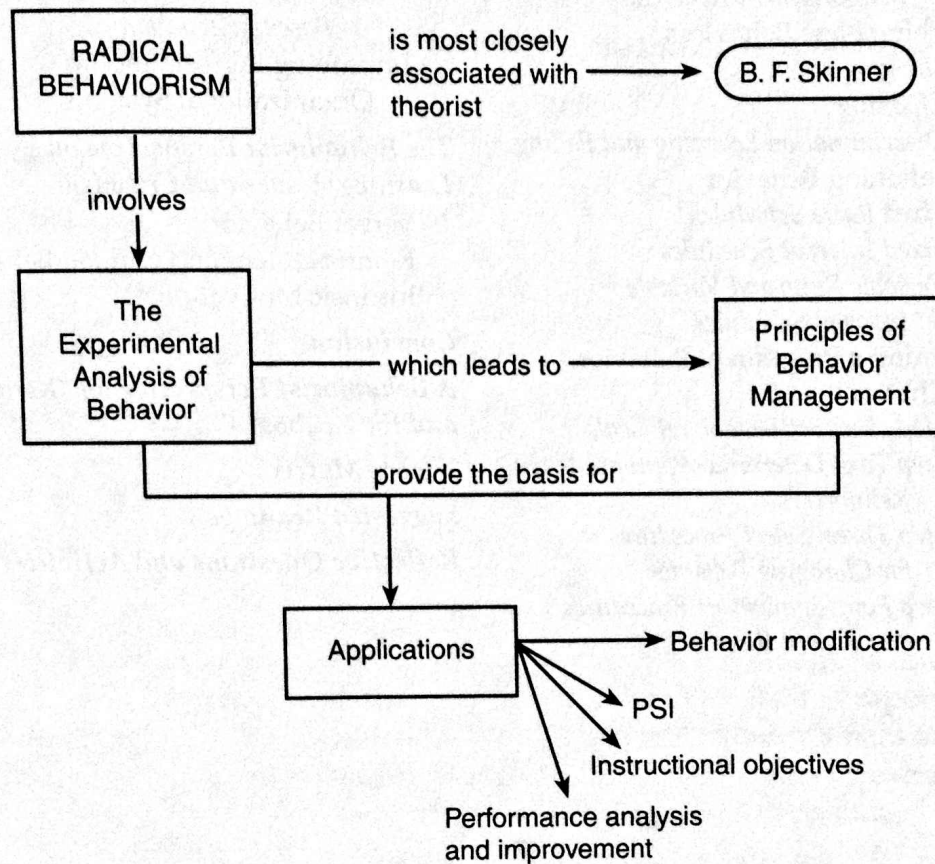


Part II: Learning and Behavior

2

Radical Behaviorism



Step Five: Evaluate Progress and Revise as Necessary

Contributions of Behaviorism to Instruction

Changing Behavior Through Behavior Modification

Managing Learning and Behavior in Instructional Systems

Classroom Management

Instructional Objectives

Contingency Contracts

Personalized System of Instruction (PSI)

Teaching Machines to Computer-Based Instruction

Improving Performance in Organizational Systems

The Behaviorist Perspective on Learning: Issues and Criticisms

Verbal Behavior

Reinforcement and Human Behavior

Intrinsic Motivation

Conclusion

A Behaviorist Perspective on "Kermit and the Keyboard"

Theory Matrix

Suggested Readings

Reflective Questions and Activities

The Experimental Analysis of Behavior
 Respondent and Operant Behavior
 Contingencies of Reinforcement
 Principles of Behavior Management
 Strengthening or Weakening Operant Behaviors

Operant Behaviors
 Strengthening a Response:
 Positive Reinforcement

Strengthening a Response:
 Negative Reinforcement

Weakening a Response: Punishment

Weakening a Response:
 Reinforcement Removal

Teaching New Behaviors
 Shaping

Cloning
 Discrimination Learning and Fading

Maintaining Behavior
 Fixed Ratio Schedules

Variable Ratio and Variable
 Interval Schedules

Planning a Program of Behavior
 Change

Step One: Set Behavioral Goals
 Step Two: Determine Appropriate Reinforcers

Step Three: Select Procedures for Changing Behavior

Step Four: Implement Procedures and Record Results

Consider these scenarios.

• *Department X*

As part of an organization-wide quality improvement effort, the head of a department sends her office manager and staff to training on the use of electronic mail. In addition to procedures such as logging on to the organization intranet to receive and send mail, the training included procedures for accessing the World Wide Web and locating and downloading information from the department's web page.

Within weeks after the training, the office manager routinely checks and reads her e-mail messages, but she continues to use paper memos and office mail to correspond and conduct business. The department is large, and some days the office manager puts as many as a half-dozen memos—each only a few lines long—in people's mailboxes.

• *Mr. Tanner's Class*

Mr. Tanner's fourth grade class reflects the ethnic diversity of his rural neighborhood—part Anglo American, Native American, Inuit, and African American. There are about as many boys as girls, and the range of their abilities is considerable. As in most classes, the students work at different rates, a few rarely participate in group assignments, and some seem to chronically misbehave.

Posted on the bulletin board in the class are these five rules (Evertson et al., 1994):

1. Be helpful and polite.
2. Respect the property of others.
3. Listen while others speak.
4. Respect all people.
5. Obey school rules.

At the beginning of each school year, Mr. Tanner discusses the rules with the students, and together he and the students determine what the consequences will be for failure to follow them.

• *Boot Camp*

Recruits are quick to learn at Boot Camp, USA. Besides doing assigned chores in their barracks, they get in shape with daily 5-mile runs and calisthenics. They learn to load, fire, dismantle, and clean their weapons. Performing their duties well can lead to privileges such as a day's pass to town, but breaking the rules inevitably leads to such consequences as extra push-ups, more miles to run, or forfeited time off.

It may not seem at first that these scenarios have much in common. Yet all of these situations illustrate (or will, with some fleshing out) the basic tenets of radical behaviorism.

The notion of behaviorism was introduced into American psychology by John B. Watson (1913). Watson promoted the view that psychology should be concerned only with the objective data of behavior. The study of consciousness or complex mental states, Watson argued, is hampered by the difficulty of devising objective and functional indicators of these phenomena. At some point, one is forced to consider the facts of behavior. These, at

least, can be agreed upon because they are observable by anyone. To illustrate, suppose, in the scenario Boot Camp, that Private Johnson draws barracks duty one week, which consists of mopping and waxing the barracks floor each day. For completion of the task with no demerits (which means those floors were spotless!), she earns commendations every day and is awarded a pass to go off base Friday night. What can we conclude from this scenario? Did Private Johnson do such a good job because she looked forward to a fine meal at a local Italian restaurant instead of army food for one night? Or maybe she just takes pride in her work. The fact that any number of inferences are possible when we attempt to understand Private Johnson's mental state and the reasons for her behavior is precisely the problem Watson noted. Stick to the facts of behavior: She completed the assigned task, the results were spotless, she earned commendations, she was awarded a pass.

B. F. Skinner, a major proponent of radical behaviorism, followed Watson's lead in emphasizing behavior as the basic subject matter of psychology (Skinner, 1938, 1974). But Skinner's work differed in a fundamental way from Watson's and others' work contemporary with and immediately following Watson. In the early days of behaviorism, the concept of association permeated theories about learning. It was assumed that a response (R) came to be established, or learned, by its association with an environmental stimulus (S). Edwin R. Guthrie, for instance, believed that, "Stimuli which are acting at the time of a response tend on their reoccurrence to evoke that response" (1933, p. 365). This has been called one-trial learning because, according to Guthrie, it is the very last stimulus before a response occurs that becomes associated with that response.

Whereas Guthrie's ideas were never fully elaborated, Clark L. Hull's S-R theory of behavior became "fearsomely complex" (Leahy & Harris, 1997). Hull believed that responses become attached to controlling stimuli, but some of these stimuli must be internal because it was not always possible to observe an external stimulus for all responses. Thus, Hull proposed intervening variables such as habit strengths and argued that observed behavior was a function of these as well as environmental variables such as degree of hunger (drive), size of reward (stimulus-intensity dynamism), and so on.

Finally, E. C. Tolman believed that behavior was guided by purpose, which led to his being called a purposive behaviorist. According to Tolman (1948), organisms do not acquire S-R bonds simply by contiguity or reward; they selectively take in information from the environment and build up cognitive maps as they learn. This helped to account for latent learning, in which rats who explored a maze for several trials found the food on a subsequent trial as quickly as rats consistently reinforced in the maze.

Tolman's cognitive maps and Hull's habit strengths, however, smacked of mentalism to Skinner. One cannot directly observe cognitive maps in a rat's mind; they must be inferred from the rat's behavior. Likewise, one



B. F. Skinner

cannot directly observe habit strengths; they must be inferred from the rat's persistence in a learned behavior. Skinner argued that such inferences were neither necessary nor desirable.

B. F. Skinner's approach to the psychology of learning was to set out in search of functional relationships between environmental variables and behavior. In other words, he believed that behavior could be fully understood in terms of environmental cues and results. Cues serve as antecedents to behavior, setting the conditions for its occurrence. Results are the consequences of behavior which make it more or less likely to reoccur. What might go on in the mind during learning, then, is immaterial to understanding or describing it.

Skinner's approach to understanding learning and behavior is commonly described using the metaphor of a black box (Figure 2.1). That is, the learner is a black box and nothing is known about what goes on inside. However, knowing what's inside the black box is not essential for determining how behavior is governed by its environmental antecedents and consequences.

Consider Private Johnson again, for example. It may well be that she thought of Italian food while mopping floors, but explaining her behavior does not require making reference to those thoughts. Skinner went so far as to argue that theories of learning simply get in the way of collecting empirical data on behavior change (Skinner, 1950). He denied, in fact, that radical behaviorism should even be thought of as a theory; rather, it is an experimental analysis of behavior (Skinner, 1974).




