Operant Conditioning

Module 17

What can you learn from a pigeon in a box? Plenty, as we shall see when we look at the work of B. F. Skinner and others who studied the principles of operant conditioning.

Learning Goals

- **17-1** Define operant conditioning.
- **17-2** Explain the law of effect.
- **17-3** Describe the different kinds of reinforcement.
- **17-4** Explain how punishment influences behavior.
- **17-5** Describe how shaping, discrimination, and extinction apply to operant conditioning.
- **17-6** Identify the advantages and disadvantages of different schedules of reinforcement.
- **17-7** Explain the role of cognition and biological predispositions in operant conditioning.

The Nature of Operant Conditioning

17-1 What is operant conditioning?

A few years ago, I was in the grocery store on a busy preholiday afternoon. In front of me in the checkout line were a mother and her preschool-age daughter. When the girl spotted the candy bar display, she asked if she could have one. Mom, clearly tired and running out of patience on a hectic day, said, "No." The girl, a bit cranky herself, started to fuss, and the battle was on. It quickly escalated to a full-fledged tantrum on the girl's part and yelling on the mom's. A short while later, the mother apparently decided the girl was not about to settle down while they waited to pay for their groceries. With an exasperated "Here, just take it and be quiet," she tossed a candy bar toward the girl, who instantly went from tantrum mode to a bright smile as she tore into the treat.

You may be thinking, "Aha! That little girl certainly had her mother trained." In fact, both the mother's and the daughter's behaviors were affected by **operant conditioning,** a type of learning in which the frequency of a behavior depends on the consequence that follows that behavior. In operant conditioning, how often a behavior occurs (its frequency) depends on the event that follows the behavior (its consequence). By giving her daughter the candy bar, the mother was inadvertently using operant conditioning to teach her daughter to throw tantrums. At the same time, by stopping her tantrum, the daughter was unknowingly using operant conditioning to teach her mother to buy her candy (see **Figure 17.1**). You decide who was the more effective psychologist!

Operant conditioning is straightforward in its most basic form. Suppose your parents handed you a $100 bill each evening if you cleared your own dinner dishes

operant conditioning
A type of learning in which the frequency of a behavior depends on the consequence that follows that behavior.

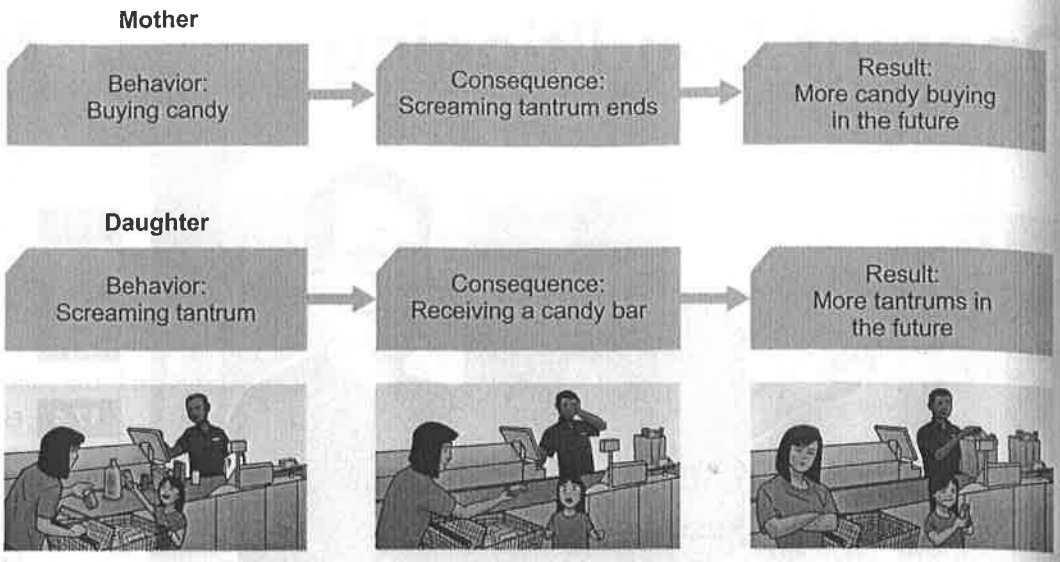

FIGURE 17.1 Operant Conditioning for Better or for Worse This child doesn't know it, but she is using operant conditioning to train her mother to buy candy. Her mother is unwittingly operantly conditioning the child to have tantrums.

from the table after completing your meal. Would this consequence influence your dish-clearing behavior? Most likely, you would begin to clear your dishes more regularly and more quickly than before. What if your parents somehow wired the television so that you received a painful electric shock every time you touched the power button on the set or the remote control? The shock is a consequence, too. Would it influence how often you touched the power button?

Wouldn't it be great if we could influence others—parents, siblings, friends, teachers, employers—to behave the way we want them to? The pair in the grocery store stumbled upon a powerful behavior-changing tool. Let's take a closer look at how operant conditioning techniques work.

MAKE IT STICK!

1. In operant conditioning, the frequency of a behavior is determined by the behavior's _____.
2. Which of the following is an example of operant conditioning?
   a. A boy who watches his brother ride a bicycle learns how to ride a bicycle.
   b. A girl practices spelling Mississippi until she learns to spell it perfectly.
   c. A baby learns to startle every time the doorbell rings.
   d. A dog that is given a piece of food after begging at the dinner table learns to beg more frequently.


Sarin Images/The Granger Collection, New York

EDWARD THORNDIKE (1874–1949) Author of the law of effect, the principle that forms the basis of operant conditioning.

The Law of Effect

17-2 What is the law of effect?

Operant conditioning developed from research conducted by **Edward Thorndike,** an early U.S. psychologist. Thorndike's work led him to describe the *law of effect,* which simply states that behaviors with favorable consequences will occur more frequently and behaviors followed by unfavorable consequences will occur

less frequently. Another U.S. psychologist, **B. F. Skinner,** built his life's work on this idea, developing the fundamental principles and techniques of operant conditioning (see Thinking Like a Psychological Scientist: B. F. Skinner). Two of the more important concepts used in operant conditioning are reinforcement and punishment.

- **Reinforcement** is any consequence that *increases* the future likelihood of a behavior.
- **Punishment** is any consequence that *decreases* the future likelihood of a behavior.

reinforcement Any consequence that increases the future likelihood of a behavior.

punishment Any consequence that decreases the future likelihood of a behavior.

Keep in mind that the learner, not the teacher, casts the vote that determines whether a consequence will be a reinforcement or a punishment. When my two children were younger, they had different feelings about broccoli. Carl really liked broccoli, and if I had given him some after he took out the garbage, he would have been more likely to take out the garbage in the future. For him, broccoli was a *reinforcement.* But Eric did *not* like broccoli. If I had given Eric broccoli after he took out the garbage, he probably would never have done that chore again. Broccoli was a *punishment* for Eric. My feelings about broccoli (which I like, especially with cheese sauce) make no difference.

