

7

Teaching by Giving Productive Feedback

Chapter Outline

A Response Learning Task

The Law of Effect

How Do Classroom Management Techniques Affect Classroom Behavior?

How Do Rewards Affect Classroom Activities?


How Does Feedback Affect Response Learning?

How Does Feedback Affect Skill Learning?

How Does Deliberate Practice Affect Complex Skill Learning?

Chapter Summary

When a learner is practicing a cognitive skill, a teacher can provide feedback to help promote meaningful learning. This chapter focuses on the educational implications of Thorndike's famous law of effect, the idea that learning occurs when a learner does something and receives feedback. Since the early part of the 20th century, and even continuing on through today, classroom instruction has probably been more strongly influenced by the law of effect than by any other idea in the psychology of learning. To help you better understand this principle and its relation to instruction, this chapter investigates how the law of effect can be related to classroom management and classroom activities. Then we explore the role of feedback in response learning and skill learning. For each topic, the chapter contrasts a behaviorist versus cognitive view of the learner and the learning process. Providing feedback is a long-standing method of instruction that dates back to the start of educational psychology. The premise of this chapter is that all forms of feedback are not equally useful in promoting meaningful learning: Meaningful learning can be promoted when feedback is presented as information intended to guide the learner's construction of knowledge and instill intrinsic motivation, but meaningful learning is not promoted when feedback is presented as reinforcement intended to automatically increase or decrease a response.





A Response Learning Task


Let's try a simple learning task. You will need about 10 pennies (numbered 1 through 10) and a small square of paper (e.g., about 3 inches by 3 inches). Place the paper flat on the floor about 3 feet away from you. Then sit down with the pennies near you. With your eyes closed, toss each penny toward the square, trying to get the penny as close to the square as possible. After you have tossed 10 pennies, take a ruler and measure how many inches each penny landed away from the square. (If a penny landed on the square, its score is 0.) Now fill in the left portion of Figure 7-1. Place a dot indicating how far away the first penny landed from the target, how far the second penny landed, and so on for all 10 pennies. This is your *baseline* performance (i.e., your performance before any training).

Now, let's see whether you can teach yourself to be more accurate. Collect the 10 pennies and sit down where you sat before. With your eyes closed, toss a penny toward the square. Then open your eyes and measure how far off you were. Close your eyes and toss the next penny; then open your eyes and measure how far off you were. Repeat this procedure for 20 tosses. Now fill in the middle portion of Figure 7-1. Place a dot indicating how far away the first penny landed, and so on for all 20 pennies. This is your *training* performance (i.e., your performance during the training period).

Finally, repeat the procedure you used for the baseline. Sit down and, with your eyes closed, toss each of the 10 pennies. Then measure how far off each one was and fill in the results on the right-hand portion of Figure 7-1. This is your *extinction* performance (i.e., your performance after training, when feedback is no longer given for each response).

Are you more accurate in tossing pennies during the training period than during the baseline period? If so, we might be tempted to say that you are learning! Why does your performance increase during the training phase? One possible answer is that you had more practice. However, if you conducted the baseline part of the task for many tosses, you might find no

FIGURE 7-1
Fill in your performance on a simple motor learning task



Source: Adapted from Mayer, R. E. (1997). *Incorporating problem solving into secondary school curricula*. In G. D. Phye (Ed.), *Handbook of academic learning*. (pp. 474-492). San Diego, CA: Academic Press.

improvement in your performance. Practice—by itself—might not be enough to enable learning. Thus, a second explanation is that during the training period you had practice with feedback. The feedback tells you how far away each penny landed from the target. Thus, knowing the effects of your behavior seems to be a potent force in helping you change your behavior.

Does your accuracy decrease during the extinction period? If so, we can say that feedback is needed for you to maintain your performance. If you are able to maintain a high level of accuracy even during the extinction period, we can say that you have learned to provide intrinsic feedback. That is, you have learned to be able to tell whether your performance is accurate without external feedback.

How does feedback affect learning? This chapter explores two views of feedback—feedback as reinforcement and feedback as information. According to the *reinforcement* view, which is based on behaviorist approaches to learning, feedback serves to strengthen or weaken responses or, more specifically, the association between a stimulus and a response. In short, rewards make the association to a response stronger, whereas punishments make the association weaker. The process of response strengthening or weakening is automatic, not requiring any cognitive interpretation by the learner. In contrast, the *information* view, which is based on cognitive approaches to learning, holds that feedback is information that the learner interprets in order to make sense of a learning episode. Learning depends not on the feedback but rather on the learner's interpretation of the meaning of the feedback. If you are interested in promoting meaningful learning, feedback is useful to the extent that it provides information that guides the learner's construction of knowledge.

In this chapter, you will learn how feedback can be used to promote meaningful learning in each of several kinds of learning situations: classroom management of disruptive behaviors, students' selection of classroom activities, response learning, and skill learning. Classroom management refers to decreasing students' disruptive behaviors, such as shouting out or pushing other students. Classroom activities refers to promoting students' selection of appropriate activities, such as choosing to do an art project during a free period. *Response learning* refers to the acquisition of a new response, such as tossing a coin so that it lands on a target. In response learning, the outcome of learning is a change in a single aspect of the learner's behavior. Academic examples of response learning include learning to say "six" when you are presented with a flash card that states " $3 + 3$," learning to raise your hand before you speak in class, or learning to turn a handle on a water fountain to get a drink. *Skill learning* refers to the acquisition of a new procedure—or set of steps—for doing some task. In *simple skill learning*, the outcome of learning is being able to carry out a procedure, like being able to compute answers for two-column multiplication problems like $35 \times 67 = \underline{\hspace{1cm}}$. In *complex skill learning*, learners develop expertise in performing authentic tasks such as playing a piano piece.




The Law of Effect

THORNDIKE'S THEORY

To get a better understanding of response learning, let's consider an early study performed by E. L. Thorndike (1898, 1911). Figure 7-2 shows a puzzle box used by Thorndike. As you can see, the puzzle box is a cage containing a door that can be opened from inside the

FIGURE 7-2
A puzzle box

"When put into the box the cat would show evident signs of discomfort and of impulse to escape from confinement. It tries to squeeze through any opening; it claws and bites at the wire; it thrusts its paws out through any opening outside but seems simply to strive instinctively to escape from confinement. . . . The cat that is clawing all over the box in her impulsive struggle will probably claw the string or loop or button so as to open the door. And gradually all the other unsuccessful impulses will be stamped out and the particular impulse leading to the successful act will be stamped in by the resulting pleasure, until, after many trials, the cat will, when put in the box, immediately claw the button or loop in a definite way."




Source: From Thorndike, E. L. (1898). *Animal intelligence: An experimental study of the associative processes in animals*. Psychological Review, Monograph Supplement, 2(8), 13.

cage. A hungry cat was placed inside the puzzle box with a dish of food outside. In a typical experiment, the cat had to learn to unlatch the door by pulling a string. The first time the cat was put into the puzzle box, the cat engaged in many irrelevant behaviors, including scratching the bars, meowing, and pouncing on the floor. Eventually, the cat pawed at the string and accidentally pulled it down; this action opened the door and allowed the cat to escape and eat the food in the dish. On subsequent trials, the cat tended to spend less time engaging in irrelevant behaviors such as pouncing and meowing and tended to require less time before pulling the string. After many trials, the cat would immediately pull the string upon being placed into the puzzle box. Figure 7-3 shows the time (in seconds) the cat spent in the puzzle box before pulling the string for each of 20 trials. As you can see, the cat required a minute or more on the first two trials but required just a few seconds by the 20th trial.

The change in the cat's performance suggests that the cat learned something. According to Thorndike, the cat learned to form a strong association between the stimulus—being in the puzzle box—and a response of pulling the string. At the start of the experiment, the cat had many responses associated with the stimulus, including pouncing, meowing, and scratching. However, each time the cat engaged in one of these behaviors, the result was not positive; thus, according to Thorndike, the association between the stimulus

FIGURE 7-3 Time spent in the puzzle box on each trial



Source: Adapted from Thorndike, E. L. (1898). *Animal intelligence: An experimental study of the associative processes in animals*. Psychological Review, Monograph Supplement, 2(8).