FIGURE 2.1 The Black Box Metaphor of Behaviorism

The Experimental Analysis of Behavior

By systematically observing behavior and manipulating environmental variables surrounding it, Skinner set about to discover the laws that govern learning. He defined learning as a more or less permanent change in behavior that can be detected by observing an organism over a period of time. Suppose, for instance, that the office manager in Organization X is seen logging on to the office intranet once a day. Over time, her incidence of retrieving e-mail messages increases to once every half-hour or so. From observations of her behavior, it can be said that the office manager has learned to access e-mail on a regular and frequent basis.

Respondent and Operant Behavior

Skinner distinguished two classes of behavior, respondent and operant, and it is the latter that drew most of his attention. **Respondent behavior**, studied by Pavlov in his famous classical conditioning experiments, refers to behavior that is elicited involuntarily in reaction to a stimulus. Pavlov's dogs salivating to food is one example, as is a child's startled reaction to a loud noise. By contrast, **operant behavior** is simply emitted by an organism. Skinner contended that all organisms are inherently active, emitting responses that operate on their environment. Most behavior is of this type. Birds pecking at insects in the grass, circus animals performing tricks in the ring, and students raising their hands in class are all examples of operant behavior.

Contingencies of Reinforcement

To understand why some operants are expressed while others are not, Skinner argued that we must look at the behavior in relation to the environmental events surrounding it. That is, we should look at the antecedents and consequences of behavior. Although antecedents set the context for respondent

ing, the consequences of a response are critical in determining whether it ever occurs again. If a dog puts its nose in a bee's nest and gets stung, for example, you can be sure the dog will be wary of repeating the behavior. What Skinner proposed, then, was a basic S-R-S relationship, as shown below:

S ——— R ——— S
 (discriminative stimulus) (operant response) (contingent stimulus)

This relationship provides the framework from which all operant learning laws are derived. Because the nature of the contingent stimulus determines what happens to the response, whether it is reinforced or lost, Skinner referred to learning principles as the contingencies of reinforcement (Skinner, 1969).

The concept of reinforcement, central to Skinner's behaviorism, was initially expressed by E. L. Thorndike as the Law of Effect:

When a modifiable connection between a single situation and a response is made and is accompanied by a satisfying state of affairs, that connection's strength is increased. When made and accompanied by an annoying state of affairs, its strength is decreased. (1913, p. 4)

Put simply, behavior is more likely to reoccur if it has been rewarded, or reinforced. Similarly, a response is less likely to occur again if its consequence has been aversive. In order to understand learning, then, one must look for the change in behavior that occurred and determine what consequences of the behavior were responsible for the change. In the case of the dog, for example, the consequence of putting its nose in a bee's nest was aversive, and so it learned not to do that anymore. As for the office manager, she learned to retrieve e-mail messages frequently during the day. What could be the consequence responsible for strengthening that behavior? Suppose the manager received at least one message every time she logged on and the content of the messages was information important to her job. It is likely that both the receipt of the messages and their content comprised the stimulus that was reinforcing the office manager's behavior.

It is useful at this point to re-emphasize the functional nature of Skinner's contingencies of reinforcement. That is, reinforcement as a consequence of behavior functions to enhance the probability of that behavior reoccurring. But if this probability has not been enhanced, then reinforcement cannot be said to occur. In the same vein, anything that does enhance this probability functions as a reinforcer. To illustrate, consider the following two examples:

1. E-mail is sent to the office manager throughout the day, but she never logs on to retrieve any of the messages from the intranet.

- The office manager checks e-mail with increasing frequency during the day, but she receives either no messages or ones that were directed to her by mistake.

In the first example, even though praise was contingent on the act of logging on and checking e-mail, the office manager does not increase her logging-on behavior. In this case, although receiving messages is presumed to be reinforcing, it does not function as a reinforcer. In example 2, on the other hand, the office manager's logging-on behavior does increase, but because of what consequence? In this example, it is likely that the reinforcing consequence (receiving pertinent messages) occurs at irregular times, so that the behavior of logging on is reinforced only some of the time. (The usefulness of intermittent reinforcement is discussed later in this chapter.)

Sometimes, what serves as a reinforcer is counterintuitive, as when a child keeps misbehaving despite the parent's disapproving actions. This happens because we tend to think of reinforcement as reward, and reward has generally positive connotations.

The point is, reinforcement is defined in terms of its function, its effect on behavior. Thus, we must be wary of everyday language usage of Skinner's principles, which may not precisely match his scientific meanings.

Principles of Behavior Management

Through systematic experimental manipulation of the contingencies of reinforcement, Skinner formulated learning principles to account for the strengthening or weakening of existing behaviors as well as the learning of altogether new ones. In addition, he studied reinforcement schedules to determine how learned behaviors are maintained over time. Although Skinner conducted most of his own research with animals, his principles of reinforcement have held equally well where human behavior is concerned. Since these principles are as often applied to the management of learning and behavior as to their understanding, it is perhaps easiest to discuss them in detail from that perspective.

Strengthening or Weakening Operant Behaviors

The basic principles of reinforcement describe the simple strengthening or weakening of a response already in the repertoire of the learner. That is, observation reveals whether the learner is not displaying some desired behavior often enough or is exhibiting some undesired behavior all too often. In the first instance, the desired behavior becomes a target for strengthening; in the second, the goal is to weaken the undesired behavior. As has already been discussed, the nature of the stimulus contingent on the response is an important factor in the behavior's occurrence.

But Skinner discovered a second factor that was also important. The contingent stimulus could be presented immediately after a response to influence the reoccurrence of that response, as in the receipt of e-mail causing the office manager to log on more frequently during the day. Or the contingent stimulus can be removed following a response, with a subsequent effect on the reoccurrence of the response. This would be the case, for example, if the office manager learned to delete messages regularly to avoid overloading her mailbox and causing her system to crash.

Crossing the presentation or removal of the contingent stimulus with the nature of that stimulus—whether satisfying or aversive—yields a set of basic principles for strengthening or weakening behavior, as shown in Figure 2.2. Let us consider, first, those principles that strengthen a response, followed by those that weaken it.

Strengthening a Response: Positive Reinforcement. Positive reinforcement refers to the presentation of a reinforcer (satisfying stimulus) contingent upon a response that results in the strengthening of that response. Several examples of positive reinforcement have already been discussed. Receiving e-mail reinforced the office manager's use of the intranet; commendations and an off-duty pass reinforced Private Johnson's completion of her daily floor-mopping task. Other examples of positive reinforcement can be readily observed in classrooms, at home, in social situations, or on the job. Dog trainers, for instance, reinforce "at attention" behavior with dog treats. Employers reinforce beyond

	Satisfying S	Aversive S
S presented contingent upon R	<p>Positive Reinforcement</p> <p>Example: Worker earns bonus for ideas that improve company performance</p> <p>(R strengthened)</p>	<p>Punishment</p> <p>Example: Sailor earns night in the brig for fighting on duty</p> <p>(R weakened)</p>
S removed contingent upon R	<p>Reinforcement Removal</p> <p>Example: Driver must pay stiff fine for parking in a restricted area</p> <p>(R weakened)</p>	<p>Negative Reinforcement</p> <p>Example: Student exempts weekly quizzes by performing well on daily homework</p> <p>(R strengthened)</p>

FIGURE 2.2 Basic Principles of Reinforcement

quota production on an assembly line with bonus pay. I reinforce my hands with chocolate bars for cleaning the bathtubs each week. One question that all these examples raise, however, is what precisely may serve as reinforcers? And how is one to determine which reinforcer to choose for a given situation?

Types of Reinforcers. A **primary reinforcer** is one whose reinforcement value is biologically determined (Figure 2.3). Food, for example, is a biological requirement of all living organisms, and hungry animals will exhibit all sorts of behavior to obtain it. In the well-known Skinner box (Skinner, 1938), food-deprived rats learned to press levers in order to activate a food magazine which dispensed small food pellets. Although primary reinforcement does not function extensively in human learning, it has proven quite useful in some cases. Wolf, Risley, and Mees (1964) reported using bits of food to reinforce wearing his glasses by an autistic boy.

More important in accounting for human learning is the concept of conditioned reinforcers. **Conditioned reinforcers** are those that acquire their reinforcement value through association with a primary reinforcer. Thus, they have been conditioned to be reinforcing. Examples of conditioned reinforcers include gold stars, money, and praise. Praise is a special case of conditioned reinforcement, in that it is not a tangible item that can be saved up or used in trade, like money or baseball cards. For that reason, it has been termed a social reinforcer and shown to have powerful effects on human behavior. Ludwig and Maehr (1967), for example, demonstrated that making simple statements of approval regarding students' performance in a physical education class led to their making many more positive statements about themselves. Likewise, psychology students discovered that the incidence of seat-belt use dramatically increased when grocery store checkers said to customers, "Be sure to buckle up. Remember, [store name] cares about your safety, too" (J. Bailey, personal communication).

Primary Reinforcer	Conditioned Reinforcer
... a stimulus whose reinforcement value is biologically determined	... a stimulus that acquires reinforcement value through association with a primary reinforcer
Examples: Food Sleep	Examples: Money Gold Stars

FIGURE 2.3 Types of Reinforcers

The Relativity of Reinforcers. In reviewing the conditions under which positive reinforcement influences behavior, David Premack (1959) demonstrated that behaviors in which learners already engage to a high degree may be used to reinforce low-frequency behaviors. *This procedure of making high-frequency behaviors contingent upon low-frequency behaviors in order to strengthen the low-frequency behavior* has come to be known as the **Premack principle**. It is simply a type of positive reinforcement, and one effectively exploited by parents everywhere. "You can watch TV (high-frequency behavior) as soon as you finish your homework (low-frequency behavior)."

Choosing a Reinforcer. The Premack principle illustrates well the need to observe learners in order to determine what reinforcer is likely to be most effective. In the case of the Premack principle, there is an empirical basis for selecting the reinforcer: The behavior serving as reinforcement is one the learner has been observed doing frequently. In other cases, it is often a matter of an educated guess on the basis of what is observed. Young children seem to like colored stickers and gold stars. Soldiers go off base when given the opportunity. Many adults appear to work hard, or take on additional tasks, in order to earn more money. These all have the potential, then, of serving as effective reinforcers. But only by selecting one—whatever seems most appropriate, given the learner and the behavior to be reinforced—and applying it, can one be absolutely sure of its effect. If it works, use it; if it does not, try another.