Parents and other authority figures don't always understand this aspect of consequences. A parent will sometimes yell at a child for misbehavior, thinking that yelling is a form of punishment. However, a child who is usually ignored may actually crave being yelled at because it's about the only parental attention the child gets. In this case, the law of effect predicts the behavior that preceded the parent's yelling—misbehavior that the parent was trying to prevent—will actually be *more* likely to happen again. Likewise, a school administrator might suspend a student for skipping class, not realizing that for a student who doesn't care for school, this punishment is not a punishment at all: It may even make future skipping more likely!

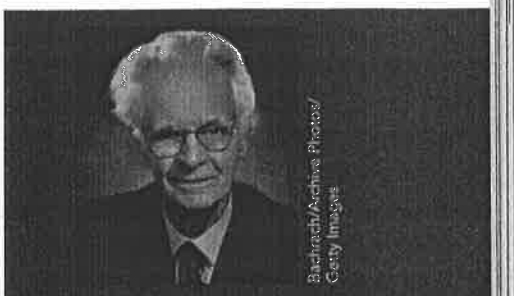

Bachrach/Archive Photos/Getty Images

B. F. SKINNER (1904–1990)
Behavioral psychologist who developed the fundamental principles and techniques of operant conditioning and devised ways to apply them in the real world.

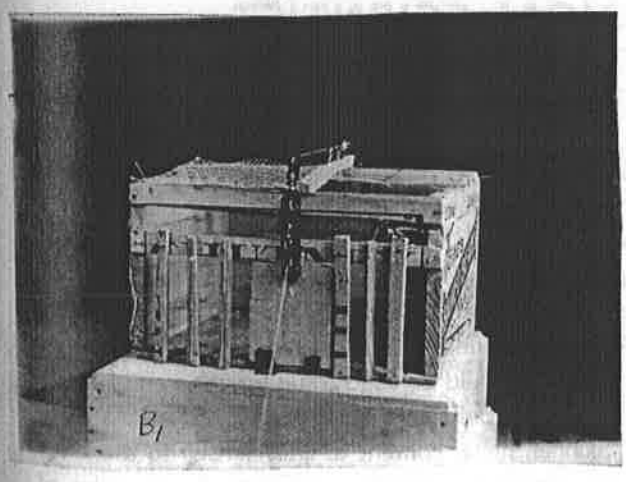

Macmillan Learning

Thorndike's Puzzle Box
Thorndike constructed puzzle boxes such as this one to study how cats learned to escape. At first, a cat's ability to find the lever that had to be pressed to open the door depended on trial and error, but later the animal had a tendency to repeat the behaviors that had been successful before. This research led to the development of the law of effect.

Charlotte van Oyen Witvliet

Reinforcement and the Law of Effect
Fishing crews in the Cayman Islands have cleaned their catch in this bay for years, thus providing an easy meal for the stingrays that live in the surrounding waters. The rays began to hang around, and then divers and swimmers began feeding them by hand and even petting them. Can you see how the law of effect has been at work in this situation?

How Do You Get What You Want? This little child enjoys sleeping with Mom and Dad. What behaviors on the child's part might have led to this reinforcement? Is it possible Mom and Dad are reinforcing behaviors they really don't want the child to engage in?

skynesher/E+/Getty Images

MAKE IT STICK!

1. _____ is any consequence that decreases the future likelihood of a behavior.
2. _____ is any consequence that increases the future likelihood of a behavior.
3. Your dog, Pavlov, loves your attention. Lately, Pavlov has been barking a lot, and every time he does, you yell, "Pavlov, be quiet!" Which of the following is true?
   a. Pavlov will bark less in the future because you reinforced him.
   b. Pavlov will bark less in the future because you punished him.
   c. Pavlov will bark more in the future because you reinforced him.
   d. Pavlov will bark more in the future because you punished him.

THINKING LIKE A PSYCHOLOGICAL SCIENTIST

B. F. Skinner

Few have done more than B. F. Skinner to advance the notion of psychology as a scientific discipline. In one study, almost 2000 psychologists ranked Skinner as the most eminent psychologist of the twentieth century.[1] Through his research and his writing, Skinner spent his career developing a behavioral technology that did not rely on references to unseen thought processes. His goal was to understand and control the actions of other organisms. His new behavioral technology was operant conditioning.

Skinner believed that all behaviors in all species are governed by the same principles. To identify these principles, Skinner studied simple behaviors, mostly of rats and pigeons. He and his assistants taught rats to press a lever with their paws and pigeons to peck a disk with their beaks. The rats and pigeons performed these acts in an invention that Skinner called an *operant chamber*, but that most others call a *Skinner box*. The Skinner box gives the experimenter an opportunity to control the environment and precisely record an animal's responses. Using this arrangement, Skinner identified principles of operant conditioning that he felt could be used to understand and control complicated behaviors in the real world.

Skinner, never content to limit himself to the laboratory, became a ceaseless advocate for his point of view. He was often the center of controversy because of many books and articles written by and about him. Skinner's public fame started when in a magazine article he detailed how he and his wife were using a climate-controlled "air crib" for their daughter (which some felt was too similar to a rat's or pigeon's operant chamber).[2]

(*Continued*)

THINKING LIKE A PSYCHOLOGICAL SCIENTIST (Continued)

Everett Collection Inc/Alamy

B. F. Skinner's Influence in the Laboratory
This pigeon in a Skinner box, or operant chamber, has been taught by Skinner to peck a key for food reinforcement.

Bernard Hoffman/The LIFE Picture Collection/Getty Images

The "Baby in a Box"
B. F. Skinner's daughter spent time in this specially designed environment. Some thought this was awful; others thought it was less of a problem than a standard baby crib with bars.

Skinner stayed in the public eye for the rest of his career. In the late 1940s, he published *Walden Two*, a novel presenting his ideas for a perfect community based on principles of operant conditioning. In the book, Skinner criticized democracy, the nuclear family, the use of money, and religion. Nevertheless, he inspired a number of groups to start communities that avoided punishment and used reinforcement to encourage desirable behaviors. Some are still in existence, and I once had a class write Skinner to learn what he thought of them. He replied that his principles were applied most appropriately at Los Horcones in Sonora, Mexico.[3] In 1971, Skinner published *Beyond Freedom and Dignity*, describing his belief that human freedom was an illusion. This sparked enough public discussion to land him on the covers of national news magazines.

Skinner loved the debate and was passionate about defending his positions and the science on which they were based. His final speech, delivered at the American Psychological Association convention only 8 days before he died of leukemia in 1990, was a spirited defense of behaviorism and critique of cognitive science, which he believed wasn't scientific at all.[4]

Skinner's viewpoints were extreme, but his work does have a number of successful, practical applications. As you will see throughout this module, operant conditioning principles affect our behavior at home, in school, in the world of sports, and at work.[5] Skinner may not have always won people's hearts and minds, but he did contribute importantly to our understanding of learned behavior and the development of psychology as a science.