(i.e., being in the puzzle box) and the response (e.g., meowing) was weakened. Similarly, each time the cat pulled the string, the result was positive; thus, according to Thorndike, the association between the stimulus (i.e., being in the puzzle box) and the response (i.e., pulling the string) was strengthened. At first, the response of pulling the string was very weakly associated with the stimulus of being in the box, but by the end of the experiment, the association was strong. At first, several irrelevant behaviors were strongly associated with being in the box, but by the end of the experiment, those associations were weaker. Thus, according to Thorndike, learning to make an appropriate response involves strengthening the relevant association and weakening the other associations. The principle underlying this learning is called the *law of effect*, because learning depends on the effect of each response. The law of effect is the idea that if a behavior is followed by a pleasing state of affairs, it is more likely to occur again in the future under the same circumstances, and if a behavior is followed by a displeasing state of affairs, it is less likely to be given again in the future. Here is how Thorndike (1911) summarized the law of effect:

Of several responses made to the same situation, those which are accompanied or closely followed by satisfaction . . . will, other things being equal, be more firmly connected to the situation, so that, when it recurs, they will be more likely to recur; those which are accompanied or closely followed by discomfort . . . will, other things being equal, have their connections with the situation weakened, so that when it recurs, they will be less likely to occur. The greater the satisfaction or discomfort, the greater the strengthening or weakening of the bond. (p. 244)

The law of effect is widely recognized as one of the consistent pillars of the psychology of learning, and no principle of learning has had a greater impact on education than the law of effect. However, as you will see, in spite of the demonstrated effectiveness of practice with feedback, the preponderance of research evidence shows that something is terribly wrong with Thorndike's explanation of how learning happens. Throughout this chapter, you should be seeking to answer the question, "Is feedback (such as rewards and punishments) a form of reinforcement that automatically strengthens or weakens responses (as Thorndike posited), or is feedback a form of information that the learner consciously interprets to make sense of the learning episode (as cognitive theory posits)?"

SKINNER'S THEORY


Skinner (1938, 1953, 1957, 1968, 1969) conducted extensive research that served to improve on Thorndike's methodology and modify his theories. For example, with respect to methodology, Skinner developed devices that have been called *Skinner boxes*. Figure 7-4 shows a Skinner box for a white laboratory rat. The Skinner box consists of a metal cage, a metal grid floor, a bar that can be pressed down, and a food tray that is connected to a machine that delivers food pellets. First, let's put a white lab rat into the Skinner box. The rat will engage in many behaviors but will rarely press the bar. Thus, the baseline rate of bar pressing is very low, as shown in the left portion of Figure 7-5. Now let's adjust the food pellet machine so that it delivers a food pellet to the food tray each time the bar is pressed. At first, the rat—like Thorndike's cat—engages in many irrelevant activities. Eventually, the rat—like Thorndike's cat—will accidentally produce the desired behavior. As soon as the rat presses the bar, a food pellet appears in the tray and the rat presumably will eat the pellet. Over the course of this learning session, the rate of bar pressing will increase dramatically, as shown in the middle portion of Figure 7-5. The Skinner box is an improvement over Thorndike's puzzle box because the animal does not have to be placed back into the box after each successful response.

To describe the learning process, Skinner developed *reinforcement theory*. The main idea in reinforcement theory is that reinforcers can control behavior. A reinforcer is any stimulus that, when contingent on a response, serves to increase the rate of responding. This definition has two components:

1. *Contingency*. The occurrence of the reinforcer depends on (e.g., follows immediately upon) the occurrence of the learner's response.
2. *Rate of responding*. The reinforcer serves to increase the learner's rate of responding.


The food pellet fits the definition of a reinforcer because (1) the food pellet is delivered immediately after the bar is pressed (i.e., the delivery of the food pellet is contingent on the learner's pressing the bar), and (2) the rate of bar pressing increases when bar pressing is reinforced.

FIGURE 7-4
A Skinner box



Courtesy of Ralph Gerbrands Co., Arlington, MA.

FIGURE 7-5 Rate of bar pressing for each 10-minute time block in the Skinner box



IS REWARD MORE EFFECTIVE THAN PUNISHMENT?

A reinforcer may be positive or negative. A positive reinforcer is anything that, when presented contingent on a response, tends to increase the probability or rate of that response. A negative reinforcer is anything that, when taken away contingent on a response, tends to increase the probability or rate of that response. Presenting a food pellet to a hungry rat is an example of a positive reinforcer, whereas stopping an electrical shock is an example of a negative reinforcer. Thus, a response can be reinforced by either presenting a reward or taking away a punishment.

In contrast, two kinds of stimuli may serve to decrease the rate or frequency of responding: (1) presenting an aversive stimulus contingent on a response, and (2) taking away a rewarding stimulus contingent on a response. You can consider these to be forms of punishment. For example, the rat might decrease its rate of bar pressing if a shock follows each bar press or if pressing a bar causes a machine to stop supplying food pellets for 60 seconds.

Although the original law of effect as presented in the preceding section gave equal status to reward and punishment, this version did not survive Thorndike's lifetime. Toward the end of his scientific career, after hundreds of studies, Thorndike (1932) felt compelled to revise the law of effect by downplaying the role of punishment in changing behavior:

In the early statements of the Law of Effect, the influence of satisfying consequences of a connection in the way of strengthening it was paralleled by the influence of annoying consequences in the way of weakening it. . . . I now consider that there is no such complete and exact parallelism. In particular, the strengthening of a connection by satisfying consequences seems in view of our experiments . . . to be more universal, inevitable, and direct than the weakening of a connection by annoying consequences. (p. 276)

Skinner's later application of reinforcement theory to instruction was also based on the power of rewarding appropriate behavior through positive reinforcement rather than punishing incorrect behavior. Thus, the remainder of this chapter focuses mainly on the uses of reward in school-related tasks and how feedback affects meaningful learning.



How Do Classroom Management Techniques Affect Classroom Behavior?

CLASSROOM MANAGEMENT

As you may have noticed, the law of effect (or reinforcement theory) suggests many classroom applications, including classroom management—techniques for maintaining a productive learning environment in the classroom. In any classroom, students will exhibit some behaviors that are disruptive and some behaviors that are productive for learning. One goal of a classroom management program is to decrease the frequency of disruptive behaviors and increase the frequency of productive behaviors. As you can see, this situation seems similar to that of Thorndike's puzzle box or the Skinner box because we have some target behavior that we want to increase or decrease. Thus, a major question is whether the principles of reinforcement developed from animal research can be applied successfully to humans. In this section, we explore two examples of how the law of effect can be applied to the task of classroom management—contingency contracting and token economies—after briefly exploring the predictions of behaviorist and cognitive theories.

BEHAVIORIST VERSUS COGNITIVE THEORIES OF CLASSROOM MANAGEMENT

According to a behaviorist view, classroom management techniques automatically change classroom behaviors: Behaviors that are punished will gradually decrease, and behaviors that are rewarded will gradually increase. According to cognitive theories, classroom management techniques serve as information that the learner interprets and uses for building useful mental representations of the situation. Punishments serve to reduce a behavior when a learner understands the cause-and-effect chain leading to negative consequences (i.e., "if I do this, the result will be that") and wants to avoid a negative consequence. Rewards serve to increase a behavior when a learner understands the cause-and-effect chain leading to positive consequences (i.e., "if I do this, the result will be that") and wants the positive consequence.

RESEARCH ON CLASSROOM MANAGEMENT


Contingency Contracting *Contingency contracting* is one application of reinforcement techniques to a classroom situation. In contingency contracting, the student and teacher make a sort of contract that specifies which student activities will lead to which rewards or

punishments. For example, Sulzbacher and Houser (1968) conducted a study involving children in a special education class. The first step was to define the behavior to be changed; in this study, the teacher attempted to decrease the frequency of children making a disruptive gesture referred to as the "naughty finger." The next step was to determine a contingent event that might serve to decrease the rate of responding; in this study, the teacher used a loss of 1 minute of recess time from a 10-minute recess. As you can see, this contingency seems to involve both punishment and reward; for each inappropriate behavior, a rewarding stimulus (1 minute of recess) is taken away, but for appropriate behavior, no time will be lost from recess. This procedure is called *response cost* because inappropriate behavior costs the student in terms of diminishing a reward. Another frequently used contingency is called *time-out*—for each case of inappropriate behavior the child is moved to a quiet place for a short time. These techniques seem to be more useful than traditional types of punishment, such as verbal disapproval or threats. The third step in classroom management is to establish a clear procedure for administering the program, with clearly articulated rules.

For 9 days, the teacher simply kept track of the frequency of occurrences of the naughty finger. This is the baseline period, and as you can see in Figure 7-6, the rate of responding for the class was about 15 occurrences per day. At the beginning of the 10th day, the teacher began the training period by making the following announcement:

From now on there will be a special ten-minute recess at the end of the day. However, if I see the naughty finger or hear about it I will flip down one of these cards, and you will have one minute less of recess whenever this happens. Remember, every time I flip down one of these cards, all of you lose a minute from your recess. (Sulzbacher & Houser, 1968, p. 88)

FIGURE 7-6
Number of
disruptive behaviors
per day in a
special education
classroom



Source: From Sulzbacher, S. I., & Houser, J. E. (1968). A tactic to eliminate disruptive behaviors in the classroom: Group contingent consequences. *American Journal of Mental Deficiency*, 73, 88-90.

