Heterogeneous chain Chain of behaviors that are not of the same topography. For example, the rat Pliny performed a chain consisting of pulling a string, carrying a marble, and dropping it. (See *homogeneous chain.*)

Homogeneous chain Schedule of reinforcement composed of a chain of behaviors in which the behaviors in each member of the chain are topographically similar. For example, a chain FR 2 VR 8 schedule requires that two lever presses occur in the presence of one stimulus, which leads to a change in the S^D, and the pressing an average of eight times then leads to reinforcement. The response required in both members (for example, lever pressing) is the same, unlike in a heterogeneous chain.

Induction Actually means "generalization." Skinner used this term, after Sherrington, to refer to induced changes in responding to one stimulus, as the result of reinforcement for responding to a similar stimulus.

Instinctive drift Name given by the Brelands (1961) to the disruptive effect of species-specific (instinctive) consummatory behavior on the learned (operant) performances of their animal subjects.

Limited hold A requirement that may be added to a VI or FI schedule, such that a response must occur within some set period of time (the LH value) after the schedule has made reinforcement available. For example, on a FI one-minute LH two-second schedule, reinforcement is available for two seconds after the passage of a minute. If a response does not occur, that reinforcement is lost.

Matching law Also called the molar law of effect. The matching law holds that relative response rates match relative reinforcement frequencies. For an explanation of what this means and what it signifies, see Chapter 9.

Methodological behaviorism Name given by Skinner (1945) to the view that holds that we can deal objectively only with observable behavior and that mind exists, but cannot be meaningfully studied. Skinner opposed this view and labeled his contrary position radical behaviorism.

Mixed schedule Schedule of reinforcement in which two or more different schedules are in effect for set periods of time, these periods appearing in sequence. For example, a mixed FI 2 VI 3 schedule might alternate two-minute periods, in which FI 2 and then VI 3 schedules were in force. If these periods are signaled by discriminative stimuli, we have a multiple schedule.

Multiple schedule A mixed schedule in which different S^Ds signal the schedules that are in effect.

Negative punishment This occurs when a response decreases in frequency when its occurrence is followed by the offset of something. For example, a bit of misbehavior may be followed by the turning off of a television set. Like negative reinforcement, the behavior removes some stimulus. Like other cases of punishment, this produces a decrease in the frequency of the behavior which causes this consequence.

Negative reinforcement An increase in the frequency of a response that produces the offset of something. Like negative punishment, negative reinforcement involves the termination of something; like positive reinforcement, it leads to an increase in responding. When we pull a window shade to stop the sun from shining in our eyes, the negative reinforcement that results makes it more probable that we repeat that act the next time the sunlight annoys us.

Operant Class of responses that vary together in strength as a function of the consequences produced by members of the class. Pressing a lever acts as an operant, as does creative behavior, going to the store, and a myriad of other behaviors. We can identify an operant class only after we have observed that a given behavior is influenced by its consequences. It may take some time to identify most or all of the responses that make up an operant.

Operant conditioning The process whereby an operant class is shown to become more frequent (that is, to increase in strength) as a function of the consequences it produces. Thus, a rat may more frequently press a lever when presses are reinforced with food, and an infant may increase its emission of vocalizations when vocalizations are followed by praise and attention.

Positive reinforcer A consequence of behavior that produces an increase in the frequency of that class of behavior.

Punishment The decrease in frequency of an operant behavior as a function of the consequences it produces. Skinner argued that punishment is not a basic effect and that effects attributed to it depend upon side effects of the aversive events used as punishers. We now know that he was mistaken and that punishers seem to work in a way opposite to the effects of reinforcers.

Radical behaviorism Position described by Skinner in 1945 and 1963, which proposes a philosophy for a science of psychology, independent of specific theories of learning, whether Skinner's or anyone else's. According to this view, the entire subject matter of psychology may be treated as activity (behavior) and therefore mental activity is essentially the same in kind as physical activity; we may speak of thinking and seeing as behaviors, just as we do when we speak of walking and talking. Radical behaviorism argues against the usefulness and the existence of intervening variables and especially of internal copies of the world. Some writers have pointed out the similarities between this point of view and that of modern European phenomenologists, such as Merleau-Ponty and Sartre. There is little doubt that Skinner named this view after the radical empiricism of William James.

Reflex For Skinner, refers to order. The discovery of a reflex is the discovery of an orderly relationship between the world and behavior. Skinner extended a masterly analysis of the concept of the reflex to conditioned reflexes and operants in the 1930s.

Respondent Skinner's term for reflex and conditioned reflex behavior, which is elicited by an identifiable stimulus. *Respondent conditioning* is his term for Pavlovian conditioning, or conditioning of Type S.

Response class The set of behaviors that change in frequency together. This applies to reflex, conditioned reflex, and operant behavior. The response class may be composed of members that do not intuitively seem to go together. For example, aggression is no doubt a number of response classes, some controlled by external stimuli, and others by consequences. Each class may contain members discoverable only after long observation of behavior under a variety of conditions.