THINK ABOUT . . . Psychological Science

1. How did Skinner identify the principles of operant conditioning?
2. Why was Skinner's air crib controversial?
3. Explain why you would or wouldn't consider living in a community built on the ideas Skinner explained in *Walden Two*.

Reinforcement

17-3 What are the different kinds of reinforcement?

Reinforcement procedures strengthen responses by making them more likely to occur again. There are two basic ways to reinforce a behavior: positive reinforcement and negative reinforcement (see **Figure 17.2**).

positive reinforcement In operant conditioning, anything that increases the likelihood of a behavior by following it with a desirable event or state.

negative reinforcement In operant conditioning, anything that increases the likelihood of a behavior by following it with the removal of an undesirable event or state.

Positive reinforcement is anything that increases the likelihood of a behavior by following it with a desirable event or state. For example, if a student earns an A in psychology and his mother pays him $10 for that, then she has positively reinforced his behavior.

Negative reinforcement is anything that increases the likelihood of a behavior by following it with the removal of an undesirable event or state. Negative reinforcement is a tricky concept. Note that the behavior is a means of either escaping or avoiding an undesirable situation. Notice that the words *positive* and *negative* in this context do not mean "good" reinforcement and "bad" reinforcement, but rather are used in a mathematical sense. Here, *positive* simply means that something desirable is *added,* and *negative* means that something undesirable is *subtracted*—in fact, it may be helpful to think of the math symbols + and – when you see the words positive and negative. For example, Roshni's headache is undesirable. If taking ibuprofen provides relief from the headache, Roshni's behavior of taking ibuprofen has been negatively reinforced. Negative reinforcement, like all reinforcement, *strengthens* a behavior. So, Roshni becomes more likely to take ibuprofen to escape a headache in the future.

The concept of negative reinforcement can be confusing at first. Sometimes students think it is simply a more technical term for punishment, but this is not the case. Punishment, as you will see soon, weakens a behavior. Negative reinforcement always strengthens a behavior that removes an undesirable stimulus. Here are two more examples of negative reinforcement:

1. The girl in the grocery store at the beginning of this module used negative reinforcement to teach her mom to buy her candy. The mother found the tantrum undesirable, and she escaped it by buying her daughter the candy she wanted. The end of the tantrum negatively reinforced the mother's candy-buying behavior, which is likely to happen again in the future.

FIGURE 17.2 Reinforcement Strengthens Behavior Positive and negative reinforcement work differently, but both types of reinforcement make a behavior more likely to happen again.

2. Hitting the snooze button on my annoying alarm clock is negatively reinforcing because the behavior allows me to escape from the alarm. This strengthens the behavior and helps ensure that I will hit the snooze button in the future.

Immediate Versus Delayed Reinforcement

Which affects our learning more—immediate rewards or delayed rewards? If psychologists had designed the warning label on cigarette packs, it might say, Warning: If you smoke these cigarettes, your breath will smell awful for the rest of the day! Psychologists know that we are more likely to respond to immediate consequences (bad breath) than to delayed consequences (the long-term risk of lung or heart disease). In other words, *immediate reinforcement* is more effective than *delayed reinforcement*. This also helps explain why it is difficult to *quit* smoking cigarettes. The desirable consequence—the rush produced by the chemicals in tobacco—is immediate. The undesirable effects on the lungs and cardiovascular system are longer term. It's no surprise, then, that in addition to the physiological effects that drugs have on the brain, those that produce the most immediate reinforcement, like nicotine and cocaine, are the most addictive.[6] You can see the same relationship in those who overeat. The taste of fattening foods provides immediate positive reinforcement, but the effects of obesity are delayed.

Rats and pigeons, like people, prefer immediate reinforcement, and they seem to require it for learning. A rat, for example, will not learn to press a bar if the reinforcement (usually food) for that behavior is delayed by 30 seconds or more. Humans, however, have a great ability to adapt, and one of the things we learn as we develop is that delayed reinforcers are sometimes worth the wait. Paychecks may not be issued until the end of the month and grades aren't given until the end of the grading period, yet they still influence us. In fact, the ability to delay gratification is a real advantage. For example, children who prefer a big reward in the future over a smaller reward now—perhaps by saving their allowance for a desired toy rather than spending it each week on candy—are likely to become higher-achieving adolescents than children who prefer immediate gratification.[7] It is possible that the instant gratification provided by social media and other screen-based activities is making it difficult for people to develop an appreciation for the long-term rewards provided by more time consuming pursuits like nurturing face-to-face relationships and reading books.

Primary Versus Secondary Reinforcement

Primary and secondary reinforcement are similar in that they both affect the frequency of behaviors, but they differ in one important way. A **primary reinforcement** is something that is naturally rewarding, such as food (if you are hungry), warmth (if you are cold), and water (if you are thirsty). A **secondary reinforcement** is something you have *learned* is rewarding because it has been paired with a primary reinforcer (see **Figure 17.3**).

primary reinforcement
Something that is naturally reinforcing, such as food (if you were hungry), warmth (if you were cold), and water (if you were thirsty).

secondary reinforcement
Something that you have learned to value, like money.

Money, for example, is a secondary reinforcer because you have learned you can use it to purchase various forms of primary reinforcement, such as pizza and clothes. But money itself is not naturally rewarding. If you give a $100 bill to a 4-month-old baby, she will probably put it in her mouth to see if it tastes good. (She's checking for a primary reinforcer!) When the baby discovers the money doesn't taste good, she will spit it out and show no further interest. Is this how you respond when given a $100 bill? Not likely! For you, that piece of paper has

FIGURE 17.3 Primary and Secondary Reinforcements Primary reinforcers, such as food, are naturally rewarding. Secondary reinforcers are rewarding because we have learned that they are associated with primary rewards.

PRIMARY REINFORCEMENT

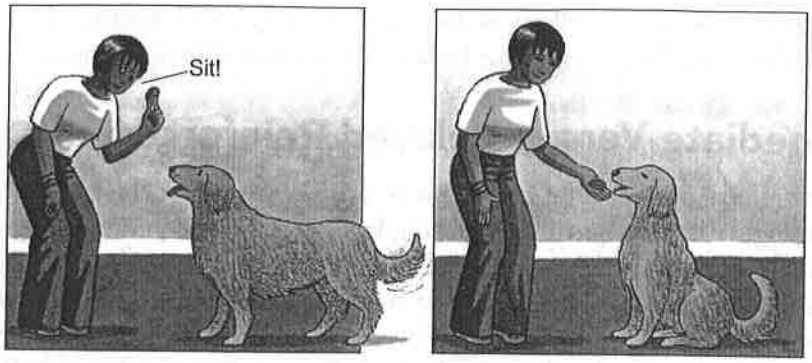

Food is a primary reinforcer for a dog.