Cueing a Learned Behavior. Sometimes, a learned behavior is not exhibited, and therefore not available for reinforcement, until it is cued in some way. The case of the office manager offers a good example. Although she reads her e-mail, she doesn't send any, despite having learned how to do so during training. To evoke the appropriate response, the department chair sends the office manager, from another location, a message that requires an immediate reply. This is the discriminative stimulus. Unable to provide that reply in any way other than by e-mail, the office manager sends a return message supplying the requested information. Her response is promptly reinforced by the department chair's follow-up message, which says, "Thanks for the information. It was very helpful!"

Strengthening a Response: Negative Reinforcement. Refer to Figure 2.2. Note that in two cells, which are diagonal to one another, the behavioral principle results in the response being strengthened. Both principles are known as reinforcement, and reinforcement always results in behavior increases. In contrast to positive reinforcement, though, **negative reinforcement strengthens a response through the removal of an aversive stimulus contingent upon that response**. Remember that positive reinforcement was the presentation of a satisfying stimulus following a response.

The principle of negative reinforcement was initially discovered in experiments with rats in a Skinner box. The rats learned to press a lever, not the food this time, but to turn off a shock that was being delivered through bars on the floor of the cage. Thus, bar-pressing, a behavior that increased in frequency, was negatively reinforced by removal of the aversive stimulus shock.

Examples of negative reinforcement are harder to find than examples of positive reinforcement. As a result, its applicability is not as easily evident. Consider, however, one of the principles behind seat belts. In most cars a bell chimes or a buzzer sounds until the driver fastens the seat belt. Fastening the belt turns off the sound (which, in my car, is quite irritating). An increase in seat belt fastening, then, can be said to be negatively reinforced by the removal of the sound.

Other examples of negative reinforcement include the student who sits closer and closer to the front of the room in order to see the blackboard, and the child who finally starts brushing his teeth regularly so that his mother will stop nagging. In the first instance, sitting in front leads to the cessation of fuzzy vision. In the second, teeth-brushing brings an end to nagging.

Negative reinforcement is commonly confused with the behavioral principle of punishment, which is described next. The confusion appears to result from the connotations of the term *negative*. If something is negative, then it must be bad. If it's bad, then it must result in a decrease in behavior rather than the increase that comes with true negative reinforcement. A typical example of this confusion occurred in an article about saving sea turtles that appeared in the *Tallahassee Democrat* on November 29, 2003. The article reports that conservationists were sprinkling habanero pepper powder around sea turtle nests to deter predators. "State sea-turtle protection officials said they are aware of the pepper strategy and that it didn't appear to interfere with turtle nests. They described it as 'negative reinforcement for predators'" (*Tallahassee Democrat*, p. 58). Rather than increasing a behavior, however, this strategy is aimed at reducing it, through the application of an aversive stimulus. This is a classic example of punishment for predators of negative reinforcement for predators.

Weakening a Response: Punishment. As illustrated in the sea turtle example, punishment is the presentation of an aversive stimulus contingent upon a response that reduces the rate of that response. No doubt other examples of punishment immediately spring to mind. A father spanks a child for taking something that did not belong to him. The drill sergeant hollers, "Yacht, more push-ups! Let's go!" to the hapless recruit grousing in the back row of the formation. A teacher yells at the student who was talking with a neighbor instead of studying. In all instances, the individual administering punishment for some misbehavior does so with the expectation that the behavior will stop and not be repeated.

Although punishment has the effect of stopping behavior, and in fact is so-called because it has that effect, it also appears to have unfortunate side effects. First, its effectiveness tends to be short-lived. That is, the behavior being punished may come to an immediate halt at the time punishment is administered, but this does not mean it has been necessarily forgotten. The student may quit talking in class when yelled at, only to do it again at another time, perhaps more surreptitiously. A dog I once had provides another good example of this. Shadow was not permitted to jump on the furniture, and she was smacked with a rolled-up newspaper if she tried. My husband and I thought we had stopped this behavior altogether (and proud we were of our success in using behaviorist principles!). But one day when I was home alone, I walked into the living room, and although there was no dog in sight, the rocking chair was rocking furiously!

Azrin and Holz (1966) discussed other, more serious, problems with the use of punishment to reduce undesirable behavior. When punishment involves a particularly aversive stimulus or induces pain, it can lead to undesirable emotional responses being conditioned. If fear is elicited, then avoidance or escape behavior may be negatively reinforced inadvertently (Skinner, 1938). Running away and truancy are good examples. A child does poorly in school, is punished severely, and then manages to escape or avoid the punishment by leaving home or cutting class.

The emotional side effects of punishment that is painful are not limited to fear, however. Aggression and anger may result, particularly in individuals who are characteristically aggressive (Azrin, 1967). Moreover, punishment can serve as a model for aggression. In a series of studies examining aggressive behavior in children, Bandura, Ross, and Ross (1961, 1963) demonstrated that those who observed others being aggressive were more likely to be aggressive themselves. This is further supported by evidence from studies of abusive families; by and large, parents who are abusive were themselves abused as children (Steinmetz, 1977; Strauss, Gelles, & Steinmetz, 1980).

Finally, a long history of punishment may cause physical or psychological harm. Especially in situations where the aversive stimuli cannot be avoided or escaped from, the phenomenon of **learned helplessness** may result. This refers to the passive acceptance of events seemingly beyond one's control, a phenomenon first demonstrated in a now classic experiment conducted by Seligman and Maier (1967). In their study, conducted in two phases, unpredictable and painful shocks were administered to dogs. For some of the dogs, escape from the shock was possible through a panel in the cage. For the others, escape was not permitted, no matter what they did. In the second phase of the study, the dogs were placed in one of two compartments of a box. A tone sounded to warn of impending shock in that compartment, which the dog could escape by jumping the barrier into the second compartment. The dogs who had been allowed previously to escape the shock

learned quickly to jump the barrier each time they heard the tone. The dog who had previously been prevented from escaping the shock. The dog made little attempt to escape under these new conditions.

When individuals perceive that their actions have little effect on given events, they, too, begin to exhibit symptoms of learned helplessness. In the context of learning, experiencing repeated failure or constant belittlement of their efforts can lead students to say, "I can't do this. I'm not a good reader" (or writer, or test-taker, or what have you).

With so many problems associated with punishment, under what conditions can it be useful? Azrin and Holz (1966) suggested that punishment has an advantage over other procedures when there is a need to stop a behavior quickly. For example, if a child is about to injure herself by picking up a hot iron, a fast slap on the wrist or loud "NO!" may be the most effective way to gain her attention and stop her in the act. Similarly, Corte, Wolf, and Locke (1971) found punishment to be the most effective procedure for eliminating self-injurious behavior in retarded children.

Finally, when used sparingly, punishment has the advantage of conveying information about what behaviors are considered appropriate or inappropriate in given situations (Azrin & Holz, 1966; Walters & Grusec, 1977). Sometimes, individuals simply are not aware that their behavior is unacceptable; it may be that the rules are different from what they have been accustomed to. This may happen particularly in multicultural situations, when, for example, ways of interacting that are socially acceptable at home or in one's neighborhood are not acceptable at school. It is for these situations that some behaviorists also recommend a warning precede punishment and reasons accompany it to explain why certain behaviors are not tolerated (Walters & Grusec, 1977).

Weakening a Response: Reinforcement Removal. Whereas one way to reduce the frequency of behavior is to present an aversive consequence, another, perhaps more effective means is to take away reinforcement when the behavior occurs (see Figure 2.2). Removing reinforcement can be done with the principles of response cost and timeout. However, a special case of reinforcement removal, which involves the absence of reinforcement, is discussed first.

Extinction occurs when previously existing contingencies of reinforcement are taken away, thereby causing a reduction in the frequency of a response. In other words, reinforcement that has been maintaining some behavior is simply stopped. For example, a teacher stops paying attention to a student making waving his arm in the air, and he eventually gives up. Or, a pet owner ignores a dog's whining and the dog eventually stops.

When extinction is used as a procedure for weakening some undesirable behavior, the key to its success is persistence. As most pet owners have undoubtedly experienced, the dog that is being ignored will redouble its ef-

forts for attention at first. Woe to the owner who gives in at this point, however! Delaying attention simply serves as an intermittent schedule of reinforcement, which we see later in the chapter has the effect of greatly strengthening behavior. With extinction, it is important to consistently withhold reinforcement; eventually the behavior will lessen. As with punishment, it is also useful to reinforce some alternative, desirable response concurrent with extinguishing the undesirable behavior. In that way, learners are being rewarded for something even while they have lost reinforcement for something else.

Response cost, like extinction, involves the removal of reinforcement contingent upon behavior. But in the case of response cost, this is done by exacting a fine, requiring the offender to give back some previously earned reinforcer. It can have a strong and rapid effect on reducing certain behaviors for some people, depending on the history of the person and the value of the fine (Weiner, 1969). In society, for example, the fine for minor infractions of the law is usually monetary. To be effective, fine amounts should be set high enough to reduce the likelihood of repeat behavior, but it is certainly true that, no matter what the fee, it may have less effect on a rich person or one who has been successful at avoiding payment.

Response cost applied in a school setting can be seen in the following example. On a class field trip, Ms. Johnson was in charge of the six third grade boys most likely to cause trouble. The morning of the trip, she told them what rules of conduct they were to follow, that they would earn stickers for good behavior, but that they would have to give back a sticker every time they broke a rule. After warning one boy twice for the same behavior, Ms. Johnson said, upon the third occurrence, "You know what the rules are, right?" The little boy said yes and tearfully handed her the only sticker he had earned so far. The happy outcome to this story is that the boy behaved without incident the rest of the day and earned the big treat Ms. Johnson had been saving for last.