As you can see in the middle portion of Figure 7-6, the number of occurrences of the naughty finger fell to fewer than five per day. After 18 days of training, the teacher told the class that the recess policy no longer was in effect. As you can see in the right portion of Figure 7-6, during the first 9 days of this extinction period, the rate of responding increased. By the 9th day, the rate was back up to where it had been before training!

Two aspects of Figure 7-6 merit special attention. First, note that the rate of responding fell dramatically on day 10. This fall cannot be attributed to the training, because training had just begun; instead, the students seem to be able to adjust their performance based on the verbal instructions of the teacher. According to a behaviorist interpretation of the law of effect, the change in behavior should be gradual and should occur only after the punishments are administered. According to the cognitive interpretation of the law of effect, students use the information presented to build a mental model of the situation, which allows them to change their behavior. As you can see, the results are most consistent with the cognitive rather than the behaviorist interpretation. Second, note that the suppression of the disruptive behavior is short-lived. Once the contingencies are removed, the disruptive behavior seems to return to its former strength. Thus, punishment was not very effective in permanently changing behavior.

Token Economies Token economies are another application of reinforcement theory to a classroom situation. In a token economy, students are given tokens (e.g., check marks by their names on a blackboard) for appropriate behaviors, and these tokens can be traded in for prizes or privileges. Alternatively, tokens can be given for inappropriate behaviors, and these tokens are later exchanged for various levels of punishment. The first step is to describe clearly the behavior to be changed. For example, suppose a teacher wishes to reduce the amount of disruptive behavior in a classroom. For purposes of our example, let's define disruptive behavior as any one of the following behaviors: (1) motor behaviors (e.g., walking around the room), (2) aggressive behaviors (e.g., hitting or kicking another person), (3) disturbing someone else's property (e.g., grabbing another person's book or tearing another person's paper), (4) making noise (e.g., stamping feet or clapping hands), (5) verbalizations (e.g., blurting out answers or talking out of turn), (6) turning around (e.g., looking to the rear of the class when the teacher is in front of the class), or (7) inappropriate tasks (e.g., drawing a picture during a spelling lesson) (O'Leary, Becker, Evans, & Saudargas, 1969). The second step is to find a contingent stimulus that will serve as an effective reinforcer for the children who produce these behaviors. Weil and Murphy (1982) suggested three categories of rewards for classroom management:

1. *Social rewards* (e.g., smiles, praise, or hugs)
2. *Material rewards* (e.g., stickers, stars, and awards)
3. *Tokens* (e.g., tickets, passes, or check marks on a chart that can be redeemed for valued prizes or privileges)

For our example, let's rely on tokens. In our classroom token economy, a child receives from 1 to 10 points after each 30-minute lesson for each of four lessons on each afternoon. The points (or tokens) are placed in a small booklet on each child's desk. On redemption days, the child can trade in points for prizes, such as 25 points for a level 1 prize, 35 points for a level 2 prize, and so on. The third step is to implement the


program clearly and consistently. The rules are made clear to the student. For example, the teacher could write the following rules on the blackboard:

We sit in our seats; we raise our hands to talk; we do not talk out of turn; we keep our desks clear; we face the front of the room; we work very hard; we do not talk in the hall; we do not run; and we do not disturb reading groups. (O'Leary et al., 1969, p. 5)

The teacher reviews the rules regularly. In addition, the teacher puts the token system into effect during short structured afternoon activities, such as spelling or science lessons in which the entire class participates. As part of the token economy, the teacher praises appropriate behaviors and ignores inappropriate behaviors. Threats such as "If you're not quiet by the time I count to three . . ." and commands such as "Sit in your seat" have been eliminated.

The example just outlined is similar to a token economy program for second graders that was carefully studied by O'Leary et al. (1969). In a classroom of 21 students, the teacher identified seven students who tended to frequently produce disruptive behaviors such as those previously described. As part of the study, observers noted whether each of the seven children was engaged in appropriate or disruptive behavior during various 20-second periods throughout the school day. Figure 7-7 shows the percentage of time that the seven children were engaged in disruptive behaviors at different points in the study. First, there was a baseline period during which no changes were made in the classroom. The children were disruptive on more than half of the observations made. A few

FIGURE 7-7
Percentage of disruptive behaviors per day during the afternoon in a second-grade classroom




From O'Leary et al. (1969).

days later, the teacher introduced the classroom rules and wrote them on the board. This step did not result in any major change in disruptive behavior. A few days later, the classroom afternoon structure was changed into easily identifiable 30-minute lessons. Structure had no major impact either. A few days later, the teacher began a policy of praising appropriate behavior and ignoring disruptive behavior during the afternoon. The result was a modest increase in disruptive behavior.

Finally, a few days later, the teacher began the token economy system of presenting points to each child after each of four 30-minute lessons each afternoon. As Figure 7-7 shows, the token economy procedure resulted in a dramatic reduction of disruptive behavior. Then the token system was stopped. Disruptive behavior returned almost to its original baseline level. However, when the token economy was reestablished, disruptive behavior fell again. Finally, during a follow-up period, the teacher stopped giving tokens but continued using rules, structured lessons, and praise. In addition, the teacher gave stars for good behavior—with each day beginning with a clean wall chart for each child's stars. The final panel of Figure 7-7 shows that disruptive behavior increased but was still maintained at well below the baseline level. These results support the cognitive theory because behavior changed immediately.

The preceding discussion is based only on performance during the afternoon (i.e., the time when the token economy was in effect). During the morning, no token economy was ever in effect. Did the reduction in disruptive behavior during the afternoon transfer to the morning? Figure 7-8 summarizes the percentage of disruptive behavior during the morning for each phase of the study. As you can see, the token economy procedure did not cause a

FIGURE 7-8
Percentage of disruptive behaviors per day during mornings in a second-grade classroom



From O'Leary et al. (1969).

large change in disruptive behavior during the morning even though there was a large change in disruptive behavior during the afternoon. Apparently, the children were able to discriminate between conditions under which appropriate behavior resulted in rewards (i.e., afternoon) and conditions during which appropriate behavior did not result in rewards (i.e., mornings).

In a subsequent study involving six disruptive ninth graders, Main and Munro (1977) were able to replicate many of the results of the O'Leary et al. (1969) study. As in the O'Leary study, providing structured lessons and praising appropriate behavior while ignoring inappropriate behavior had only a slight effect on reducing disruptive behavior. Use of tokens greatly reduced disruptive behaviors; withdrawal of tokens produced an increase in disruptive behavior. During a final follow-up period, the token system was reestablished but tokens were gradually thinned out, leaving a contract system and teacher praise.

This follow-up procedure was successful in maintaining a relatively low level of disruption. Overall, there are many success stories involving the use of rewards and punishments to influence how children behave in classrooms (Schloss & Smith, 1998). However, it is useful to examine how and when such procedures work. From the examples reviewed in this section, we can begin to see the characteristics of a successful classroom token economy. First, the systems have clearly stated rules and structure. Students know what is expected of them and what will happen if they behave in certain ways. Second, the systems tend to use reward rather than punishment. For example, Bandura and Walters (1963) found that when a model (e.g., a teacher) uses punishment on a child, the child is more likely to use aggression as a way of controlling others. Madsen, Becker, and Thomas (1968) found that verbal punishments (e.g., saying "Stop that!") tend to increase disruptive behaviors. Similarly, ignoring disruptive behavior does not help to reduce it. However, a time-out—such as having to sit for 3 minutes watching a sand timer—was effective in reducing disruptions in a physical education class (White & Bailey, 1990). Third, the systems are administered in a way that always allows for fairly rapid recognition of appropriate behavior. Students receive feedback within a few minutes of their behavior, and prior inappropriate behavior cannot cancel out rewards for current appropriate behavior. In the first example, feedback was given to the entire class; in the second example, feedback was given to each student separately. Overall, classroom management techniques require that the students understand why they are being rewarded or punished, so more is involved than simply dispensing rewards and punishments.