SECONDARY REINFORCEMENT

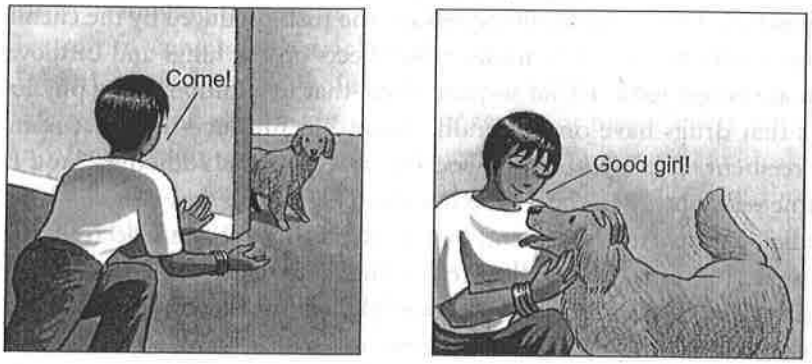

An owner's words can become secondary reinforcement when they're associated with petting and approval.

value—so much value that you would work hard to get more. And if you really are hungry, the money becomes even more valuable to you—in one experiment, participants were less likely to donate to charities if asked when they were hungry.[8]

Grades are a major influence on student behavior. Are they a primary or a secondary reinforcement? If you said secondary, you're right. You had to learn the value of grades. Without this learning, grades have no value. (Try training your dog by giving a B plus for sitting or heeling, instead of a primary reinforcement in the form of attention or a dog biscuit.) Some students never learn to associate much value with school grades, so grades have little effect on their behavior.

MAKE IT STICK!

1. Explain the distinction between positive and negative reinforcement.
2. Which of the following is an example of negative reinforcement?
   a. A husband gets a kiss on the cheek from his wife after he takes out the garbage.
   b. A third grader reads five books for a gold star award at school.
   c. A football player goes to the weight room every day to gain strength.
   d. A twisted ankle hurts less after taking Tylenol.
3. True or false? Both positive and negative reinforcement make a behavior more likely.

Punishment

17-4 How does punishment influence behavior?

The Process of Punishment

We noted earlier in this module that punishment *weakens* a behavior, or makes it less likely to occur again in the future. Punishment can take either of two forms (see **Figure 17.4**):

- The first, positive punishment, occurs when the behavior leads to something undesirable. Positive punishment is anything that decreases the likelihood of a behavior by following it with an undesirable event or state. For example, if a toddler puts her hand on a painfully hot stove burner, the behavior of touching the burner is punished because it leads to an undesirable event: getting burned (Ouch!). Because stove touching has been punished, that behavior is less likely to happen in the future.
- The second form of punishment, negative punishment, occurs when the behavior ends something desirable. Negative punishment is anything that decreases the likelihood of a behavior by following it with the removal of a desirable event or state. Let's say a young boy pulls his sister's hair while watching television and his father takes away TV privileges for the rest of the day. The behavior of hair pulling has ended something desirable—watching television. The loss of privileges should make the boy's hair-pulling behavior less likely to occur in the future. Parking tickets are another example of punishment that removes something desirable. If I engage in the behavior of illegal parking and then have to pay some of my desirable money, I should be less likely to park illegally in the future.

This is exactly parallel to what we learned about positive and negative reinforcement. *Positive* and *negative* are again being used in the mathematical sense and do not imply that one type of punishment is good while the other one is bad. Positive punishment simply means that something undesirable is *added*. Negative punishment means that something desirable is *subtracted*.

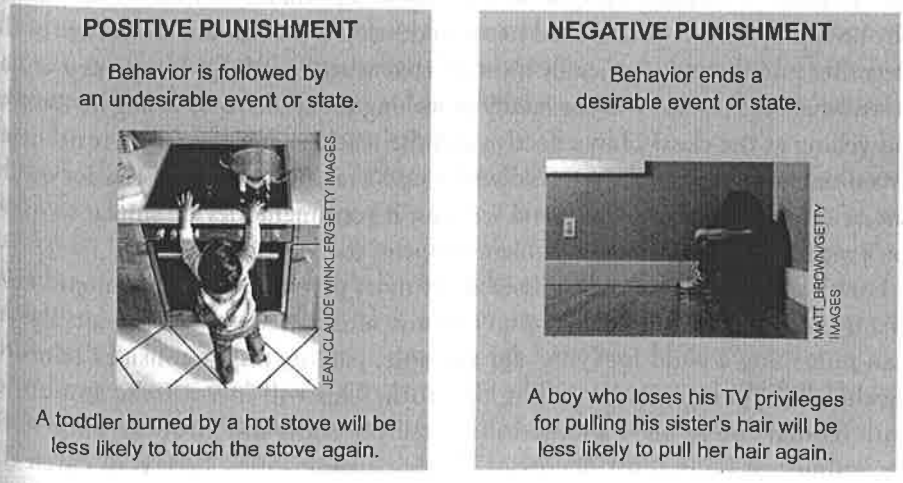

FIGURE 17.4 Punishment Weakens Behavior Here are two forms of punishment, positive and negative. In one, the punished behavior is followed by an undesirable event, which is *added*. In the other, the punished behavior is followed by the loss of a desirable event or state, which is *subtracted*. Although the types of punishment differ, each decreases the likelihood that the behavior will happen again.

Problems With Punishment

Many learning experts oppose the use of punishment, especially physical punishment, to control behavior.[9,10] They feel that punishment is likely to backfire in the long run for a variety of reasons. For starters, punishment does not end the desire to engage in a behavior. Children punished for using inappropriate language often continue to use the bad language—just not in the presence of the one who punished them for it. Likewise, adults punished for speeding may simply purchase a radar detector rather than drive moderately.

wavebreakmedia/Shutterstock.com

Pros and Cons of Punishment
Punishment, under the right circumstances, can decrease behavior, but it also has several undesirable side effects. It can produce fear and anxiety.

Punishment can also lead to fear and anxiety. Children or animals punished frequently may learn to engage in *avoidance* behaviors: Harshly punished children may run away from home, and harshly punished students may drop out of school. A final criticism of punishment is that when adult role models use aggression to solve their problems, children can learn to model that aggressive behavior as a problem-solving strategy. This may help explain why abusive parents tend to come from abusive families (although, impressively, most abused children do not go on to become abusive parents).[11]

So, despite all these problems, why is punishment used so often? One reason seems to be that when a punished individual stops misbehaving, even for just a few minutes, this consequence negatively reinforces the behavior it followed—which in this case is punishment. This consequence in turn makes the punishing behavior more likely to happen in the future. The result is a vicious cycle: Punishment leads to *temporary* suppression of misbehavior, which reinforces the punishment, which is then even more likely to be used when the suppressed misbehavior inevitably returns, which leads to another reinforcing, temporary suppression, and so on. For example, a teacher scolds a child who likes to talk in her first-grade classroom, and then the child quiets down for a few minutes—this negatively reinforces the teacher for scolding. However, the effects of the scolding do not last long, and the child begins talking again before you know it. The teacher, having been previously reinforced for scolding, will likely scold again. Thus, it goes round and round like a dog chasing its tail!