The final principle involved in reducing behavior, **timeout**, does so by removing the learner, for a limited time, from the circumstances reinforcing the undesired behavior. In some situations, it is very difficult to determine precisely what consequence is responsible for maintaining some behavior. It may be the case, moreover, that several events follow a behavior and all have some reinforcing effect. In a typical classroom, for example, a student's acting out, accompanied by "Watch me!," may cause the teacher to stop class and the other students to laugh, both of which may contribute to its recurrence. Stopping the behavior, then, may take more than simply ignoring it (extinction). Yet other conditions may not make response cost an appropriate alternative.

In cases such as these, individuals may be removed altogether from the sources of reinforcement. Wolf, Risley, and Mees (1964) used timeout to virtually eliminate temper tantrums thrown by an autistic boy. Every time a

tantrum occurred, they isolated him in a room by himself for slightly longer than the tantrum had been. Solnick, Rincover, and Peterson (1977) added further evidence to the effectiveness of timeout, but noted that it can be reinforcing instead of punishing in some circumstances. Imagine, for example, a noisy classroom. The disruptive behavior of one student causes the teacher to put him out in the quiet hallway with his assignment. The next time the class is noisy, this student acts out again, with the same result. What appears to have happened is not timeout at all; rather, the disruptive behavior has been negatively reinforced by the student's escaping the noisy classroom for the quiet hallway.

Sulzer and Mayer (1972) suggested that for timeout to be most effective, the following conditions should be met. Timeout should not be used from an aversive situation (illustrated in the example above). It should provide for removal of all reinforcement, it should be consistently maintained, and the time period should be kept short (a general rule of thumb is one minute for each year of the learner's age). Finally, like extinction and punishment, time-out should be used with other procedures that reinforce alternative, desirable behaviors.

Depicted in Figure 2.4 is a concept tree for principles of behavior management. It illustrates in a visual way what attributes are shared by certain principles (e.g., those that strengthen behavior) and what attributes are unique to each one (e.g., high-frequency behavior as reinforcer is unique to the Premack principle). The tree also includes, for each principle, an example illustrating its use or occurrence.

Teaching New Behaviors

The principles discussed in the previous section concerned behaviors that were already present to some degree in the learner's repertoire. One might say that the learner already knew the behavior; what was learned seemed to be the frequency with which the behavior was performed. But how are behaviors learned that are not already present in the organism's repertoire? Bar-pushing, for example, is not a behavior that rats do in their natural environment. Similarly, one could watch a pigeon in a Skinner box for a long time without ever seeing it turn around in a complete circle. In Mr. Tanner's class, students are unlikely to spontaneously exhibit a complex behavior such as the foxtrot. If the students, the rat, and the pigeon never exhibit the behavior targeted for reinforcement, how does it come to be acquired? Behaviorists have defined three principles for teaching new, and in many cases, complex behaviors: shaping, chaining, and fading.

Shaping. Shaping refers to the reinforcement of successive approximations to a goal behavior. It involves positive reinforcement in that a reinforcer is presented contingent upon desired behavior. But in the case of shaping, the desired

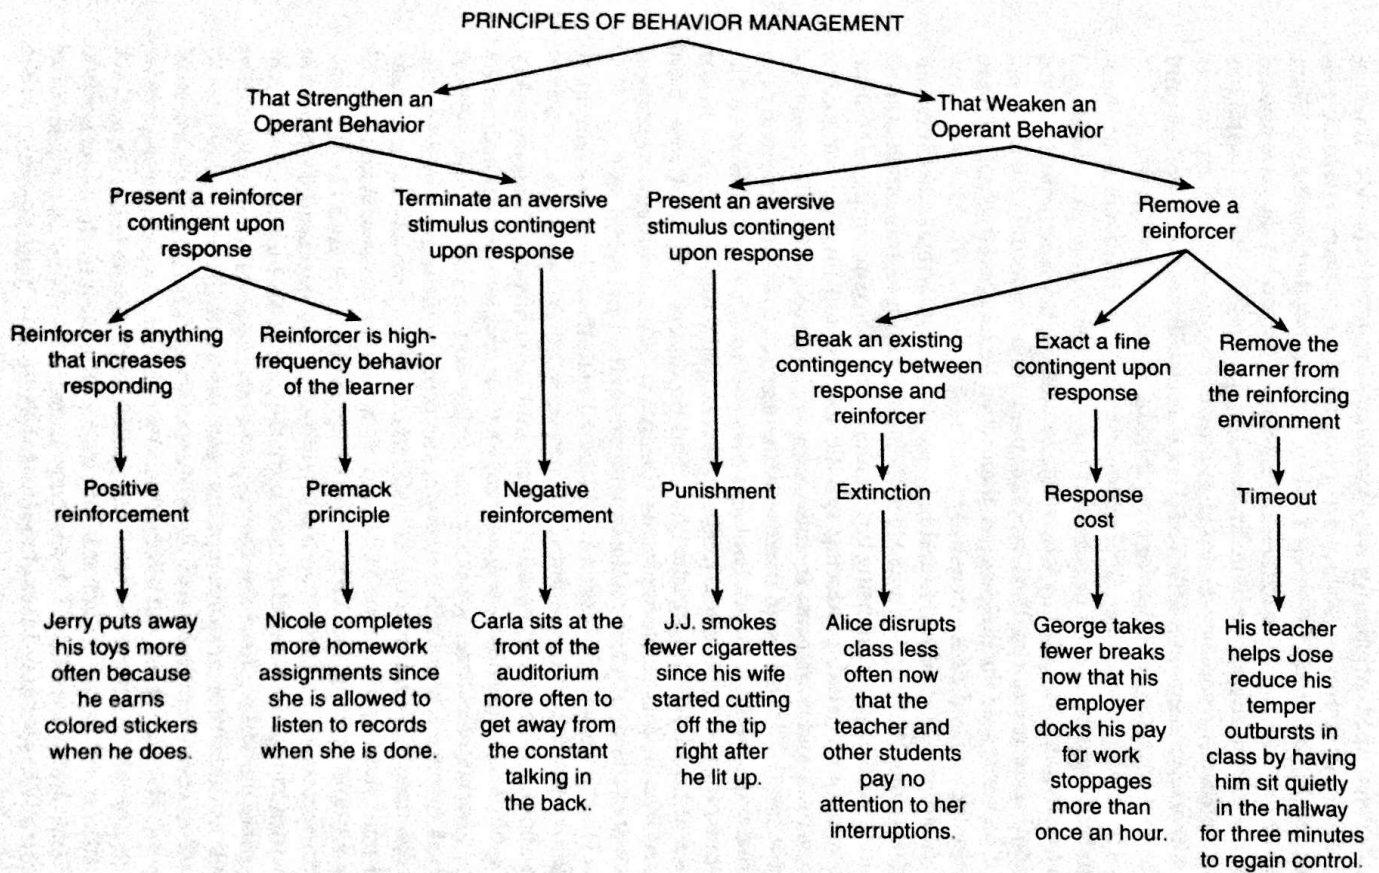


FIGURE 2.4 A Concept Tree for Principles of Behavior Management

behavior reinforced each time only approximates the target behavior. And successively closer approximations are required for the reinforcement to be presented (Reynolds, 1968). To teach a rat to press a bar, then, one might first reinforce proximity to the bar, then raising a paw, extending the paw toward the bar, touching the bar, and finally, pressing the bar. As soon as the rat has made the correct response—in this case, pressing the bar—then the principle of positive reinforcement is followed. That is, each bar press is reinforced until the desired frequency of behavior is exhibited.

Harris, Wolf, and Baer (1967) demonstrated the effectiveness of shaping to teach new behaviors to children. They selected climbing on the jungle gym for shaping in a little boy who spent no time on it. Teacher attention was the contingent reinforcer. Thus, teachers paid attention to the little boy first when he went near the jungle gym, then when he touched it, climbed on it, and finally, climbed on it extensively.

Shaping has also been found to be particularly effective in teaching autistic children. Wolf, Risley, and Mees (1964), for example, trained an autistic boy in speech acquisition, using bits of food to reinforce making eye contact, producing any sound, producing specific sounds, and finally saying complete words and sentences. In this example, however, as in the previous ones, it could still be argued that the learners were capable of producing the desired response; they just did not. Bar-pressing, in other words, is not a difficult response. Climbing on a jungle gym was well within the capabilities of the small boy. Even the autistic boy could produce sounds that were then shaped into language. Is shaping as effective with truly difficult responses, which are not initially within the capabilities of the learners?

That the answer is yes can be illustrated with the following example. A waiter at a Moroccan restaurant served tea with dessert by raising the teapot high above his head and pouring the tea into tall, narrow glasses on a very low table, where we diners were sitting on floor cushions. He spilled nary a drop, and so, of course, we marvelled at his skill and asked how he had learned to pour tea in such a manner. His reply went something like this: "Well, naturally, I couldn't do it at first without spilling tea all over the place. So, I tried holding the teapot only slightly above the glasses. When I could pour without spilling, I moved the teapot up a few inches. And I kept repeating this process until I could do it with the teapot over my head." Successive approximations had been reinforced until the goal behavior was achieved. In this case, the ability to make the response at one level of approximation served as the reinforcer to attempt the response at one level of approximation.

The above example also illustrates a factor critical to the success of shaping. The waiter did not attempt a more difficult approximation until he had mastered the easier one. Similarly, in shaping any new behavior, a closer approximation to the goal should not be reinforced until the previous one has been firmly established. If too large a step is expected of the learner at once, the behavior may break down and shaping may have to resume at the

point where the learner has repeatedly demonstrated success. Finally, it is also important in shaping to ensure that reinforcement is delivered immediately contingent upon the desired response. Any delays can result in some random behavior being reinforced and becoming conditioned.