Is reward more effective than punishment in changing disruptive classroom behaviors? To help answer this question, Conyers et al. (2004) visited a preschool classroom consisting of 4- and 5-year-olds. During a 30-minute baseline period, every 10 seconds they recorded whether they observed a disruptive behavior—such as screaming, crying, throwing objects, using objects as weapons, or refusing to comply with a teacher's request. On average, there was at least one disruptive behavior during 64% of the 10-second intervals in the baseline period. Then, starting on day 11, the teachers instituted a token economy by putting up a board with 15 stars next to each child's name. Students were told that they would get a prize (such as a choice among jelly beans, chocolate, or Gummi Bears®) if they had 12 stars by the end of the session (which lasted 15 minutes). Students were shown that a timer was started, and were told that they would lose one star every time they engaged in disruptive behavior. The timer was set for about 1 minute and when the timer rang, students who had engaged in disruptive behaviors were shown that they had lost a star and were told what they had done wrong. Then the timer was set again, and the process was repeated for a total

of 15 minutes. At the end of the session, children who had 12 stars could get a prize. This treatment is a form of punishment called *response cost* in which students lose a star for each disruptive behavior they produce. On the next day, students were shown a board with 15 spaces next to each child's name, and were told they would get a star added for each period in which they produced no disruptive behaviors. The timer was set for about 1 minute and when the timer rang, students who had not engaged in any disruptive behaviors saw a star added next to their name. The timer was set again and the process was repeated for a total of 15 minutes. At the end of the session, children who had 12 stars could get a prize. This treatment is a form of reward in which students earn a star for each interval in which they do not misbehave. The teachers alternated between the response cost (i.e., punishment) treatment and the reward treatment for 34 days, and then for the next 5 days (i.e., extinction period) they stopped giving prizes for stars.

FIGURE 7-9
Percentage of
10-second
intervals with
disruptive
behaviors during
the baseline,
treatment, and
extinction periods




Figure 7-9 shows the percentage of 10-second intervals with disruptive behaviors during the baseline, treatment, and extinction periods. As you can see, the rate of disruptive behaviors was high during the baseline period and was low during the treatment period—both for days in which response cost was used and days in which reward was used. Overall, it appears that both reward and punishment were effective in decreasing disruptive behavior.

On the surface, you might conclude that the law of effect works as Thorndike proposed: Responses are weakened when they are punished (such as losing a star for shouting) and responses are strengthened when rewarded (such as gaining a star for talking quietly). Let's look at two aspects of the data a little more closely—what happens on the first few days of the treatment (when the contingencies are in effect) and what happens during the extinction period. According to the behaviorist interpretation of the law of effect, the learning should be gradual during the treatment period because responses are strengthened or weakened as a consequence of reinforcement. However, as you can see, the rate of disruptive behaviors falls greatly in the first days that the treatments are in effect, suggesting that students can adjust their behavior based on their understanding of the new classroom rules. This pattern of a rapid decline in disruptive behavior is more consistent with the cognitive interpretation of the law of effect, in which people change their behavior based on their interpretation of the situation. In addition, during the extinction period, the rate of disruptive behavior immediately shoots back up to somewhat high levels. Again, this pattern is not consistent with the behaviorist interpretation of the law of effect and is more consistent with the cognitive interpretation of the law of effect.

IMPLICATIONS OF RESEARCH ON CLASSROOM MANAGEMENT

The limitations of classroom management systems are clear from the preceding examples. It is difficult to establish a program that works (i.e., a program that eventually can end the constant use of material rewards and punishments). Students may come to rely heavily on external rewards and punishments as a guide to how they should behave; instead, a goal of many classroom management systems is to help children develop intrinsic or self-motivated methods of behavior control. When the goal is to promote transfer—that is, to encourage productive changes in behavior that will last—classroom management techniques should be seen as information that is used by the learner to understand how the classroom community works.



How Do Rewards Affect Classroom Activities?

HIDDEN COSTS OF REWARD FOR CLASSROOM ACTIVITIES

Suppose a teacher wants to encourage students to engage in certain learning activities such as reading books or drawing pictures? A straightforward approach based on the law of effect would be to offer rewards for participating in the targeted activities. If a teacher wants a student to read, she can give the student a gold star for each book the student reads. If a teacher wants a student to draw a picture, the teacher can offer a certificate for completing an art project. What's wrong with such a straightforward application of the law of effect? I explore some hidden costs of reward when using contingency contracting and token economies to promote classroom activities that are already valued by the students.

BEHAVIORIST VERSUS COGNITIVE THEORIES OF REWARD

There are many how-to books for classroom management. However, to effectively use the techniques discussed in these books, it is important to understand the theory of how reinforcers (e.g., tokens) affect behavior. In this chapter, we explore two kinds of theories of reinforcement—behaviorist theories and cognitive theories. Behaviorist theories are based on Thorndike's and Skinner's interpretations of the law of effect, namely, the idea that a reinforcer (such as a token) serves to stamp in (reinforce) a response (such as appropriate behavior) and its link with a stimulus (such as the classroom environment). The strength of the stimulus-response association is automatically increased each time the response is reinforced; the process does not require active awareness or interpretation by the learner. In contrast, cognitive theories are based on an alternative interpretation of the law of effect, namely, the idea that the learner actively thinks about and interprets the feedback and its relation to the learner's response. Thus, feedback in cognitive theory serves as information that the learner uses in building a plan for responding.


RESEARCH ON HIDDEN COSTS OF REWARD

Hidden Costs of Contingency Contracting Let's examine what these two theories predict in a simple reward situation. Suppose we observe a preschool to see how children spend their free time. We find that many of the children spend a large portion of their free time in a drawing activity—using colored marking pens and large sheets of drawing paper. Let's take those children who like to spend free time drawing and assign them to one of three groups. In the expected reward condition, we will make a contingency contract with each child: The child agrees to produce some drawings in exchange for an extrinsic reward (i.e., a certificate with a gold seal and ribbon). In the unexpected reward condition, no contract is made, but the child is given the same reward (i.e., a certificate) after drawing some pictures. However, the child does not know in advance that a reward will be given for drawing. Finally, in the no-reward condition, the child does not receive a reward (i.e., certificate) for engaging in the drawing activity.

To see how rewards affect behavior, let's examine how our preschoolers spend their free time after the rewards are given. Let's come back to the preschool a week or two later and observe the percentage of free time spent in drawing. What do the behaviorist and cognitive theories predict? The behaviorist theory would predict that children in the expected reward condition and the unexpected reward condition should spend more free time drawing because they were rewarded for drawing. In contrast, the cognitive theory predicts that the expected reward condition children will spend less time drawing than children in the other two groups. One cognitive theory—called *overjustification* (Lepper & Greene, 1978)—states that the children are interpreting or justifying their own behavior and extrinsic rewards. If a child engages in a behavior, that behavior must be justified by the child. Thus, children who draw without expecting to receive a reward (no reward and unexpected reward conditions) can justify their behavior by saying they draw because they like to draw; in contrast, children who draw with the expectation of receiving a certificate (expected reward condition) can justify drawing by saying they draw because they get something for it. When rewards are no longer given, the justification for engaging in the drawing behavior is reduced for these children.

In a classic study, Lepper, Greene, and Nisbett (1973) carried out the procedure just described. Figure 7-10 summarizes the percentage of free time spent in the drawing

FIGURE 7-10
Some hidden costs
of reward for
preschoolers



Adapted from Lepper, M. R., Greene, D., & Nisbett, R. E. (1973). *Undermining children's intrinsic interest with external rewards: A test of the overjustification hypothesis*. *Journal of Personality and Social Psychology*, 28, 129–137.

activity for children who had experienced the three treatment conditions. As you can see, the results conflict with the behaviorist theory, which states that the two rewarded groups should behave similarly—but is consistent with the cognitive theory, which states that children in the expected reward condition (or expected award condition) should spend less time drawing than children in the other groups. Apparently, when children are rewarded for engaging in a behavior that is already interesting to them, the effect can be to diminish their intrinsic motivation for engaging in that behavior. Lepper and Greene (1975) refer to this process as “turning play into work.”

Hidden Costs of Token Economies This research shows that there can be some “hidden costs of reward” (Lepper & Greene, 1978). Based on their findings, Lepper et al. (1973) recommend caution in using token economies, especially when rewards are given out for behaviors that are already intrinsically satisfying to children:

It has already been recommended by some thoughtful proponents of token economies that their use be limited to circumstances in which less powerful techniques have been tried and found inadequate—in other words, only when they are necessary. . . . The present study provides some empirical evidence of an undesirable consequence of the unnecessary use of extrinsic rewards, supporting the case for the exercise of discretion in their application. (p. 136)

Greene, Sternberg, and Lepper (1976) obtained similar results in a classroom token economy. Fourth and fifth graders were given rewards for spending time on certain math activities (such as creating geometric designs), some of which were already interesting to the students. Once the reward phase of the study ended, students could spend their math time as they wished. Students tended to spend less time with the previously rewarded activities than did a control

group that received no rewards. In most cases, students who had been rewarded for engaging in certain math activities spent less than half as much time on those activities during the withdrawal phase as they had during the baseline phase. Again, the hidden cost of reward can be a drop in students' intrinsic motivation for rewarded activities. Similar results have been obtained by Ross (1975), Lepper and Greene (1975), and Deci (1971).