There is a role for punishment in learning, but it is a limited one. Swift, sure punishment can effectively control certain behaviors, especially if the punisher's goal is to protect a child from a dangerous situation.[12] For example, if a toddler has developed the bad habit of running into the street, a harsh reprimand or swat on the behind may be appropriate. A young child needs to develop some fear and avoidance of the street. But punishment is generally most effective when used least. Have you ever had a class where the teacher was constantly punishing the students by losing her temper and yelling at the class? How effective was the teacher's behavior after it occurred several times? Compare that scenario with a class's reaction to angry behavior on the part of a teacher who rarely yells and screams. If your memories are similar to mine, you'll see that the less often punishment happens, the more effective it is.

For all the reasons we've discussed here, most psychologists recommend *reinforcing an incompatible behavior* as an effective alternative to punishment. Rather than punishing a child for lying, for example, parents might consider reinforcing the child with praise for telling the truth. This will increase the amount of truth telling, and because a child cannot tell the truth and lie at the same time, the amount of lying must decrease. The basic philosophy here is to catch the child being good and reinforce accordingly. This approach will lead to a more

gradual change of behavior, but the change will be more permanent than the temporary suppression of behavior that follows punishment. Reinforcement also has two other benefits. First, it tends to lead to *approach* behaviors that draw people together (one reason kids are often so eager to see their mostly reinforcing grandparents) rather than feelings of fear and anxiety. Second, children who model positive reinforcement are more pleasant to be with than children who model aggressive behaviors.

MAKE IT STICK!

1. Positive punishment _____ the frequency of behavior. Negative punishment _____ the frequency of behavior.
2. What are two reasons psychologists usually oppose the use of spanking?
3. True or false? One of the advantages of punishment is that it usually permanently changes behavior.

Reinforcement Procedures

17-5 How can you use operant conditioning to teach a new behavior or make an operantly conditioned behavior stop?

shaping Reinforcement of behaviors that are increasingly similar to the desired one; the operant technique used to establish new behaviors.

Now that you know a little bit about how reinforcement and punishment work, we can begin to explore some other procedures that make operant conditioning so useful. In this section, we show how you can use shaping to establish new behaviors and how discrimination and generalization can fine-tune when behaviors will occur.

Shaping

The law of effect says that reinforcing a behavior makes the behavior more likely to occur in the future. But how can you apply operant conditioning to a behavior that hasn't yet occurred? To do this, Skinner developed a technique called **shaping,** reinforcement of behaviors that are increasingly similar to the one you want to occur. Shaping is the operant conditioning technique used to establish new behaviors. When you shape a behavior, you positively reinforce behaviors that move ever closer to the target behavior.

B. F. Skinner, for example, demonstrated how to train a pigeon to turn in clockwise circles. He started by providing a food reward every time the pigeon turned its head to the right. Pigeons turn their heads frequently, and the law of effect says that reinforcing this behavior will cause the pigeon to turn its head to the right more often. Now the trick was to gradually extend how far the pigeon must turn its head before Skinner gave it a food reward. After a series of increasingly longer turns, the pigeon finally turned all the way around. By breaking the circling behavior into a succession of gradual steps, Skinner easily shaped it.

Many other examples of shaping happen in everyday life. Remember when you learned to ride your bicycle without training wheels? Chances are, someone held

Shaping
When this child falls, as she surely will, Dad will praise her attempt. Parents who give such praise know that it is important to provide reinforcement following first attempts, even if they are failures, to encourage a child to ride ever farther.

Marc Romanelli/Getty Images

discrimination The ability to distinguish between two similar signals or stimuli and produce different responses.

extinction In operant conditioning, the loss of a behavior when no consequence follows it.

the seat, ran beside you until you were reasonably well balanced, and then let g You probably managed to roll several feet on your own before falling, at whic point you were rewarded with a hearty, Good job! It really wasn't that good job—you only made it a few feet—but for a first attempt it deserved reinforc ment. Gradually, as your riding skills improved, your trainer made you ride fa ther and farther before giving you a compliment. You were being shaped!

Discrimination and Extinction

Shaping is useful for training behaviors that otherwise probably wouldn't happen Other issues that make operant conditioning useful include *discrimination* and *extinc tion*. **Discrimination** is the ability to distinguish among similar signals or stimuli **Extinction** is the loss of a learned response when a consequence no longer follows it

Life would be fairly chaotic if we made the same response to all stimuli tha were similar. For safety reasons, students and teachers need to learn to *discriminat* between class bells and fire alarms—two kinds of fairly similar stimuli. We do thi by learning the difference between the similar signals. What signal tells you it's time to leave the classroom at the end of the period? How does it differ from the fire alarm signal? And false-alarm fire drills—where students are sometimes called back to class before they've even left the building—need to be kept to a minimum to prevent the extinction of the evacuation response.

Discrimination This trainer is teaching the dolphin to discriminate between different hand signals. Reinforcement—a fish treat—is provided only when the animal performs the proper behavior for that signal.

DOLPHIN INST./Science Source

Sometimes, however, *extinction* is a good thing. If the mom in the checkout line at the beginning of this module had managed to ignore her child's tantrums, those behaviors would eventually have died out on their own. Remember that without reinforcement, behaviors learned through operant conditioning will eventually disappear.

Without discrimination, we wouldn't know when to answer the phone and when to answer the door. We wouldn't know whether to say, Hi, Jill or Hi, Jane when a friend comes into view. Without extinction, we wouldn't stop repeating the same unsuccessful chess strategy or stop flirting with someone who doesn't respond to our interest. These operant conditioning concepts can help us understand why certain behaviors thrive and others die out. Your reading behavior has obviously not extinguished, so let's continue!

Cartoon Stock

"Will whoever has the doorbell ringtone please set your phone to vibrate?"

Signals Can Be Powerful Can you explain why the speaker wants the dogs in the audience to avoid the doorbell ringtone?

Learning New Habits

If you could learn or develop one new habit, what would it be? Eating healthier? Exercising more? Keeping track of your stuff better? Positive psychologists are learning how to best change a behavior you'd like to improve. More specifically, they're learning how to grow the skills you need to create the habits you want in life.