Skinner (1948) called this superstitious behavior and demonstrated its inducement by delivering noncontingent reinforcement to pigeons. That is, he offered food at random intervals, not dependent upon the animal's behavior. Whatever the pigeon happened to be doing at the moment reinforcement arrived, however, became more likely to reoccur because of the reinforcement. As a result, Skinner observed the inadvertent conditioning of all sorts of weird behavior, and he argued that the simple contiguity between response and stimulus could account for the learning of superstitious behavior in humans. For example, you buy a new pen with which to take a particular test, and you score well on the test. Scoring well rewards your use of that pen, and so you begin to attribute good performances to the causally irrelevant pen when, in fact, good performance was contingent upon your study behavior.

Chaining. Whereas shaping is used to teach new behaviors that are relatively simple and continuous in nature, **chaining serves to establish complex behaviors made up of discrete, simpler behaviors already known to the learner.** A typical example of chaining in human behavior is learning a new dance. Each dance step may be acquired through shaping, but then the steps are strung together in sequence through forward or backward chaining. In other words, one might begin by practicing the last step in the dance and then progressively add the steps that precede it (backward chaining). Or one could start with the first step and progressively add steps that follow until the entire dance can be performed (forward chaining).

Memorizing long passages of prose is another typical example of forward chaining. Sentences are added in succession until the entire passage can be repeated without error. Finally, reassembling their weapons after cleaning is a behavioral chain that is probably acquired through forward chaining by the soldiers in the scenario, *Boot Camp*.

Discrimination Learning and Fading. To this point very little has been said about the control the setting has over learning except in terms of the consequences of behavior. Behaviors are acquired and exhibited because they are reinforced; nonreinforced behaviors tend not to occur, at least in the setting where they have been ignored or punished. This is an important distinction. Individuals are clearly able to distinguish between settings in which certain behaviors will or will not be reinforced. A playful slap on the back may produce grins from the guys in the gym, but it is likely to have a quite different effect on one's commanding officer or teacher. Thus, something besides the behavior itself must be learned, and these are the cues, or discriminative

stimuli (S^Ds), which signal to the learner when and where the behavior should be performed.

Most learning in formal instructional situations is accompanied by cues. School bells signal the end of classes; getting up to leave before the ring is a behavior likely to be punished. Thus, staying in one's seat is reinforced before the bell rings; moving about the halls is reinforced after the rings. The bell simply acts as a cue to indicate what behavior is appropriate and will be reinforced (or conversely, what behavior is inappropriate and will be punished).

Discriminations are often learned, then, by a behavior being reinforced in the presence of one stimulus and being punished in the presence of another. Alternatively, a different behavior may be reinforced in the presence of the second stimulus. Motor vehicle drivers, for example, must learn to stop at a red light and go on the green light. Thus, the S^D for stopping is a red light, and the S^D for going is a green light. In either case, however, errors can sometimes be extremely costly, so that applying the simple principles of positive reinforcement and punishment may not be the most effective for establishing the discrimination.

In his studies with pigeons, Terrace (1963a, 1963b) demonstrated that almost errorless discrimination performance could be achieved with fading. He first taught the pigeons to peck a red key, so that red became a discriminative stimulus for pecking. Then he turned off the key, which caused the pigeons to stop pecking, and gradually lengthened the intervals during which the key was dark. The darkened key then became the discriminative stimulus for not pecking. Finally, Terrace slowly faded in a green light in place of the darkened key. Since the pigeons never pecked the dark key, and the fading was so gradual from darkened key to well-lit green key, the green key came to be established as the S^D for not pecking.

The concept of fading as it has been applied to human performance has come to refer to the fading out of discriminative stimuli used to initially establish a desired behavior (Sulzer & Mayer, 1972). In other words, the desired behavior continues to be reinforced as the discriminative cues are gradually withdrawn. A classic example of fading used in instruction can be seen in Skinner and Krakower's (1968) *Handwriting with Write and See* program. In this program, children trace letters in an instructional workbook. Gradually, portions of the letters, which serve as the discriminative stimuli for forming the right shapes, are faded, thus requiring the children to compare increasingly more of each letter. Reinforcement is accomplished through a special chemical reaction between the pens used by the children and the paper. They form a black line when their letters are correct but the paper turns orange when the pen moves from the prescribed pattern.

Other examples of fading can be seen in the gradual reduction of verbal cues given by a laboratory instructor as students work through a set of procedures for staining slides or in the withdrawal of physical cues given by a

golf pro showing a beginner how to hold and swing a golf club. Job aids in industrial settings are also good examples of fading. As employees become more proficient in their assigned duties, they rely less and less on the cues provided by the aid.

Maintaining Behavior

If we consider that the job of instruction is not only to bring about desired changes in behavior, but to maintain them as well, then we must determine what conditions will be most effective for behavior maintenance. A typical behaviorist approach to the question would be to find some high-frequency, persistent behavior occurring naturally and to study the consequences responsible for its maintenance. One good example is people playing the slot machines at Las Vegas or Reno. Some will stand there for hours, doing nothing but pumping coins or tokens into the machines and pulling the handle. Every so often, the player receives a payoff, accompanied by flashing lights and ringing bells. So what is going on here?

Skinner was apparently in search of a means to economize on the costs of feeding his experimental subjects when he made an interesting discovery (Leahey & Harris, 1997). When he reinforced only some of the bar-pressing responses made by his rats, rather than reinforcing every response, the behavior became much more resistant to extinction. In other words, continuous reinforcement, while necessary to establish a response in the first place, was not essential to maintaining that response. In fact, intermittent reinforcement worked much better for that purpose. By systematically investigating schedules of reinforcement, Ferster and Skinner (1957) were able to determine what pattern of reinforcement gave rise to what sort of behavior maintenance.

Although behaviorists have investigated reinforcement schedules and invented new ones since Ferster and Skinner's original experiments, four basic schedules remain. These are determined on the basis of whether reinforcement is contingent upon a given response (called a ratio schedule) or upon the passage of time (called an interval schedule). In addition, reinforcement can occur regularly, after a fixed amount of time or number of responses, or it can occur irregularly, after a variable amount of time or number of responses. Taking these characteristics together, we have four possible schedules, as shown in Figure 2.5: fixed ratio, fixed interval, variable ratio, and variable interval.

Fixed Ratio Schedules. Continuous reinforcement, i.e., reinforcing every desired response, amounts to the same thing as a fixed ratio schedule of one (FR1). Ratio schedules of reinforcement are those in which the reinforcer is delivered contingent upon the response made by the learner. A fixed ratio schedule, therefore, requires the learner to make so many responses before

Reinforcement is contingent upon:		
	Responses	Time
Reinforcement occurs consistently	Fixed Ratio Example: Students earn points for every skill mastered	Fixed Interval Example: Worker is paid every two weeks
Reinforcement occurs intermittently	Variable Ratio Example: Slot machines pay off after a random number of pulls on the lever	Variable Interval Example: Drill sergeant makes spot-check inspections

FIGURE 2.5 Types of Reinforcement Schedules

reinforcement is delivered. Quota systems on factory assembly lines are examples of fixed ratio schedules. For every fifteen widgets produced (FR15) or for every 300 chickens inspected (FR300), employees earn a standard wage credited toward their pay. This type of reinforcement schedule tends to produce a response pattern like the one shown in Figure 2.6A. In other words, responding occurs at a high and steady rate, since the more employees produce, the quicker they earn more money. Animals responding on a fixed ratio schedule also show a tendency to pause immediately following reinforcement. While this phenomenon has not been demonstrated consistently with humans, studies have shown that it can occur. For example, I typically put myself on an FR15 schedule when grading undergraduate assignments, getting up for a snack or a short walk after each fifteen papers graded. Getting started again after the break, however, generally entails a pause before I am fully focused on the task once again.

Fixed Interval Schedules. As indicated, time is the determining factor for an interval schedule of reinforcement. For a fixed interval schedule, then, reinforcement is delivered after some fixed period of time, such as 5 minutes (FI 5 min) or 10 days (FI 10 days). A commonly cited example of this type of schedule is the procedure by which many professors are tenured and promoted. Although tenure and promotion are ostensibly tied to performance, they are typically awarded, or become available for award, at particular times, such as after so many years in rank. As a result, performance over time may take on the characteristic “scallop” typically produced by a fixed interval schedule (Figure 2.6B). In other words, responding becomes more frequent as the time for reinforcement nears. Weekly quizzes can produce a




FIGURE 2.6 Response Patterns Produced by Different Types of Reinforcement Schedules

similar pattern, with students spending more time studying as the time for the quiz draws near.

Variable Ratio and Variable Interval Schedules. In variable schedules, the time or number of responses required for reinforcement is varied from reinforcement to reinforcement. Thus, a VR5 schedule means that, on the average, reinforcement is delivered for every five responses, but one time it may be given after the second response and the next time after the eighth response. Similarly, a variable interval schedule of 5 minutes (VI 5 min) means that reinforcement may be given after 3 minutes, then after 7 minutes, then after 4 minutes, and so on, creating an average interval of 5 minutes.

Variable schedules typically produce the highest and steadiest rates of responding, with variable ratio schedules producing the highest of all (Figure 2.6, C and D). The slot machine example provided earlier demonstrates the effect of a variable ratio schedule; typically, payoffs are scheduled to occur after some average number of pulls on the lever. (This average, by

the way, is set high enough so that the money taken in is always more than the money paid out.) In a classroom setting, teachers can assure steady rates of studying or homework completion by administering pop quizzes on the average of once a week (VI 5 days), or by collecting and spot-checking assignments (e.g., every third assignment, on the average, for a VR3 schedule).

Planning a Program of Behavior Change

To this point, principles have been discussed that relate the incidence of behavior to its environmental cues and consequences. Learning has been described as a relatively permanent change in behavior, and schedules of reinforcement have been presented that are useful for maintaining such changes. The question that remains is, How can these principles be systematically applied in order to bring about specific, desired changes in behavior? What follows are five essential steps in implementing a behavior change program (see Table 2.1 for a summary). Evaluating the success of such a program will be discussed last.