Schedule of reinforcement Rule by which reinforcers are delivered. The rule may include response requirements (as in ratio schedules), temporal requirements (as in interval schedules), or both. The first major analysis of reinforcement schedules and their effects was published by Ferster and Skinner in 1957.

Shaping Commonly used term for the method of successive approximations. This method involves the selective reinforcement of some subset of a class of operant responses. This leads to extinction of the nonreinforced members and a consequent increase in the variations of behaviors emitted. The requirement for reinforcement may be progressively restricted until the final product is a set of behaviors very different from the original behavior class.

Static laws In Skinner's early papers on the identification of reflexes, static laws referred to changes in responsiveness visible with single elicitations of the reflex. For example, we may see an increase in the magnitude of the response as we apply stronger eliciting stimuli. On a given occasion, a stronger stimulus produces a stronger response, irrespective of when the last stimulation was given.

Stimulus class The set of stimuli that may be shown to control a reflex, conditioned reflex, or operant class of responses. Like the response class, the stimulus class may be composed of elements that may not seem intuitively obvious. What we call "concepts" are names for stimulus classes.

Stimulus control The effect of external stimuli on behavior. The study of stimulus generalization and discrimination learning is now called the study of "stimulus control."

Tandem schedule Two or more schedules of reinforcement arranged in sequence, so that the requirement of one schedule must be met before the next schedule begins. Food or other reinforcement is delivered only after all schedules in the sequence have run. If different stimuli signal the successive schedules, it is a chained schedule.

Theoretical law of effect Attempts to explain the fact that reinforcers work as they do by postulating some underlying process. For example, Hull's suggestion that reinforcers work by reducing biological drives was a theoretical law of effect. One could suggest that all reinforcers promote survival, produce pleasure, or share some other characteristic. All such attempts have failed, and Skinner believed that we only waste time and effort by trying to work out a theoretical law of effect. His reasons are the same as those he uses to argue against theories of any kind.

Theory For Skinner, a translation of terms. For example, if we attribute intelligence to properties of the brain, information processing mechanisms, a "smarts" center (to use Stephen Gould's term), or the like, we are proposing a theory. More generally, theories involve the use of intervening variables, rather than the independent and dependent variables we should be interested in. Acceptable explanations refer only to the phenomena to be explained and the concrete conditions that influence them. The translation of these basic terms into hypothetical entities (habits, motives, and so on) should be avoided.

Token economy Method of psychotherapy originated by Ayllon and Azrin in the early 1960s. Patients, for whom other methods of therapy had failed, were reinforced for grooming, working, eating in an acceptable manner, and other activities.

Type R Skinner's term for behavior that is sensitive to its consequences—that is, operant behavior. Type R conditioning is therefore operant conditioning.

Type S Skinner's term for classical conditioning, in which the emphasis is placed on the eliciting CS rather than upon the consequences of the elicited behavior. Classical conditioning is conditioning of Type S.

Unfinished causal sequences Skinner's term for the common practice of explaining behavior and experience by reference to some hypothetical inner state or process. Hence, we may explain unruly behavior as the product of aggressiveness and the ability to recite well as the result of a good memory. In both cases, we have done no more than name the behavior involved and, unless we explain aggressiveness and memory, we are left with unfinished causal sequences, not real explanations.

Variable-interval schedule Schedule of reinforcement, symbolized as VI. This schedule provides reinforcement for the first response that occurs after some mean interval of time has passed since the last reinforcement. For example, a VI six-minute schedule would reinforce the first response occurring after an average of six minutes. Some interreinforcement intervals could be as short as a few seconds and others could be ten minutes long or longer.

Variable-ratio schedule Schedule of reinforcement, symbolized as VR. This schedule provides reinforcement after the completion of some number of responses. The number varies from reinforcement to reinforcement, and the value of the schedule is the mean number of responses required. Thus, a VR 25 schedule provides reinforcement after a varying number of responses, the average requirement being 25.

RECOMMENDED READINGS

Verplanck, W. S. (1954). Burrhus F. Skinner. In Estes, et al. (Eds.), *Modern learning theory*. New York: Appleton-Century-Crofts.

This is still considered the most perceptive evaluation of Skinner's theory. It was written by a former colleague.

Skinner, B. F. (1972). *Cumulative record* (3rd ed.). New York: Appleton-Century-Crofts.

This is a collection of Skinner's papers selected from publications since 1930. The reader will find many topics covered, dealing with theory, application, and others of personal interest to Skinner.

Skinner, B. F. (1974). *About behaviorism*. New York: Knopf.

Skinner explains what is meant by radical behaviorism in this clear and entertaining little book.

Skinner, B. F. (1983). *A matter of consequences*. New York: Knopf.

This is the third and last volume of Skinner's autobiography, covering the period beginning with his return to Harvard in the late 1940s to the 1980s. It is more a professional than a personal autobiography and provides many insights into Skinner's treatment of psychology.

Skinner, B. F. (1984). Canonical papers of B. F. Skinner. *Behavioral and Brain Sciences*, *7*, whole number 4.

This special issue of this journal is composed of six statements by Skinner dealing with basic issues, each followed by commentators' statements. Skinner replies to approximately 150 such commentators.