Not surprisingly, learning a new habit involves self-control (often called *self-regulation* by psychologists). Self-control is the process of exerting control over your thoughts, feelings, and actions.[13] The less self-control you have, the harder it is to learn and develop a new habit. So, if self-control is not one of your strengths, does this mean you're out of luck when it comes to learning a new habit? *No*. It turns out that there is a fairly simple way to increase your self-control.

Joseph Forgas, Roy Baumeister, and Dianne Tice (2009) have learned that one of the better things you can do to gain the self-control you need to create a new

habit is to practice self-control over *any* area of your life.[14] They explain what this means by using the results from two studies. In the first, participants were asked to exert self-control over spending money for several months. In the second, participants were taught an exercise routine and instructed to follow the routine every day for 2 months. The researchers found that while saving money increased in the first study and participants lost weight in the second, there were also unanticipated side benefits: The self-control used on finances and exercise showed up in *other* aspects of the participants' lives. Not only had they saved money or improved their health, but they also reported studying more, watching less television, and doing more household chores that they usually avoided. The narrowly focused self-control of the participants was like a muscle that was made stronger through regular use. It made it easier for these individuals to exert self-control more widely in other aspects of life.

So, if you want to build up the self-control skills you need to change bad habits or to learn new good habits, start by trying to control one thing that doesn't turn your world upside down. That is, start with something that only requires a little self-control. Tell yourself you're only going to check Facebook twice a day, and for no longer than 30 minutes ever. As psychologist Senia Maymin puts it, "Structure something concrete into your life. That's the best way to develop self-control."[15]

MAKE IT STICK!

1. Describe how you'd use shaping to teach a young child how to play catch with a Nerf ball.
2. We know to stop at a red light and go when the light turns green because of _____.
3. Extinction occurs when a behavior
   a. is punished.
   b. indicates discrimination has been learned.
   c. is shaped.
   d. is no longer followed by reinforcement.

Schedules of Reinforcement

17-6 What are the advantages and disadvantages of different schedules of reinforcement?

What do buying food from a vending machine and playing the lottery have in common? They are both behaviors maintained by positive reinforcement. However, you're more likely to continue playing the lottery after having purchased a losing ticket than you are to put more money into a vending machine that has just failed to produce the desired potato chips. This is because the two examples illustrate different schedules of reinforcement—continuous reinforcement for the vending machine and partial reinforcement for the lottery.

Continuous Reinforcement

In **continuous reinforcement,** a reward follows every correct response, just as a vending machine is supposed to operate. The vending machine trains you to behave in a certain way—inserting your money in the machine—by continuously reinforcing your behavior. If you put the money in properly, you will be reinforced every time by receiving a bag of potato chips.

continuous reinforcement
In operant conditioning, a schedule of reinforcement in which a reward follows every correct response.

Continuous reinforcement is most useful for establishing new behaviors. Lots of reinforcement is often necessary when you are trying to teach someone to do something new, such as to speak a new language. One problem with behaviors that have been continuously reinforced, however, is that they are quite easy to extinguish. If the learner is used to being reinforced for each correct behavior and the reinforcement stops, extinction will occur rapidly. Think about how you behave when you put money in a vending machine and the machine doesn't dispense your product. Do you quickly put more money in? Probably not! When the goal is to establish behavior that is resistant to extinction, one of the partial reinforcement schedules works better.

Partial Reinforcement

partial reinforcement schedule In operant conditioning, a schedule of reinforcement in which a reward follows only some correct responses.

In **partial reinforcement schedules,** a reward follows only some correct responses. When our behavior is reinforced *intermittently* (only some of the time), hope springs eternal and we are reluctant to give up. If a vending machine is a good example of continuous reinforcement, a lottery is a good example of partial reinforcement. People don't expect to win every time they buy a ticket. Therefore, they will continue to buy tickets even if they don't win. As lottery commissions know, partial reinforcement schedules produce behavior that is hard to extinguish.

There are four partial reinforcement schedules.[16] Two of the partial reinforcement schedules, called *interval schedules,* focus on the time that elapses between reinforcements. The other two partial reinforcement schedules, called *ratio schedules,* focus on the number of responses required before reinforcement occurs. Let's take a closer look at each of these four partial reinforcement schedules.

fixed-interval schedule In operant conditioning, a partial reinforcement schedule that rewards only the first correct response after some defined period.

Fixed-Interval Schedule

A **fixed-interval** partial reinforcement schedule rewards only the first correct response after some defined period of time has passed. For example, a researcher might always reinforce the first time a rat presses a bar after 60 seconds have passed. After receiving a food pellet (a reinforcement) for that response, the rat has to wait 60 seconds before it will be reinforced for another correct response. The interval (60 seconds) is unchanging, and there is no way the rat can earn reinforcement until the end of that 60-second interval—thus the term *fixed-interval schedule.*

A rat experienced with a fixed-interval schedule learns not to respond during the first part of the fixed time interval, when there is no way to earn a reinforcement. Toward the end of the interval, the rat starts pressing the bar, checking to see if the time is up. The rate of checking increases as the end of the fixed interval approaches. The result is the response pattern you can see in the yellow line in **Figure 17.5**.

Do fixed-interval schedules happen in real life? You bet! Have you ever had a class with a quiz every Friday? If so, you were being reinforced (with a good grade) for your behavior (studying) on a fixed-interval (once a week) schedule. Did you study much for the quiz on Monday, Tuesday, or Wednesday? Many students don't. Instead, they pack the main part of their responding (studying) into the end of the interval (Thursday night), just as the rats did in Skinner's research on fixed-interval schedules.

variable-interval schedule In operant conditioning, a partial reinforcement schedule that rewards the first correct response after an unpredictable amount of time.

Variable-Interval Schedule

A **variable-interval** partial reinforcement schedule rewards the first correct response after an unpredictable amount of time has passed. The amount of time changes after each reinforcement, so a bar-pressing rat has no way to know how long the interval will be. The rat must keep checking