In a review of research, Cameron and Pierce (1994) found evidence that was consistent with the results of Lepper et al. (1973): Across 50 comparisons, students who received expected rewards averaged lower intrinsic motivation than nonrewarded students—as indicated by spending less of their subsequent free time on the rewarded task. Cameron and Pierce (1994) conclude: “Expected tangible rewards produce a decrease in intrinsic motivation as measured by free time on task when they are given to individuals simply for engaging in the task” (p. 394). In another large-scale review of research, Deci, Koestner, and Ryan (1999) found that “tangible rewards . . . typically decrease intrinsic motivation for interesting activities” (pp. 658–659). In examining previous findings, Cameron, Banko, and Pierce (2001) concluded that “negative effects are found for high-interest tasks when rewards are tangible, expected (offered beforehand), and loosely tied to level of performance” (p. 1), all conditions that were present in Lepper, et. al's original study.

These analyses should not be taken to mean that reward always hurts intrinsic motivation. For example, in a recent review, Cameron and Pierce (2002) provide consistent evidence that giving rewards does not decrease intrinsic motivation in some situations, including when the reward is not expected or when rewards are offered for successful performance rather than participation. Sometimes, offering rewards for high task performance can even increase intrinsic motivation (Cameron & Pierce, 2002; Eisenberg, Pierce, & Cameron, 1999). Instead of asking whether rewards increase or decrease intrinsic motivation, researchers are now seeking to understand the conditions under which rewards can affect intrinsic motivation (Cameron & Pierce, 2002; Eisenberg et al. 1999; Lepper, Keavney, & Drake, 1996; Pintrich & Schunk, 1996).

What does this research on rewards tell us about teaching for meaningful learning? An important way to promote teaching for meaningful learning is to instill a sense of intrinsic motivation that propels the learner when confronted with new academic tasks. Research on reward shows us that there are serious limitations in the classic view of rewards as ways of strengthening specific behaviors. Instead, this line of research shows that rewards (and punishments) constitute information that is interpreted by the learner. These interpretations affect the nature of what is learned and therefore the potential for transfer.

For example, when a child receives an expected reward for doing something he or she already likes to do, the child's interpretation of the reward may be, “I really do not like doing this after all.” The result can be a decrease in motivation to perform the task and hence an impediment to transfer. As another example, when a child receives an unexpected reward for doing something he or she likes to do, the child may interpret the reward to mean “I really like doing this, and I get recognized for it.” The result can be an increase in intrinsic motivation, or at least no decrease, hence no negative consequences for transfer. Finally, when a child is rewarded for achieving a challenging goal, the child may interpret the reward as “I was able to do something that was hard, and the teacher thinks I'm good at it.” The result can be an increase in intrinsic motivation, which leads to the potential for transfer. In short, rewards don't affect learning, but the learner's interpretation of rewards does. Similarly, getting a wrong answer should be interpreted not as a failure but rather as a challenge to try harder. When the instructional goal is to promote

transfer, special consideration must be given to helping the learner interpret rewards (and punishments) in ways that increase rather than decrease intrinsic motivation.

IMPLICATIONS OF RESEARCH ON THE HIDDEN COSTS OF REWARD

The instructional prescription "use positive reinforcement" is one that needs to be applied in light of research on the hidden costs of rewards. Indiscriminate use of rewards to reinforce all desired behaviors can backfire. Students should not be rewarded for doing things that they already enjoy doing. For example, if a student enjoys reading novels and spontaneously does so, it would not always be appropriate to enter into a contingency contract in which the student receives prizes for reading and reporting on a certain number of books. Similarly, in many cases where rewards are used, the ultimate goal should be to encourage the student to develop intrinsic motivation. Motivation is explored in Chapter 14.



How Does Feedback Affect Response Learning?

RESPONSE LEARNING IN THE CLASSROOM

Another major instructional application of Thorndike's work has been an emphasis on drill and practice with feedback. For example, consider the following dialogue:

TEACHER: How do you spell *behavior*?

STUDENT: B-e-h-a-v-e-r.

TEACHER: No, you have the first two syllables correct, but the last syllable is misspelled.

STUDENT: B-e-h-a-v-i-o-r.

TEACHER: Right.

In this case, the teacher presents a question (i.e., a stimulus) and the student spells a word (i.e., a response). If the response is correct, the teacher says "right," and if the response is wrong, the teacher says "no" and gives a hint.

BEHAVIORIST VERSUS COGNITIVE THEORIES OF RESPONSE LEARNING

Let's examine two contrasting interpretations of how feedback affects response learning. First, according to a reinforcement interpretation, feedback serves to reinforce (or strengthen) the association between the response and the stimulus. In the spelling example, the reinforcer is "right," and this serves to strengthen the student's tendency to emit the correct spelling for behavior. Feedback automatically stamps in the correct response without the learner's having to actively be aware that learning is taking place. Second, according to a cognitive interpretation, feedback serves as information to the learner. The learner can interpret this information and use it as a key in generating responses. In the spelling example, the feedback concerning the location of the error is helpful information that the learner can use in generating a better response. Feedback is actively interpreted by the learner so that its effect depends on how the learner thinks about it. In summary, the controversy concerning how feedback affects learning centers on


the question, "Does feedback serve as a reinforcer or does feedback serve as information that is interpreted by the learner?"

FIGURE 7-11
Motor Learning
with Feedback

RESEARCH ON RESPONSE LEARNING

Practice With and Without Feedback Early work by Thorndike (1931) clearly demonstrated that feedback improved response learning in humans. In one study, subjects were seated at a desk with a pencil and a large pad of paper. The subjects were asked to close their eyes and draw a line 4 inches long and to continue trying to draw 4-inch lines, keeping their eyes closed throughout. They were asked to perform this task day after day until 3,000 lines had been drawn. The results of this tedious study were clear: Repetition of the response 3,000 times caused no learning. The performance of the subjects at the end of the study was no better than the performance at the start of the study. However, in another experiment, Thorndike gave feedback after each line drawing: "right" if the line was within $\frac{1}{8}$ inch of the target length and "wrong" if it was not. Under these conditions, subjects' performance did improve. After approximately 4,000 trials, performance increased from 13% correct to 25% correct. Based on research studies like these, Thorndike concluded that practice alone would not promote learning but practice with feedback would. His argument was so compelling that it served to change school practices; instead of forcing students to practice without feedback, schools began to make use of praise and reward as techniques for increasing learning. Feedback was viewed as a way of reinforcing correct responses.

Quality of Feedback Subsequent research suggests that feedback may serve as information rather than reinforcement. If feedback serves mainly as a reinforcer, then the important aspect of feedback is to tell the learner whether the response was correct. If feedback serves mainly as information that is actively interpreted by the learner, then detailed feedback should be more effective than simple "right-wrong" feedback. An early study by Trowbridge and Cason (1932; reported in Adams, 1976) helped test these predictions. Students were blindfolded and asked to draw one hundred 3-inch lines. Some students were told "right" for each line within $\frac{1}{8}$ inch of the target and "wrong" for all other lines (right-wrong feedback group); some students were told how many eighths of an inch too long or too short each line was ("how much" feedback group); some students received no feedback (no-feedback group). As you can see in Figure 7-11, the no-feedback group showed no improvement, thus confirming Thorndike's observation that practice alone does not enhance learning. Moreover, the right-wrong group showed some improvement, ending up with an average error of about $\frac{1}{2}$ inch. This result is also consistent with Thorndike's research showing that right-wrong feedback aided learning. However, the main new result in Figure 7-11 is that the "how much" feedback group showed dramatic improvement from the start, ending up with an average error of less than $\frac{1}{8}$ inch. Thus, the group given detailed feedback learned more rapidly and more completely than the group given only right-wrong feedback. This study suggests that learners can use the information in feedback to help revise their plans for how to generate lines. It is interesting to note that the most useful information came on errors; in other words, subjects in the how much group seem to have been able to learn from their mistakes. This result conflicts with reinforcement theory's assertion that learning occurs mainly when a correct response is reinforced. In summary, these results are most consistent with the feedback-as-information view rather than the feedback-as-reinforcement view of response learning.




Source: Adapted from Trowbridge, M. H., & Cason, H. (1932). An experimental study of Thorndike's theory of learning. *Journal of General Psychology*, 7, 224-258. Reprinted with permission of the Helen Dwight Reid Educational Foundation. Published by Heldref Publications, 1319 Eighteenth St., NW, Washington, DC 20036-1802. Copyright © 1932.