Step One: Set Behavioral Goals. In order to go about changing behavior, one must determine what behavior is to be changed and what the change is. Questions to consider in this step are: What is desirable behavior? How often should the behavior occur? Does the change in behavior involve its being strengthened or reduced? Is some new behavior to be taught? What are the requirements for behavior maintenance? Also essential to setting behavioral goals is knowing to what extent the targeted behavior is being exhibited relative to its desired strength. In other words, is the learner not doing enough of something, or is some behavior being exhibited too often?

In order to have an accurate answer to the question of what learners are actually doing in any given situation, they must be observed. For example, suppose Charla is a student in Mr. Tanner's class who is "always acting out," which disrupts class and wastes valuable instructional time. Observation may reveal, however, that Charla acts out only three or four times a day. The results of her behavior are so severe in terms of lasting impact on the class that it just seems she is acting out more often.

Observation, therefore, provides a baseline of behavior, a measure of behavior incidence as it occurs before any intervention is implemented.

TABLE 2.1 Implementing a Program of Behavior Management

- Step 1: Set behavioral goals.
- Step 2: Determine appropriate reinforcers.
- Step 3: Select procedures for changing behavior.
- Step 4: Implement procedures and record results.
- Step 5: Evaluate progress and revise as necessary.

From this baseline, goals for change can be determined. The teacher may still decide, for example, that Charla's acting-out behavior should be reduced because of its adverse impact on the class. A reasonable goal may be to reduce the incidence of acting out to no more than once a day. As we will also see, the baseline provides a basis against which the success of the intervention can be measured.

Although behavior management is typically implemented on an individual basis, it can be effective in group situations. Mr. Tanner, for example, has clearly set goals for appropriate behavior for all students with the rules he posts in his classroom. He and the students then jointly plan Steps Two and Three, which follow.

Step Two: Determine Appropriate Reinforcers. The choice of reinforcers for use in a behavioral change program depends on the learner, the instructor, the behavioral goals, and the practical circumstances surrounding the implementation of the program. Behavior in young children, for example, may be reinforced with colored stickers or gold stars, which would clearly not be appropriate or effective with older students or adults. Some teachers are opposed to the use of tangible rewards, preferring instead to use praise, attention, and other social reinforcers. A behavioral goal that involves reducing a behavior may call for a procedure such as response cost, which means that appropriate fines rather than reinforcers must be determined.

Finally, there will always be pragmatic considerations in choosing reinforcers. It is not always easy to determine what will be the most effective reinforcer for a particular individual, or, once a reinforcer is identified, it may not be within the control of the program designer. Peer approval, for example, can be a particularly potent source of reinforcement for teenagers (Sulzer & Mayer, 1972), but it is not something easily controlled by teachers or parents. Moreover, an ethical dilemma may arise in some situations as to whether the program professional has a right to control an effective reinforcer. In a community mental health facility, for example, money to buy cigarettes or candy has been found to have powerful reinforcing effects on the residents (Mulligan, Ogleby, & Perkins, 1980). But should the mental health professionals have control over the residents' money in order to use it as reinforcement?

Step Three: Select Procedures for Changing Behavior. The decision as to what procedure should be used obviously depends on what behavior change is desired. To strengthen an existing behavior, positive and negative reinforcement and the Premack principle are possibilities. To teach a new behavior, one might select from shaping, chaining, or fading. To maintain behavior, some schedule of reinforcement should be selected to produce the desirable pattern of performance. And finally, to reduce or weaken a behavior, punishment, response cost, timeout, or extinction could be implemented. To choose from among the options, where more than one procedure may be

appropriate to a given goal, one should consider such questions as, How important is it that I effect this change in behavior quickly? How permanent is the result of this procedure likely to be? What other unintended effects might this procedure have that I would like to avoid? Are there any additional factors that should be taken into consideration?

Step Four: Implement Procedures and Record Results. Once a plan for behavior change has been generated, it may be implemented and its results monitored. Observation again becomes important at this step, since only by looking at the behavior can any change from baseline be detected. Recording behavior incidence also helps to ensure that real, rather than imagined changes are monitored. It is easy to engage in wishful thinking, hoping for changes or thinking that changes must have occurred by virtue of the program being in place.

Step Five: Evaluate Progress and Revise as Necessary. Based on the records kept in Step Four, it should be easy to see whether, in fact, any change from baseline behavior has occurred. If the program was designed to reduce some behavior, it should have had that effect, or if it was designed to teach some new behavior, then that behavior should now be in evidence. Assuming that the desired behavior has been achieved, no change in the program established at a desirable rate, some alteration in the reinforcement schedule may be required to sufficiently maintain the behavior.

What if, on the other hand, the program has not produced the intended results? Any number of possibilities could be the problem, but according to Skinner, simple observation and systematic alteration of the program should enable you to find out which one is the culprit. It may be that another procedure would be more effective, or that a different reinforcer should be selected. Perhaps a combination of procedures should be tried, as in reducing one behavior while at the same time reinforcing an alternative to take its place. Whatever the problem, the program should be modified appropriately and implemented again. This process of monitoring results and revising as necessary should be repeated until success has been achieved.

Since radical behaviorism is the experimental analysis of behavior and behavior is assumed to be reliably, functionally related to environmental events, a behaviorist would not necessarily be content with showing that a behavior change had occurred following the implementation of some program. It would also be necessary to reverse the implementation of some program in this way can we be sure that it was the procedure, or remove the program, random set of variables, that was responsible for the change in behavior. To take an example, suppose that Mr. Tanner decides to implement time-out in order to reduce Charla's acting-out behavior. During baseline, con-

ducted over a week's period, observations revealed that Charla acted out about four times each day, for an average duration of 5 minutes each. The typical results of the acting out included the teacher stopping class, paying attention to Charla to get her under control, and then spending some minutes trying to regain the attention of the rest of the students. At the beginning of week 2, Mr. Tanner implements the timeout procedure, isolating Charla for 8 minutes each time she acts out. He does this by taking Charla by the hand without saying a word and putting her in a chair just outside the classroom door. At this point, he says, "When you can be quiet, you can return." Mr. Tanner continues class and after 8 minutes allows Charla to return to her seat.

Suppose that timeout appeared to be effective, and Charla's acting out dropped to once a day by the end of the week. To be sure that it was timeout, and not something else going on in class having the desired effect, the teacher would institute a reversal in procedure and stop using timeout during week 3. Thus, he would go back to his original reaction to Charla's behavior, which should have the effect of increasing its incidence. Finally, at week 4, timeout would again be reinstated and its results monitored. If timeout is indeed an effective procedure for reducing the undesired behavior of acting out, then the record of results should resemble that displayed in Figure 2.7.

Although this reversal process has the advantage of demonstrating the functional relationship between any behavioral procedure and behavior, it also has several disadvantages in a practical, rather than experimental, context. First, it is time-consuming to establish a reasonable estimate of baseline and to carry out each phase long enough to demonstrate a procedure's effectiveness. More importantly, however, once a behavior change has been effected, it may be extremely counterproductive, or even unethical, to return that behavior to its original rate. For example, performing tasks without demerits is desirable behavior in boot camp whether or not earning off-base passes is the only factor responsible for its occurrence. As such, most applications of behaviorist principles are considered to be successful when the goals for behavior change have been met.

Contributions of Behaviorism to Instruction

Few would argue that radical behaviorism has had a profound impact not only in psychology but on instruction as well. And its influence continues to be felt in fields ranging from clinical therapy to instructional design. Although many applications and new developments in behaviorism go beyond the scope and purpose of this book (e.g., biofeedback, treatment of clinical depression), others bear examining. The ones I have chosen to discuss pertain to




FIGURE 27 Occurrences of Acting-Out Behavior in Relation to Timeout: A Hypothetical Case

changing personal behaviors, managing learning and behavior in instructional systems, and improving performance in organizational systems.

Changing Behavior Through Behavior Modification

Application of behavioral principles in the way described so far in this chapter is essentially the same as behavior modification (also known as behavior therapy or contingency management). Typically, behavior modification is used to treat problem behaviors in social, personal, or school situations. Clinical applications include treatments for phobias, or obsessions, or eating disorders, to name a few (Bower & Hilgard, 1981). Instructional applications involve treatment of school-related problems, such as inattention, hyperactivity, temper tantrums, or any behavior that interferes with learning and the normal conduct of classroom activities. Special education teachers are typically well trained in the use of behavior modification, since they regularly deal with children who have special problems and special needs. As part of the individual education plans for individual students, teachers may

target problem behaviors, devise and implement interventions, and keep records to monitor student progress and inform changes to the original plan.

In recent developments, behavior modification methods are taught to individuals, who then use them to change their own behavior. This is an application of behavior modification known as self-control, and it has been successfully demonstrated with people who wish to lose weight, quit smoking, or improve their social skills, study habits, or concentration. Bower (Bower & Hilgard, 1981) reported that he taught a college seminar in which students teamed up with a cooperative friend in order to change some aspect of their own behavior. One of my own favorite examples of self-control came from a friend whose husband enlisted her help to quit smoking. Given his propensity toward saving money, they decided an appropriate punishment would be sending money to fly-by-night charities. Therefore, the husband wrote a series of \$25 checks and handed them to his wife with instructions to mail one every time she saw him smoking. Three checks later, he had quit smoking altogether and, to my knowledge, has never smoked since.

Managing Learning and Behavior in Instructional Systems

Whereas behavioral therapists and special education teachers generally focus on the needs of individuals, teachers in regular classrooms may have twenty to thirty students or more to manage at one time. Likewise, instructional designers may be developing instruction with goals to be achieved by individuals or groups. For teachers and designers, behavioral principles are useful for managing learning and behavior within instructional settings such as classrooms, individualized instruction, and on-the-job training.

Classroom Management. To a limited extent, teachers may apply behavior modification to change the problem behaviors of one or another student. More often, like Mr. Tanner, they set up group contingencies, i.e., a standard reinforcement given to individuals or the group as a whole for following certain rules of conduct.

One means of applying group contingencies in the classroom that some teachers find useful is the token economy (Ayllon & Azrin, 1968). In this system, tokens serve as conditioned reinforcers that can later be exchanged for objects or privileges. Tokens are earned for good conduct—whatever behaviors have been identified by the teacher for strengthening. But since tokens operate much like money, students may be fined for breaking the rules or engaging in behavior the teacher has deemed undesirable.