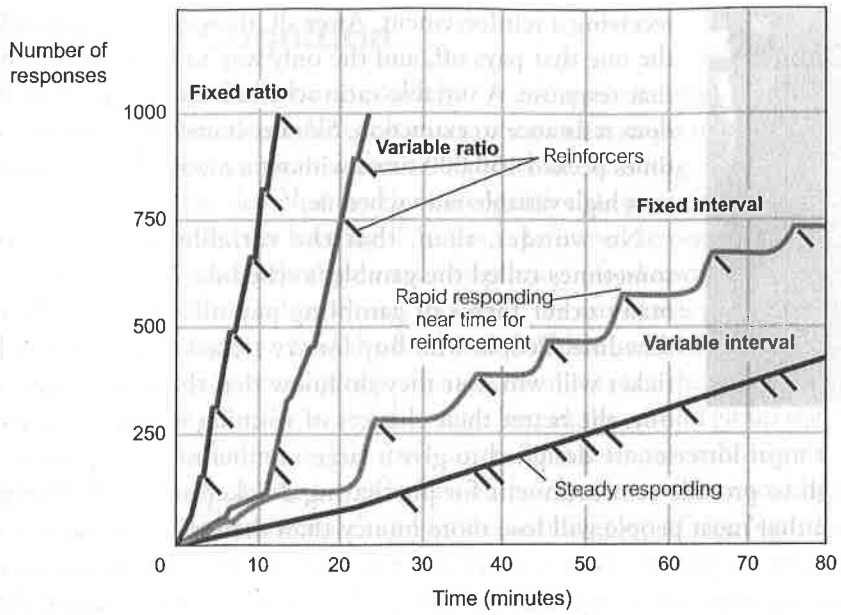

FIGURE 17.5
The Cumulative Record Skinner used graphs like this one to track the responses of pigeons and rats. The horizontal axis of the graph measures time, and the vertical axis measures total responses. The small slash marks indicate points at which the animal received a reinforcement (such as a food pellet). The slopes of the lines connecting the slashes show the rate of responding—a steeper line indicates more responses per minute. Skinner's pigeons showed these four patterns of responding for the partial schedules of reinforcement. The ratio schedules produce a faster response rate (a steeper line) than the interval schedules. (Adapted from Skinner, 1961.)

by pressing the bar to see whether anything happens. When the variable interval is up, the next correct press on the bar earns the rat its food pellet.

Recall that on the fixed-interval schedule, the rat (or the student) can sit out the first part of the interval (that is, not respond) without risking the loss of a reward. Not so on a variable-interval schedule, where any given time interval might be short. Instead, the rat learns to respond at a moderate, steady rate, as the red line in Figure 17.5 shows. Faster responding doesn't result in many extra rewards, so response rate is not very important with a variable-interval schedule.

As many psychology teachers know, pop quizzes operate on a variable-interval schedule. When a quiz can occur at any time, it's best to study a little bit each day. If there is no quiz today, there may be one tomorrow. If there is a quiz today, there may be another one tomorrow. To earn the most reinforcement (good grades on your quizzes), you must be a steady studier.

LIFE MATTERS
As a student, pop quizzes seem like the worst. However, according to The Learning Scientists, regularly scheduled, spaced out quizzes will lead to more long-term retention of information, which will come in handy during exam time.

Fixed-Ratio Schedule

A **fixed-ratio** partial reinforcement schedule provides a reward only after a certain number of correct responses. The word *ratio* in the term refers to the ratio of reinforcements to responses, such as 1 reinforcement for every 20 correct responses. Fixed-ratio schedules do place a premium on speedy responding: The faster the rat makes the required number of responses, the faster it will be fed, which means more to eat for a hungry rat. A rat with some experience on the fixed-ratio schedule will run through the required number of responses rapidly. As you can see in the blue line in Figure 17.5, it will then take a short break. After "catching its breath," the rat will run through the next set of responses as rapidly as possible.

Coffee shops that run buy 10, get 1 free specials are using a fixed-ratio schedule of reinforcement. You may have found yourself buying 10 lattes quickly to qualify for your reinforcement—the free drink. After enjoying your free coffee, you may not buy again for a while. (You're probably a little sick of coffee after drinking several a day!) However, after catching your breath for a few days, you may be back to ordering double espressos again to earn the next free drink.

fixed-ratio schedule In operant conditioning, a partial reinforcement schedule that rewards a response only after some defined number of correct responses.

Variable-Ratio Schedule

A **variable-ratio** partial reinforcement schedule rewards an unpredictable number of correct responses. The number of correct responses is unpredictable because it changes after each reinforcement. Rats on a variable-ratio schedule tend to respond fast and to continue responding after

variable-ratio schedule In operant conditioning, a partial reinforcement schedule that rewards after an unpredictable number of correct responses.

Eric Raptosh Photography/Getty Images

Schedules of Reinforcement in Real Life
This gambler is being reinforced on a variable-ratio schedule. She does not know how many times she has to play to win. The variable-ratio schedule produces a high, steady response rate, much to the delight of the casino owners.

receiving a reinforcement. After all, the next response *could* be the one that pays off, and the only way to find out is to make that response. A variable-ratio schedule also produces tremendous resistance to extinction. Skinner found that pigeons sometimes pecked 150,000 times without a reward after having been on a high variable-ratio schedule.[17]

No wonder, then, that the variable-ratio schedule is sometimes called the gambler's schedule. Lottery tickets and many other forms of gambling pay off on a variable-ratio schedule. People who buy lottery tickets never know which ticket will win, but they do know that the more tickets they buy, the better their chances of winning will be. The trouble is that most lotteries are designed to give a large number of small payouts, big enough to provide reinforcement for purchasing a ticket but small enough to ensure that most people will lose more money than they win. I stood behind a woman in the grocery store one day and watched her spend $5 on five scratch-off tickets. Four were losers, and the fifth paid off with two "free" tickets. One of these paid $2, and as she left she turned to her friend and said, "I won!" She had actually lost $3 ($5 minus her $2 win), yet her small win had reinforced her ticket-buying behavior. Do you think she was motivated to buy additional tickets the next time she was at the store?

LIFE MATTERS

Thinking about purchasing a lottery ticket? The odds of winning the Mega Millions is 1 in 302.5 million. You are more likely to get into Harvard, get killed by a vending machine, or get elected as POTUS, and yet the average American spends $207 annually on lottery tickets.

MAKE IT STICK!

1. Reinforcing only some correct responses is called
   a. negative reinforcement.
   b. partial reinforcement.
   c. primary reinforcement.
   d. continuous reinforcement.
2. If you don't know when to expect text messages from your friends, checking to see if you've received a text is reinforced on a _____ schedule.
   a. variable-interval
   b. variable-ratio
   c. fixed-interval
   d. fixed-ratio
3. A fitness tracker that produces a starburst screen every time you complete 10,000 steps in a day utilizes _____ reinforcement.
4. True or false? A variable-interval schedule of reinforcement usually produces a consistent rate of responding.

New Understandings of Operant Conditioning

17-7 How do cognition and biology affect the operant conditioning process?

Cognition (our thought processes) affects all types of learning, including operant learning. Furthermore, our biology sets boundaries for how and what we can learn.