Duration of Feedback Subsequent research also provides additional support for the feedback-as-information view. For example, let's consider what would happen if you provided feedback for part of the time and then stopped giving feedback. If feedback served mainly as reinforcement, then once the feedback is taken away we can expect extinction (i.e., performance will deteriorate). If feedback serves as information that the learner uses for building an internal plan or procedure, then once a correct plan is learned, feedback is no longer needed. To test these conflicting predictions, Newell (1974; described in Adams, 1976) gave subjects 75 tries at moving a lever $9\frac{1}{2}$ inches. Some subjects received feedback after each try (all-trials feedback group); some subjects received feedback after the first 52 tries but not the last 23 tries (52-trial feedback group); some subjects received feedback on the first 17 tries but not on the last 58 tries (17-trial feedback group); and some subjects received feedback after just the first 2 tries (2-trial feedback group). Figure 7-12 shows that the group given feedback on all 75 tries and the group given feedback on 52 tries both performed about the same. Apparently, after 52 tries the subjects had acquired an internal procedure for how to generate internal feedback concerning the appropriate response so that external feedback was no longer needed. However, when feedback was taken away early in the learning (e.g., after 2 or 17 tries), performance deteriorated, as would be expected under extinction. Apparently, subjects had not yet developed a way of generating internal feedback, so they were lost once external feedback was taken away. These results suggest that "what is learned" as a result of feedback is not always just the stamping in (reinforcement) of a response; rather, under some conditions, learners can build a cognitive plan that allows them to generate internal feedback.

FIGURE 7-12

Effects of withdrawal of feedback early versus late in motor skill learning



Source: Adapted from Newell, K. M. (1974). Knowledge of results and motor learning. *Journal of Motor Behavior*, 16, 235-244. Reprinted with permission of the Helen Dwight Reid Educational Foundation. Published by Heldref Publications, 1319 Eighteenth St., NW, Washington, DC 20036-1802. Copyright © 1932.

Finally, growing evidence indicates that practice with feedback causes qualitative differences in what is learned (Anderson, 1993; Singley & Anderson, 1989). For example, a child's computational skill goes through several stages rather than simply a series of strengthening associations. As responses become automatized through practice with feedback, the learner is able to use those automatized procedures in more complicated tasks. In summary, although feedback can be used to enhance response learning, the effective use of feedback must be based on an understanding of how the learner interprets the information in feedback.

IMPLICATIONS OF RESEARCH ON RESPONSE LEARNING

The research presented in this section highlights the crucial role that feedback plays in response learning. However, the research indicates that feedback does not simply stamp in (or reinforce) the response. Instead, students think about the feedback and use the feedback as information to help them interpret their learning. Thus, high-quality feedback is most useful (i.e., learners need to be told specifically about what they are doing that is correct or incorrect rather than just be told "right" or "wrong"). Furthermore, eventually, students may learn to give themselves feedback on a well-practiced task. Thus, teacher-provided

feedback should be withdrawn only when the learner has reached a level of automatic internal feedback, an event that may take many trials of practice but that should be the ultimate goal of instruction for response learning.



How Does Feedback Affect Skill Learning?

SKILL LEARNING IN THE CLASSROOM

Skill learning involves learning how to do something, so learning how to use a word processor is an example of skill learning. Suppose you are sitting at a computer that is running a word processor. On the screen you see a page of text, but you also have a printout of the page with some corrections that need to be made, as shown in Figure 7-13. Write a brief summary of what you would do for each correction, but write it as a set of rules that you could explain to someone else. Which of the following best corresponds to how you explained the first correction?

- _____ to carry out the first correction, press the up arrow key 17 times and press the back arrow key 65 times and then press the keys n-o-t-space-o-n-l-y-space-w-i-l-l-space-t-h-e-space-u-n-i-t-space-n-o-d-e-s-space-i-n-space-t-h-e-s-e-space-t-r-a-c-e-s
- _____ to carry out the first correction, move the cursor to the first space in the first line and type in "not only will the unit nodes in these traces"
- _____ when you want to insert text, move the cursor to where the text should be inserted and type in the inserted text

Which of the following best corresponds to how you explained the last correction?

- _____ to carry out the last correction, press the down arrow key one time, press the right arrow key 51 times, press the delete key 11 times, press the keys t-a-r-g-e-t
- _____ to carry out the last correction, move the cursor to select the word "illustrates" press delete, then type "target"
- _____ when you want to replace text, highlight the text to be replaced and type in the new text

If you focus on specific behaviors, then you probably checked the first answer (or even the second answer) for each of these questions, but if you focus on more general principles of word processing, you probably checked the third answer.

BEHAVIORIST VERSUS COGNITIVE THEORIES OF SKILL LEARNING

The role of feedback in skill learning can be analyzed from a behaviorist or cognitive point of view. According to the behaviorist view, learning involves adding specific responses to one's repertoire, and feedback is needed to help learners add appropriate responses. This view is consistent with the theory of specific transfer in which transfer occurs only when two tasks share many common behaviors. According to the cognitive view, learning involves making sense of the learning situation, so feedback is needed to help learners

FIGURE 7-13
A word processing
task

not only will the unit nodes in these traces
accrue strength with days of practice, but also
the element nodes will accrue strength. As will
be seen, this power function prediction
corresponds to the data about practice. A set of
experiments was conducted to test the prediction
about a power-law increase in ^{the} strength with
extensive practice. In one experiment subjects
studied subject-verb-object sentences of the form
(~~The~~ lawyer hated the doctor). After studying
these sentences they were transferred to a
~~furthermore, the thought prevents the study~~
sentence recognition paradigm in which they had to
discriminate these sentences from foil ~~by the mind-~~
target
sentences made of the same words as the ~~illustrates~~
sentence but in new combinations. There were 25 days of
tests and hence practice. Each day subjects were tested
on each sentence 12 times (in one group) or 24
times in the other group. There was no difference

Source: Reprinted with permission of the publisher from *The transfer of cognitive skill* by Mark K. Singley and John R. Anderson, p. 73, Cambridge, Mass.: Harvard University Press, Copyright © 1989 by the President and Fellows of Harvard College.

build broad rules or procedures. This view is consistent with the theory of specific transfer of general principles in which transfer occurs when a principle or strategy learned in one situation can be applied to a new situation.

RESEARCH ON SKILL LEARNING

Suppose we find a group of people who know how to type but who have never used a word processor. (It was easier to find such people in the 1980s than it would be today.) After explaining the commands for the word processor, we give our students 3 hours of


practice per day for 4 days on word processing tasks such as the one in Figure 7-13. Some students use the same word processor every day (Group AAAA); others use one kind of word processor on the first 2 days, but on the last 2 days they use a different kind of word processor that has different commands than the first one (Group BBAA).

When Singley and Anderson (1985, 1989) conducted this study, they essentially were looking for transfer. If skill learning involves acquiring specific behaviors, Group AAAA should outperform Group BBAA on days 3 and 4. The rationale is that quite different sequences of key presses are needed for the same tasks in the two different word processors. If skill learning involves building general strategies, both groups should perform about the same on days 3 and 4. The rationale is that both groups received practice in the same kinds of solution strategies, with only the specific commands being different.

Figure 7-14 shows the average time (in seconds) to carry out a correction for the two groups on each of the 4 days of learning. As you can see, Group AAAA showed a pattern of skill learning in which the time needed to carry out a correction decreased across the four days. Importantly, Group BBAA showed the same pattern of skill learning even though students had to switch word processors. In short, Group BBAA showed a high level of transfer from one word processor to another.

Why was there high positive transfer? According to Anderson (1993), both groups learned general strategies or rules rather than solely specific behaviors. Similar to the theory of specific transfer of general principles described in Chapter 1, transfer is

FIGURE 7-14
Average time to
make a correction
across 4 days for
two groups



Source: Adapted from Singley, M. K., & Anderson J. R. (1989). *The transfer of cognitive skill*. Cambridge, MA: Harvard University Press. Copyright 1989 Harvard University Press. Reprinted by permission.

facilitated when learners focus on general strategies rather than specific behaviors. What the word processors had in common was not specific keystrokes (such as which keystrokes to press) but rather general rules (such as how to insert text or how to replace text). Singley and Anderson (1989) conclude that "the very high level of positive transfer observed between [word processors] that shared few commands reinforces the position that superficial identical elements of the type that Thorndike advocated are inadequate" and "provides further support for an abstract representation of elements" (p. 112). According to this view, the key to transfer of cognitive skills rests in representing them as general rules rather than as specific behaviors.

IMPLICATIONS OF RESEARCH ON SKILL LEARNING

Building on earlier research on motor skill learning, Anderson and his colleagues (Anderson, 1993, 2005) offer a three-stage theory of cognitive skill learning.

Cognitive stage. The learner works from instructions and represents the steps in the skill as verbal statements—which can be called *declarative knowledge*. For example, in learning to drive using a manual transmission, you might tell yourself, "Second is directly below first."

Associative stage. The learner begins to be able to carry out a procedure more easily without needing to verbalize each step. The new skill is changed from declarative knowledge (i.e., factual statements) into *procedural knowledge* (i.e., a sort of mental program). For example, at this stage of learning to use a manual transmission, you might be able to shift gears without having to verbalize what you are doing.

Autonomous stage. The learner is able to carry out the skill rapidly and errorlessly without paying attention to it. The learner may even be unable to verbally describe the skill, such as not being able to tell someone how to tie his or her shoes.