In one of the first formal uses of a token system for reinforcing and maintaining desired behaviors, patients at a mental hospital earned tokens for appropriate behaviors in the ward (Ayllon & Azrin, 1968). With their tokens, patients could buy candy, soda, trips to town, movies, and the like. Bushell,

Wrobel, and Michaelis (1968) demonstrated the effective use of a token system with preschool children to strengthen school-related behaviors. The study took place in a regular classroom setting, and children earned tokens for such behaviors as attending to assigned tasks, being quiet, asking questions, and so forth.

When tokens can be exchanged for objects, such as books or toys, keeping a steady supply of such things can become rather expensive. Sulzer and Mayer (1972) therefore recommended that teachers consider using a variety of activities for which students can exchange tokens.

Instructional Objectives. According to the behavioral perspective, just as we can set goals for appropriate behavior, so can we express in behavioral terms the instructional outcomes we desire students to achieve. In fact, behaviorists would argue the only evidence we have of learning comes from the students' behavior; they can do something after instruction that they could not do before. It is important, therefore, to specify desired instructional outcomes in terms of clear, observable behavior. These goal statements are variously called behavioral objectives, instructional objectives, or performance objectives.

Mager (1962) made popular the three-component objective, which states the behavior to be acquired, the conditions under which the behavior is to be demonstrated, and the criteria governing how well the behavior is to be performed. Typical Mager-type objectives, for example, would include the following:

1. Given the values of two sides of a right triangle, students will be able to correctly solve for the value of the third side.
2. Handed the pieces of an unassembled M-16, the soldier will be able to assemble the weapon in no more than 2 minutes.

Although other objective formats are used (for example, the five-component formats used in the instructional design models of Merrill [1983] and Gagné, Briggs, & Wager [1992]), all specify essentially the same information. The effectiveness of instructional objectives for enhancing academic performance has been debated since the 1960s, primarily because research studies have yielded mixed results. Gagné (1985) argued for informing learners of objectives, since doing so readies them for learning. Objectives also provide a framework for studying what will eventually be tested. In a meta-analysis of research on objectives, Klauer (1984) provided evidence of a small, positive effect of objectives on learning, but also noted that objectives tend to focus learners' attention on certain information and away from other information. This would suggest that, to enhance learning, objectives must be written for all information considered important to learn.

As will be discussed again in Chapter 10, many educational and training programs today are based on objectives. However, they probably do not

include objectives for each and every skill or piece of information that students might expect to learn. Rather, objectives are written for critical skills or the minimum information deemed acceptable for a graduate of the program to know. In addition, some educators (e.g., Popham, 1988; Reiser & Dick, 1996) suggest that students be given simpler, perhaps more general, statements of objectives to guide their learning, since these are easier to understand and yet still keep students and teachers alike on the same track toward particular goals.

Contingency Contracts. An instructional application that may make use of both *behavior modification* and *instructional objectives* is the **contingency contract**. Used with individual students, the contract sets out the terminal behavior the student is to achieve, along with any conditions for achievement and the consequences for completion (or noncompletion) of the assigned task(s). The contract is negotiated between teacher and student, and both agree to its terms.

Contingency contracts are particularly useful in open educational systems, where students from several grade levels participate together in learning activities. Since students are not all at the same achievement levels, they negotiate individual contracts each week indicating their expected progress in accomplishing objectives in subject areas such as math or reading, for example. Instructors at all levels of schooling have also found contingency contracts to be a useful means of managing independent study projects. Instead of simply giving an assignment such as, "write a 10-page research paper on a topic related to behavioral psychology," instructors may negotiate with individual students on what should be included in the paper and how well it should be written.

Personalized System of Instruction (PSI). In 1968, Fred Keller proposed a whole new approach to college instruction that was based on behavioral principles (Keller, 1968). Keller noted problems with typical group instruction in the classroom—delays in reinforcing achievement, students progressing to more difficult instruction when they have not mastered basic material—that he believed could be solved with the personalized system of instruction (PSI), also known as the Keller Plan. PSI calls for course material to be broken up into units, or modules, each with a set of behavioral objectives specifying what is to be learned in that unit. Units generally correspond with chapters in a textbook, so that they are taken up in sequence. What makes PSI unique are the following characteristic features:

1. **Emphasis on individual study.** Students tackle course material on their own, often aided by study guides which provide practice on unit objectives (e.g., Johnson & Perkins, 1976). The teacher and any course aides serve as resources to students when they encounter difficulty understanding information or answering questions in the textbook or study guide.

2. *Self-pacing.* Students work at their own pace, and report to class only when they are ready to take a unit quiz. As a result, some students work quickly, finishing the course in half the semester or less, while other students who require more time to master concepts, take the entire semester to finish.
3. *Unit mastery requirement.* Students are required to meet a prespecified mastery level on each unit. When they take a unit quiz, they receive feedback immediately, and if unit mastery has not been achieved, they may take the quiz again with no penalty. Typically, three or four versions of a unit quiz are available to students, and individual records are kept, noting which version a student took at a given time.
4. *Use of proctors.* The requirement to provide immediate feedback to students regarding their quiz performance obviously means considerable work for the teacher. To alleviate this problem, proctors are used to score quizzes and provide feedback. Proctors may be advanced students who have already taken the course, or they may be students in the class who have already completed the unit they are now proctoring. Advantages to the latter arrangement include students solidifying their own knowledge of the material as well as getting to know their fellow students better.
5. *Supplementary instructional techniques.* Since the primary mode of instruction in a PSI course is self-study, lectures, demonstrations, and other modes of delivery may be used in a supplementary way to enhance motivation and transfer. Students may be motivated to reach a particular unit, for example, because mastery of the unit is their ticket to attend a special demonstration related to the next unit.

In the decade following Keller's proposal, PSI was tried in literally thousands of college courses. Kulik, Kulik, and Cohen (1979) reported that students generally liked PSI better than traditional courses, course grades were higher in PSI than in traditional courses, and student achievement on course final examinations was higher in PSI than traditional courses. There are several reasons, however, why PSI is not more popular currently.

Offsetting its effectiveness are the costs of PSI in time and resources. Preparation time is likely to be great initially, because study materials must be quiz-taking and multiple versions of quizzes written. Some arrangement for this often meant scheduling two rooms for a significant number of hours each week (which is not looked upon kindly by college administrators). Record-keeping can also be burdensome, since individual records must be kept on the progress of all students, and copies of all quizzes and keys must be accounted for.

Although problems of record-keeping and quiz-taking may be ultimately solved through the use of computers, other disadvantages of PSI are not so easily counteracted. Some students, for example, are simply unable to

meet the mastery criterion set for passing quizzes, despite repeated testing (e.g., Sussman, 1981). It may be that more moderate levels of mastery should be set (cf. Reiser, Driscoll, & Vergara, 1987), or that some students would better profit from alternative instructional presentations. Finally, self-pacing permits procrastination, which means that some students will not finish the requirements of a PSI course within the designated semester-long period. After several semesters of experience with a PSI course that I taught, I learned to reduce procrastination by limiting self-pacing. That is, quizzes were made available for a 3-week window, which essentially forced students to maintain a reasonable rate of progress.

Teaching Machines to Computer-Based Instruction. "Educational toys with feedback are to be found in patent files reaching back at least a hundred years," said Sydney Pressey in 1964 (Pressey, 1964, p. 354). So perhaps it is not entirely accurate to attribute teaching machines and programmed instruction solely to the influence of behaviorism, but certainly automation has been viewed as the solution to the problem of providing immediate reinforcement for correct responses in instruction. Although contingency contracts allow for reinforcement at task completion and PSI provides for immediate feedback on quiz performance, neither provides for sufficient reinforcement during learning. An automated teaching machine, however, has this capability.

After the early teaching machines of Pressey (1926, 1927), Skinner (1958) proposed applying behavioral principles to teaching academic skills through programmed instruction. In an instructional program, content is arranged in small steps, called frames, which progress from simple to complex and require a response from the learner to go on. Since the steps are small and increase gradually in difficulty, learners respond correctly most of the time, which means their responses are reinforced frequently. What this amounts to is shaping of complex academic skills.

A typical example of a programmed text can be found in Holland and Skinner's (1961), *The Analysis of Behavior*, an excerpt of which is shown in Figure 2.8. It should be obvious from this excerpt that early programmed instruction, despite providing immediate and frequent reinforcement, suffered from one serious flaw: It was boring. The small steps, for some students, were too small. Furthermore, all students had to work through the frames in the same order.

To improve on this linear style of program, Crowder (1960) introduced the notion of branching. In branching programs, frames are larger and are typically followed by questions with several possible answer options. Depending on how students answer a given question, they are branched to another segment of the program. In this way, students who know the material already may skip quickly ahead to new material. Likewise, students having

▽	Performing animals are sometimes trained with "rewards." The behavior of a hungry animal can be "rewarded" with _____.	7-1
Food	A technical term for "reward" is reinforcement. To "reward" an organism with food is to _____ it with food.	7-2
Reinforce	Technically speaking, a thirsty organism can be _____ with water.	7-3
Reinforced (NOT rewarded)	The trainer reinforces the animal by giving it food _____ it has performed correctly.	7-4

FIGURE 28 A Typical Example of Programmed Text

Source: Holland & Skinner, 1961, pp. 42-45.

difficulty with the instruction may be branched to remedial segments, which provide additional information and practice.

Computer-based instruction, as originally conceived, is simply programmed instruction presented via computer. The computer provides obvious advantages over text-based programmed instruction, which can be very cumbersome for both the writer of the program and the student. The computer allows for complex branching sequences and can automatically record through the instruction). Increased computer technology has also enabled program designers to include complex graphics and synthesized speech along with text. As a result, instructional software is increasingly available that provides drill and practice on various academic skills, simulations to enhance problem-solving, or tutorials in various subject matters.