The Role of Cognition

With **latent learning,** the learning occurs but is not apparent until the learner has an incentive to demonstrate it. Latent learning is a good example of how our thinking—not just whether we are reinforced—affects our learning. In a now-classic experiment, researchers demonstrated how a **cognitive map,** or mental representation of a place, can influence learning.[18] In this experiment, researchers trained one group of rats to find its way through a maze by putting a food reward in a box at the end of the maze. As the number of trials increased, the rats in this group completed the maze faster and faster to find the food (see the blue line in **Figure 17.6**). The researchers also placed a second group of rats in the maze, but they did not reward them with food for finding the end of the maze. Rats in this second group wandered through the maze, exploring it, but their times did not improve during this first phase of the experiment. In the second phase of the experiment, the researchers put a food reward in the box at the maze's end for the second group of rats. Now the performance of the second group of rats rapidly improved until they made fewer errors than the first group. The rats in the second group had developed a cognitive map as they wandered through the maze in the first phase of the experiment. They had learned, but the learning occurred cognitively before it was expressed behaviorally.

latent learning Learning that occurs but is not apparent until the learner has an incentive to demonstrate it.

cognitive map The mental representation of a place.

The **overjustification effect** provides more support for cognition's role in operant conditioning. Rewarding an already enjoyable behavior *overjustifies* it and may actually *decrease* the frequency of that behavior. This is the direct opposite of the effect Skinner's principles would predict—an increase in the rewarded behavior. Unfortunately, overjustification sometimes happens in school. Activities like reading, which should be (and are) naturally reinforcing for most young children, can be overjustified if the school provides lots of special rewards, such as gold stars, grades, and special parties for meeting reading goals.

overjustification effect The effect of promising a reward for doing what one already likes to do; the reward may lessen and replace the person's original, natural motivation so that the behavior stops if the reward is eliminated.

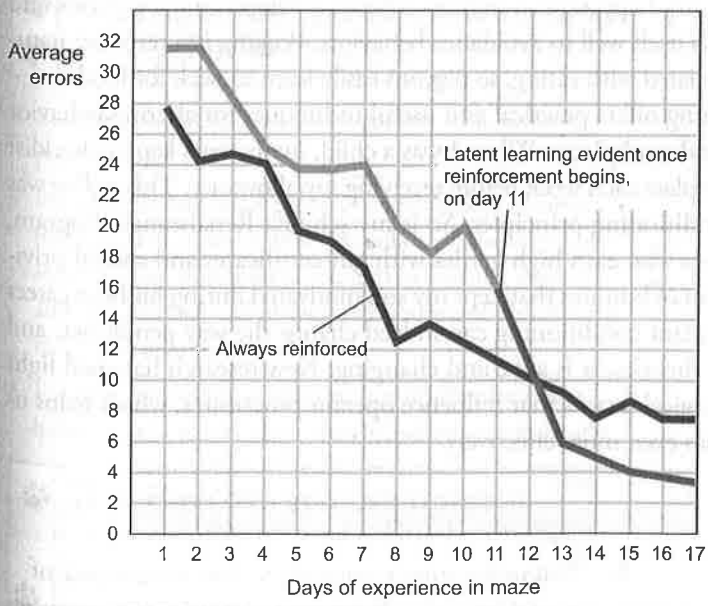

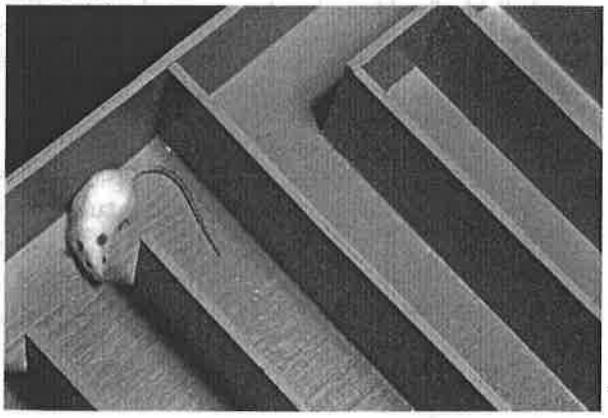

Will & Deni McIntyre/Science Source

FIGURE 17.6
Learning Without Reward
If this rat is allowed to wander through the maze on several occasions, it will develop a cognitive map of the maze—an example of latent learning. But the rat will not demonstrate that learning until researchers add some positive reinforcement for showing the knowledge of the maze paths. (From Tolman & Honzik, 1930.)

LIONEL CIRONNEAU/AP Images

Matching the Species to the Behavior
Operant conditioning works best when it focuses on behaviors that come naturally to a species, such as jumping-up behaviors in dogs.

The danger is that these rewards may begin to overwhelm the child's natural motivation to read and become the primary means of maintaining the behavior. When the rewards stop, the behavior stops as well. ("What, I don't get free pizza for reading 10 books?")

Even grades can lead to overjustification. One experiment divided fifth graders into two groups. The researchers told one group to read a passage and informed students in the group that they would be graded on how much they'd learned. They told the other group that they should read the passage but they wouldn't be graded, simply questioned to find out what they remembered. There was little difference in how much the two groups remembered, but the second group thought the passage was more interesting.[19] There seems to be a link between external environmental rewards and internal cognitive factors. If we ignore cognition, we won't get the whole picture about how operant conditioning works. (To see how cognition can influence personal habits for the better, see Learning New Habits.)

The Role of Biology

Why don't the principles of operant conditioning work equally well for all behaviors in all species?

Erik Isakson/Tetra Images/Alamy Stock Photo

Individualized Learning
Many electronic instructional programs take advantage of operant conditioning principles. This girl will have her answers shaped by a program that breaks instruction down into a series of easy steps. She will receive positive reinforcement for each correct answer.

Like cognition, biology clearly influences how we learn and what we learn, including what we learn through operant conditioning. Some species are *biologically predisposed* to learn some behaviors easily and other behaviors only with great difficulty.

Pigeons, for example, easily learn to flap their wings to avoid a mild electric shock and to peck at a disk for food. Operant conditioning principles would suggest that you could reverse these two behaviors and teach the birds to flap for food and peck to avoid shock. This turns out to be difficult, however, because it defies the biological tendencies of the species.[20] Wing flapping is a natural defense mechanism for pigeons and thus lends itself well to avoidance behaviors. Pecking is a response naturally associated with eating, so pigeons easily learn to peck for food.

Operant conditioning offers practical and useful techniques for altering behavior in families, schools, and workplaces. When I was a child, my parents kept a checklist of chores I had to complete each week before receiving my allowance. This policy was based on operant conditioning principles. So is my school's Renaissance Program, which rewards students who earn high grades with gift certificates and special privileges. So was the system of bonuses that kept my son motivated during his brief career in telemarketing. Operant conditioning can indeed change the way people act, and like all science-based theories, it is alive and changing. New research has shed light on cognitive and biological factors that influence operant procedures, which helps us employ these principles even more effectively.

MAKE IT STICK!

1. Rewarding individuals for behavior they already enjoy may cause them to do the behavior less when the rewards stop. This is called the _____ effect.

2. Latent learning demonstrates the importance of _____ in operant conditioning.
   a. extinction
   b. biology
   c. punishment
   d. cognition

3. True or false? Biology helps determine what behaviors will be easiest for an animal to learn.