What is learned that can be transferred from one skill to another? At first, in the cognitive stage, the learner acquires factual statements such as "The backward arrow key is in the lower right corner of keyboard." Next, in the associative stage, the unit becomes a collection of productions (Anderson, 1993), that is, IF-THEN pairs such as "If you want to move the cursor one space back, press the backward arrow key." Through practice, learners consolidate knowledge by collapsing several specific procedures into a single one. For example, the actions in erasing a letter may become one production: "If you want to erase a letter, move the cursor to the right of the letter and press the delete key." Finally, in the autonomous stage, the learner may create more generalizable productions by combining specific productions into a more general one. For example, the many specific procedures for replacing various texts can be generalized into a general production: "If you want to replace text, highlight the to-be-replaced text with the cursor and type in the new text." As you can see, skill learning involves building units of knowledge that are increasingly more transferable—beginning with specific verbal statements that change to specific procedures then change to more general procedures. Thus, the key to transfer of skill learning is to help students build general procedures. Anderson and his colleagues (Anderson, 2005) developed computer tutoring systems for skills ranging from writing geometry proofs to writing computer programs. In each case, teaching for transfer is accomplished by giving students practice in solving problems and providing feedback aimed at helping them build generalized solution procedures.



How Does Deliberate Practice Affect Complex Skill Learning?

DELIBERATE PRACTICE IN THE CLASSROOM AND BEYOND

What is the best way to help someone master a skill, such as becoming an expert pianist, an outstanding baseball player, or a top-notch writer? According to Ericsson (2003a, 2003b), the answer is what he calls *deliberate practice*—continually practicing a skill at increasingly more challenging levels of performance with the intention to master it. Ericsson (2003b) describes the process of deliberate practice as follows:

The acquisition of expert performance extends over years and even decades, but improvement of performance is not an automatic consequence of additional experience. Merely performing the same activities repeatedly on a regular daily schedule will not lead to further change once a physiological and cognitive adaptation to the current demands have been achieved. *The principal challenge for attaining expert performance is that further improvements require continuously increased challenges that raise performance beyond its current level.* The engagement in these selected activities designed to improve one's current performance is referred to as *deliberate practice*. Given that these practice activities are to be outside the aspiring experts' current performance, these activities create mistakes and failures in spite of the performers' full concentration and effort—at least when practice on a new training task is initiated. (p. 116)

In short, according to Ericsson's theory of deliberate practice, the development of expert performance in a domain depends on the quality and quantity of practice. Concerning quality of practice, learners need situations in which they obtain detailed feedback and encouragement, especially following failures on challenging tasks. Concerning quantity of practice, learners need to spend a great deal of time, with consensus estimates that it takes 10 years of sustained practice to become an expert in a field (Ericsson, 2001a; Simon, 1980).

BEHAVIORIST VERSUS COGNITIVE THEORIES OF DELIBERATE PRACTICE

Why would deliberate practice foster learning? If you are a proponent of the law of effect—psychology's first principle of learning—you could say that practice is an indispensable ingredient in learning. According to the behaviorist interpretation of the law of effect, practice works because it provides opportunities for the learner to produce responses and receive rewards or punishments. The rewarded responses are strengthened and the punished responses are weakened. Becoming an expert involves accumulating a very large repertoire of correct responses, and practice with feedback is the vehicle for achieving that accumulation. An implication is that students learn best when they are rewarded for giving a correct response, because in this way the association between the situation and the right response is strengthened. In short, it is best to give students simple questions that require a single response that they usually can answer correctly. A primary instructional method based on this view is *recitation*—in which a student answers a simple question and receives feedback on whether the answer was correct. For example, when you create flash cards to help you memorize the definitions of key terms, you are using a recitation method.

Recitation methods can be effective in learning specific facts or procedures, but what happens when the goal is to master a domain—such as to become an expert pianist, baseball player, or writer? This is where the cognitive interpretation of the law of effect comes in. According to the cognitive interpretation of the law of effect, learners are attempting to build cognitive representations—including knowledge of how things work (conceptual knowledge), knowledge of problem types (schematic knowledge), knowledge of how to plan problem solutions (strategic knowledge), basic facts (factual knowledge), and basic procedures (procedural knowledge) as described in Chapter 1—and they use feedback from the consequences of their performance to help them determine how to adjust their knowledge. According to this view, practice is most effective when it involves challenging and complex tasks on which the learner is likely to produce errors. This situation provides learners with rich feedback—and personal encouragement—that they can use as information to help build more effective cognitive representations. Ericsson's vision of deliberate practice is consistent with the cognitive interpretation of the law of effect. For example, when a pianist practices on increasingly more challenging pieces over the years, he or she is engaging in a form of deliberate practice.

If the cognitive interpretation of the law the effect is valid, then it may be worthwhile to shift our focus from recitation methods to deliberate practice. In the next sections, I focus on the effectiveness of deliberate practice as a way of learning to become an expert in a domain.

RESEARCH ON DELIBERATE PRACTICE

Does deliberate practice work? Let's look a simple example—trying to improve your digit span through practice. In a digit span task, I read you a list of digits—one per second—and then ask you to repeat them in order. Your digit span is the longest list you can repeat without making an error. For example, read the following list aloud, and then try to write down the digits in order (without looking at this page):

8 3 5 2 7 6 4 9 2 3


If you are like most people, you made mistakes in writing down the numbers. In fact, in a typical study the digit span of college students was about 7—that is, they could repeat a list of seven digits without error but they made errors on longer lists (Chase & Ericsson, 1982; Ericsson, Chase, & Faloon, 1980). Let's see what happens when the students practice on increasingly more complex strategies for increasing their digit span. For example, a simple beginning strategy is *chunking*, in which a list of 10 single digits is converted into three chunks such as 835-276-4923. After practicing 1 hour per day for 50 days, students improved their digit span to over 20 digits, and after 240 days, students were up to over 80 digits. Although you may not be excited about spending 240 days practicing how to memorize digit lists, this classic research demonstrates that deliberate practice can work. The features of practice in this successful training included learning to use increasingly more complex strategies for increasingly longer lists, making errors, and getting high-quality feedback.

Let's consider another domain—becoming an accomplished musician. Let's start with young children—beginning at about age 6—who show an interest in music, and see what is the difference between those who end up in a selective music school as teenagers (higher performing group) and those who do not (lower performing group). A team of

researchers in England (Davidson, Howe, Moore, & Sloboda, 1996; Howe and Davidson, 2003; Sloboda, Davidson, Howe, & Moore, 1996) interviewed 257 higher performing and lower performing music students to determine factors that may have led to their current level of musical expertise. The students had all been recognized for musical talent, had begun studying music at about age 6, and were teenagers at the time of the interview. First, the higher performing group reported more parental encouragement and support in their music lessons throughout elementary school. That is, "the most successful children had parents who were most involved in lessons and practice in the earliest stages of learning" (Davidson et al., 1996, p. 399). Similarly, students in the higher performing group were more likely to perceive their music teachers as being warm and encouraging.

Second, the higher performing group reported spending more time practicing (e.g., around 2 hours per day by age 12) compared to the lower performing group (e.g., 15 minutes per day). By age 12, the higher performing group averaged a total of 2,572 hours of practice, compared with 439 hours for the lower performing group. Figure 7-15 shows

FIGURE 7-15
How many minutes per day did higher and lower performing music students practice when they were starting to learn music?



Adapted from Sloboda et al. (1996).

the mean number of minutes of practice per day for the first 4 years of training (during elementary school) for students who had achieved higher and lower levels of musical accomplishment in high school. As you can see, the students who became high-performing musicians as teenagers engaged in progressively more practice in elementary school, whereas the students who did not achieve high levels of musical performance as teenagers engaged in far less practice when they were younger. In short, "there was a very strong relationship between musical achievement and the amount of formal practice undertaken" (Sloboda et al., 1996, p. 287).

The groups did not differ in terms of signs of early musical ability or exposure to music lessons. In short, the main differences between the higher and lower performing students was in the quality and quantity of their practice. Concerning quality, higher performing students received more encouragement from parents and teachers; and concerning quantity, higher performing students spent much more time practicing.

Other researchers have also found evidence for the role of deliberate practice in music learning. Ericsson (2001, 2003b) conducted case studies of expert musicians who had attained varying degrees of musical acclaim. The main finding was a "close relation between the amount of deliberate practice accumulated during musicians' development and the level of attained music performance" (Ericsson, 2003b, p. 114). Importantly, musicians who had reached elite status in the performance class at a musical conservatory had spent over 10,000 hours of solitary practice by the age of 20.