Improving Performance in Organizational Systems

A focus on performance improvement in organizations is the professional orientation of a fast-growing field known as performance technology (Slojovich & Keeps, 1992; Kaufman, Thiagarajan, & MacGillis, 1997). Behaviorism is commonly acknowledged as one of the primary origins of this hybrid field, contributing an emphasis on observable performance and the importance of incentives in shaping behavior.

In much the same way that behavior therapists have sought to manage individual behavior, performance technologists attempt to manage performance, usually within a team setting. "Managing performance (rather than judging or appraising it) is the key way in which managers can be successful through delegating effectively, gaining support, and building synergy with their team members" (Bell & Forbes, 1997). And several behavioral principles comprise the key to effective performance management. These include, for example, well-defined objectives for employee performance that are linked to the organization's business plan and regular feedback, consisting of knowledge of results and knowledge of progress (Spence & Hively, 1993). In the traditional behavioral paradigm, feedback is the consequence of a response, typically reinforcement for an appropriate behavior. From decades of research on feedback, however, we have learned the importance of the information value of feedback. That is, feedback not only reinforces a response, it also provides information to the learner as to how performance can be improved. Knowledge of results provides feedback as to the quality of a particular performance, and knowledge of progress provides feedback of performance over time.

In studies of performance improvement within organizations, feedback as an intervention appears to have a profound effect. For example, consultants to a senior center used public feedback to increase contributions to the center (Jackson & Matthews, 1995). Volunteers clipped coupons, stamped the name of the senior center on the back, and then put them on products in local grocery stores. When store patrons bought these products, they could choose to redeem the coupon or donate its value to the senior center. While this procedure alone brought in some donations, both the value and frequency of donations increased substantially when the stores began posting signs indicating the progress each week of the dollars donated through the "Coupons for Caring" program.

According to Dean, Dean, and Rebalsky (1996), feedback is one of several environmental factors that support or hinder exemplary performance in an organization. In a study examining perceptions about performance blocks, they found that two-thirds of employees and managers identified environmental factors such as feedback, clear guides to expected performance, resources, and appropriate incentives as their biggest performance blocks. In contrast, only one-third identified individual factors such as sufficient knowledge and motivation to do assigned tasks. Interestingly, when teachers were asked to identify performance blocks of their students, individual factors were cited more often than environmental ones (Dean, Dean, & Rebalsky, 1996). The researchers concluded that analysis of the work environment is critical for managing performance improvement.

Planning for performance improvement is a process analogous to planning for behavior management or modification (Table 2.2). Once the desired performance has been determined and the gap is identified between what it

TABLE 2.2 *Planning for Performance Improvement*

- Step 1: Determine desired performance and the gap between what is and what should be.
- Step 2: Identify appropriate rewards and incentives for performance.
- Step 3: Generate a plan for communicating performance goals and implementing incentives.
- Step 4: Carry out the plan.
- Step 5: Evaluate results and revise as necessary.

is and what it should be, appropriate rewards and incentives for performance can be selected (Zigon, 1997). Then a plan is generated, implemented, evaluated, and revised as necessary.

The Behaviorist Perspective on Learning: Issues and Criticisms

Behaviorism has contributed to a number of instructional innovations, and behavioral principles continue to be useful to professionals in a number of disciplines. But what of behaviorism's shortcomings? What aspects of learning does it fail to account for readily? What problems can be seen in the behavioral paradigm that suggest alternative theories should be explored? This chapter concludes with examination of these questions.

Verbal Behavior

The astute reader may have noticed that nowhere in this chapter has the learning of language been mentioned. Skinner maintained a long-standing interest in language, publishing an extensive operant analysis of language learning in 1957 with *Verbal Behavior* (Skinner, 1957). Skinner treated language as he did any other set of complex operant responses. He proposed that the verbal behavior of children is shaped, with appropriate verbal labels for objects and events being maintained through reinforcement as inappropriate ones are extinguished.

Skinner's position on language learning met with heavy criticism (e.g., Chomsky, 1959), and, indeed, accounting for certain kinds of utterances is difficult. Although a child's learning to call only cows by that label may be set up under operant analysis, consider a sentence such as, "I am looking for my glasses" (Leahey & Harris, 1997). Our immediate reaction to such a statement is to explain it in terms of what the person is thinking. He has an image

of his glasses, which he has misplaced and is now trying to find. But Skinner, not permitting references to thought or mind, would argue that the stimulus in control of the verbal statement is the person's observation of his own searching behavior. That is, searching behavior in the past has resulted in the person finding his glasses and stopping the behavior; so he has learned to say, "I am searching for my glasses" as a response to this stimulus situation (Leahey & Harris, 1997).

This account of language learning seems a bit weird (Malcolm, 1964), and not all modern behaviorists adhere to it. Schoenfeld (1993) argued that behaviorists and nonbehaviorists alike must agree on the objective physicality of verbal behavior and the fact that language is learned within one's sociocultural environment. What differs between them, he claims, is their explanation for how each culture does its teaching.

Reinforcement and Human Behavior

While Skinner was interested in deriving functional laws of learning, i.e., the probability of behavior is increased when it is followed by reinforcement, some researchers wondered why reinforcement operates as it does. Why are some consequences of behavior reinforcing when others are not? Shedding light on this question are results summarized by Leahey and Harris (1997) on the use of different sorts of reinforcement schedules with humans. In order for human learners to exhibit the response patterns characteristic of certain reinforcement schedules, they had to be instructed as to the schedule in effect. Moreover, when given false information about what schedule would be in effect, human subjects responded according to what they believed was going on and not according to the actual manipulation (cf. Brewer, 1974).

In attempting to explain why reinforcement works, Estes (1972) provided an important link between behaviorism and later cognitive conceptions of motivation. Estes reviewed studies indicating that humans must have an expectation of being rewarded in order for reinforcement to work, and they must value the reward. As we will see in later chapters, the concepts of expectancy and value will play major roles in social learning theory (e.g., Bandura, 1986) and motivation (e.g., Keller, 1983).

Intrinsic Motivation

Finally, problems cropping up with the behaviorist conception of reinforcement were only further exacerbated by investigations into the notion of intrinsic motivation. It seems obvious to the casual observer that learners sometimes do things without ostensibly being reinforced. Some children

spend hours reading, for example, because they "like to read." Others might spend days putting together jigsaw puzzles "just for the fun of it." Skene would explain this behavior by referring to the reinforcement history of the individual. Some time ago, he would argue, the sources of reinforcement for that behavior were undoubtedly external (e.g., the child's parents praise her for reading and spend time reading with her). Over time, however, internal reinforcers became associated with the behavior and became conditioned reinforcers to sustain it.

Skinner's account of motivation, like his ideas about language, are with criticism. Bates (1979) reviewed studies which demonstrated how intrinsic satisfaction can even be undermined by extrinsic reinforcement. When rewards were given to learners for behavior in which they had already engaged on their own (e.g., puzzle solving or creating artwork), their response rate went down. This supports the notion that reinforcement is not necessarily a straightforward affair. In Chapter 9, the topic of motivation is taken up in greater detail.

Conclusion

Perceived problems and limitations with radical behaviorism as an explanatory paradigm for learning have led many investigators to propose cognitive, neurological, developmental, and other theoretical constructs as alternative ways of understanding learning. To Skinner, reliance on internal mechanisms of learning has led psychology away from a science of behavior to "questions that should never have been asked" (Skinner, 1987, p. 75). And he argued for a return to consideration of behavior "as a subject matter in its own right" (Skinner, 1987, p. 780).

Yet, "all psychological research is essentially behavioral," claimed Bornstein (1988), "in that psychological data inevitably take the form of observable, measurable behaviors, whether those are conditioned responses, responses to questionnaire items, or descriptions of inkblots" (pp. 819-820). In a commemorative issue of *Psychological Review*, Kimble (1994) and Thompson (1994) agreed, citing modern behaviorism as a "sophisticated statement of the scientific method applied to the study of behaving organisms. From Watson to the present day, the emphasis has been on measurement" (Thompson, 1994, p. 263).

"Because behavioral data must ultimately serve as the dependent variable in all psychological research, however, it does not necessarily follow that internal states, causes, and motivations are inappropriate or misleading constructs" (Bornstein, 1988, p. 820). In the chapters that follow, these core explanatory value in understanding learning and their usefulness for planning effective instruction.

A Behaviorist Perspective on "Kermit and the Keyboard"

How might a behaviorist explain learning in the story of Kermit and the keyboard? If learning amounts to behavior change, then the first step is to look for what behaviors are being exhibited and how they have changed. Two behaviors that are easy to spot are that Kermit selects songs and then plays them. Let's examine those behaviors in detail. Are they increasing or decreasing in frequency? What is the consequence of each that might reveal the contingencies of reinforcement that are operating?

Selecting Songs

Behavior change: Increasing for some, decreasing for others
Consequence: Kermit plays some songs that he selects with ease, whereas he makes a lot of mistakes on other songs.

Contingencies of reinforcement: Kermit selects more often those songs that he can play easily and less often those on which he makes mistakes. Thus, particular song selection appears to be positively reinforced by playing well and punished by making mistakes.

Playing Songs

Behavior change: Here we see that the time Kermit spends playing appears to vary.

Consequence: As with song selection, the consequence of time spent playing is either performing well or making mistakes.

Contingencies of reinforcement: Kermit plays longer when he is playing well (positive reinforcement) but stops when he makes a lot of mistakes (punishment).

There is also evidence of shaping in this story, in that Kermit first practiced "House of the Rising Sun" very slowly and gradually increased the tempo until he could perform the song as it written without making any mistakes. The mistake that he continues to make while using the one accompaniment was reinforced by this arrangement but not by other arrangements he has tried. It is likely as well that chaining has taken place. Chaining would occur if he practiced a portion of each song individually and then put the sections together to successfully play a complete song.

Some aspects of this story are hard to explain by using behaviorist theory. Why, for instance, did Kermit choose to learn the keyboard in the first place? Motivation is usually explained in terms of reinforcement history. However, Kermit has had no prior experience with either the keyboard or a