In a related study, music students were asked to keep a diary concerning their music practice (Sloboda, et al., 1996). There was a strong positive correlation between amount of deliberate practice and level of achievement on the British Associate Board examinations of musical performance. Howe, Davidson, and Sloboda (1998) report that "achieving the highest level (grade 8) of the . . . examinations in performing music required an average of approximately 3300 hours of practice irrespective of . . . ability group" (p. 450). It was the amount of practice rather than the supposed ability of the music students that most strongly determined their level of musical performance on the exam.

Research on deliberate practice shows that the law of effect can work its wonders only if students are committed to spending a lot of time practicing—in particular, practicing at increasingly more challenging levels. Why do some students spend so much more time practicing than other students who appear to have equivalent talent? According to a review by Mayer (2003b)—for students with basic talent—positive early experiences, adult encouragement, and noncognitive factors (such as temperament) can play an important role in the learner's willingness to engage in massive amounts of practice, which in turn contributes to the learner's knowledge and ultimately the learner's performance.

In spite of the encouraging evidence presented in this section, you may have noticed that all of the research is based on correlations between practice and performance. You might well ask yourself whether the correlation could be attributed to self-selection, in which students with high talent tend to practice more and to perform better. We tried to control for this issue by comparing students who began with seemingly equivalent levels of musical talent as young children but ended up displaying different levels of performance as high schoolers. Yet, a stringent way to examine the "practice-makes-perfect" hypothesis is to examine experimental studies that compare the performance of students who do or do not receive extensive practice with feedback. This is the approach I take in the next section.

IMPLICATIONS OF RESEARCH ON DELIBERATE PRACTICE

In reviewing research on skill learning, Anderson and Schunn (2000) came to the conclusion that "high degrees of competence only come through extensive practice" (p. 17). The implications for instruction are clear: "goals for education require practice and more practice" (p. 17). One way to provide high-quality practice with immediate and personalized feedback is through computer-based tutors (Anderson, Corbett, Koedinger, & Pelletier, 1995). Students are asked to solve increasingly more complex problems—such as writing computer programs or solving geometry problems—and when they make an error they receive immediate feedback explaining why it is an error. The computer has a model of the knowledge needed to solve the problem and the appropriate solution paths. When a student's solution moves off the appropriate path, the tutor provides short and focused instruction aimed at getting the student back on the solution path. In particular, the tutor identifies "bugs" in the learner's knowledge (as discussed in Chapter 5) and explicitly points out the needed changes. In this way, the "tutor provides immediate feedback, consisting of short and directed error messages" (Anderson et al., 1995, p. 167).

Does providing practice with personalized feedback help student learning? Consider a college course in computer programming (in a language called LISP) in which some students work on programming exercises in a standard LISP environment (control group) and other students work on programming problems with a LISP tutor. Both groups "learn by doing," but the tutored group receives high-quality feedback that is customized to explain each error. Anderson et al. (1995) reported that students who received the tutor learned in 30% less time and scored 43% better (or 1 standard deviation higher) on a final test than did students in the control group. Similarly, when a geometry tutor was incorporated into a high school math class, student test scores increased by 14 points on an 80-point test.

Consistent with a tenet of Thorndike's law of effect, Anderson et al. (1995) offer the strengthening principle: "It is assumed that both declarative and procedural knowledge acquire strength with practice" (p. 169). However, consistent with the cognitive interpretation rather than the behaviorist interpretation of the law of effect, practice works because it allows the learner to receive feedback that can be used to repair faulty knowledge. The effectiveness of Anderson et al.'s tutors is based on allowing learners to make errors within the context of solving authentic problem-solving tasks and providing them with immediate, focused, and concise feedback aimed at helping them repair errors in their knowledge. In short, "tutors provide significant learning gains" because "they help guarantee that the student is engaged in deliberate practice" (Anderson & Schunn, 2000, p. 24).

As another example, consider the development of expertise in troubleshooting—diagnosing and repairing malfunctions in mechanical and electronic devices. Lesgold (2001; Gott & Lesgold, 2000) provided troubleshooting training in aircraft repair using a computer-based tutor called Sherlock. Trainees worked on simulated repair problems and received immediate and focused feedback. Lesgold (2001) reports that 20 to 25 hours of Sherlock training produced the same amount of performance improvement as about 4 years of on-the-job training. Apprentices (with an average of 33 months of experience) who received extensive practice with a more sophisticated version of Sherlock showed a much greater improvement in their troubleshooting skill than did those who practiced without Sherlock—that is, the Sherlock group outscored the control group by more than 1 standard deviation on tests of troubleshooting. In fact, the Sherlock-trained apprentices

reached the same level of performance as those who had 10 years of experience (Gott & Lesgold, 2000; Gott, Lesgold, & Kane, 1997).

Overall, there is promising evidence that learning can be achieved through carefully sequenced practice on authentic tasks, with immediate and focused feedback that is aimed at building expert knowledge.



Chapter Summary

This chapter began with a look at Thorndike's famous law of effect, the idea that feedback following a response serves to determine the likelihood that the response will be made again. Then it examined situations that seem related to the law of effect: classroom management, classroom activities, response learning, and skill learning. In each case, we compared two contrasting views of the learner. The behaviorist view is that the learner is passively being conditioned by the feedback; that is, feedback serves as reinforcement that automatically stamps in (or stamps out) responses. This view is most consistent with the theory of specific transfer described in Chapter 1, because what is learned is a specific stimulus-response (S-R) association. In contrast, the cognitive view is that the learner is actively interpreting the feedback that is given; that is, feedback serves as information that the learner uses in building a learning outcome. This view is most consistent with the theory of specific transfer of general principles described in Chapter 1, because what is learned is a general attitude, rule, or procedure. As you can see, we have moved full circle around Thorndike's law of effect—beginning with a hungry cat's trying to escape from a puzzle box based on automatic strengthening and weakening of associations based on reinforcement, and ending with eager learners trying to master their domain through deliberate practice based on using feedback to build knowledge.

For classroom management, we found that token economies and contingency contracting can serve to increase desired behaviors and decrease undesired behaviors. However, the mechanism underlying the behavior change seems to be best described by a cognitive theory (i.e., by assuming that the learner attempts to interpret or justify the rewards or punishments that are given). An aspect of teaching for meaningful learning involves helping the learner achieve an interpretation that fosters intrinsic motivation rather than dependence on external rewards.

For classroom activities, we found that rewarding a behavior that students like to do can result in a decrease in that behavior—that is, a hidden cost of reward. This finding is most consistent with the cognitive interpretation of the law of effect, in which learners use rewards as information to help guide their construction of cognitive representations.

For response learning, practice with feedback tended to enable improvements in performance, whereas practice alone did not affect performance. Again, however, the mechanism underlying the behavior change seems best described by a cognitive theory (i.e., by assuming that the learner uses feedback as information). Teaching for meaningful learning involves helping the learner build a mental program for producing appropriate behaviors (including self-feedback) rather than learning specific behaviors themselves.

For skill learning, practice with feedback resulted in improved performance, and again, a cognitive theory (i.e., that learners build increasingly general rules) was the best

explanation of the behavior change. Teaching for meaningful learning involves helping the learner compile general procedures rather than a chain of specific behaviors.

Similarly, mastering complex skills depends on deliberate practice, in which learners receive immediate and high-quality feedback on their performance in increasingly more challenging tasks.

In summary, we are left with a sort of irony. The basic findings of Thorndike and Skinner and other behaviorist-oriented researchers seem to be fairly well established: The appropriate use of feedback (or reinforcers) given after the learner's responses can serve to change behavior. However, the processes of learning do not seem to fit the theoretical description given by the behaviorists. Instead, as we have seen in this chapter, the learner appears to be actively involved in trying to interpret the feedback that is given. Any effective use of feedback or reinforcement techniques requires an understanding of the active cognitive processing by the learner. In essence, we leave with a modified view of the law of effect: It is not the feedback that changes behavior, but rather the learner's interpretation and understanding of that feedback. When the goal of teaching is to promote transfer, teachers should make sure that learners interpret feedback in productive ways. Students engage in meaningful learning when they build intrinsic motivation that supports a self-monitored change in behavior, general rules that support concept learning, and general procedures that support skilled performance. In short, teaching for meaningful learning occurs when learners use feedback to make sense out of a learning episode rather than to acquire a collection of specific responses.

SUGGESTED READINGS

- Cameron, J., & Pierce, W. D. (2002). *Rewards and intrinsic motivation: Resolving the controversy*. Westport, CT: Bergin & Garvey. (Reviews research on the hidden cost of reward.)
- Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, 125, 625–668. (Reviews research on the hidden costs of rewards.)
- Singley, M. K., & Anderson, J. R. (1989). *The transfer of cognitive skill*. Cambridge, MA: Harvard University Press. (Describes the process of skill learning.)
- Thorndike, E. L. (1913). *Educational psychology*. New York: Columbia University Press. (Presents Thorndike's famous law of effect.)