Module 20:
Operant Conditioning

Overview

Sections

- What Is Operant Conditioning?
- The Law of Effect
- Reinforcement
- Punishment
- Reinforcement Procedures
- Schedules of Reinforcement
- New Understandings of Operant Conditioning

Learning goals

Students will be able to:

1. Define operant conditioning.
2. Define the law of effect, and explain the different kinds of reinforcement.
3. Describe the effects of punishment and the disadvantages of using punishment to control behavior.
4. Explain how behaviors are influenced through shaping, discrimination, and extinction.
5. Contrast the effects of different schedules of reinforcement.
6. Discuss how cognition and biology influence the operant conditioning process.

Vocabulary Previewing Key Terms and Key People:

operant conditioning
reinforcement
punishment
positive reinforcement
negative reinforcement
primary reinforcement
secondary reinforcement
shaping
discrimination
extinction
continuous reinforcement
partial reinforcement schedule
fixed-interval schedule
variable-interval schedule
fixed-ratio schedule
variable-ratio schedule
latent learning
cognitive map
overjustification effect
Edward Thorndike (1874–1949)
B. F. Skinner (1904–1990)

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

What Is Operant Conditioning?

THINKING CRITICALLY *What is operant conditioning?*

A few years ago, I was in the grocery store on a busy afternoon. In front of me in the checkout line was 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." Actually, 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 20.1). 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 from the table after completing your meal.

Daughter	Mother
Behavior: Screaming tantrum	Behavior: Giving candy
Consequence: Receiving a candy bar	Consequence: Screaming tantrum ends
Result: More tantrums in the future	Result: More candy buying in the future

Figure 20.1 Operant Conditioning for Better or for Worse This child is using operant conditioning to train her mother to buy candy, and the mother is operantly conditioning the child to have tantrums.

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 the frequency with which 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. Take a closer look at how operant conditioning techniques work.

THINKING CRITICALLY SUMMARY *Operant conditioning is a type of learning in which the frequency of a behavior depends on the consequence that follows that behavior.*

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

B. F. Skinner (1904–1990)
Developed the fundamental principles and techniques of operant conditioning and devised ways to apply them in the real world.

The Law of Effect

THINKING CRITICALLY *What is the law of effect?*

Operant conditioning developed from an idea of **Edward Thorndike,** an early American psychologist. Thorndike's idea was the *law of effect,* which simply states that behaviors with favorable consequences will occur more frequently and behaviors followed by less favorable consequences will occur less frequently. Another American psychologist, **B. F. Skinner,** built his life's work on this idea, developing the fundamental principles and techniques of operant conditioning (see "Psychology Is a Science: B. F. Skinner," pages 388–389). Two of the more important concepts used in operant conditioning are reinforcement and punishment.

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?

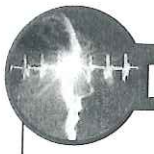

PSYCHOLOGY IS A SCIENCE

B. F. Skinner

Few have done more than B. F. Skinner to advance the notion of psychology as a scientific discipline. In one recent study, almost 2000 psychologists ranked Skinner as the most eminent psychologist of the twentieth century (Haggbloom & others, 2002). 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.

Richard Wood/The Picture Cube

B. F. Skinner's Influence in the Laboratory This rat in a Skinner box, or operant chamber, has been taught to press a bar for a food reinforcement.

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 a 1945 magazine article 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).

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.

- **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 did a chore, he would have been more likely to do the chore in

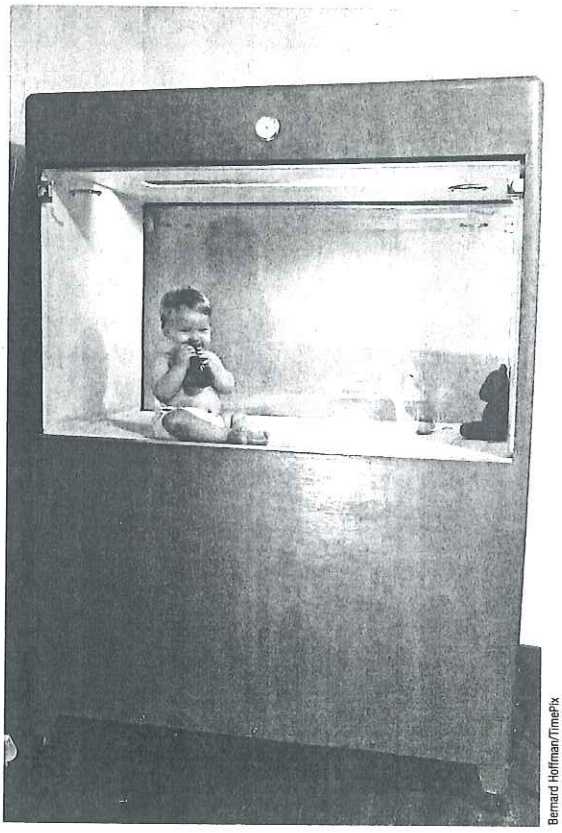

Bernard Hoffman/TimePix

The "Baby in the 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.

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, including Los Horcones in Sonora, Mexico, where Skinner believed his principles were applied most appropriately. In the 1970s, 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 (Skinner, 1990).

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, and at work. 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.

the future. For him, broccoli was a *reinforcement.* But Eric did *not* like broccoli. If I had given Eric broccoli after he did a chore, 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 view being yelled at as a form of desired parental attention. If this is the case, the law of effect predicts

Who Teaches Whom?
What behavior has been reinforced?

Hi and Lois

the behavior that preceded the parent's yelling will be *more* likely to happen again. Likewise, a school administrator might suspend a student from school 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!

THINKING CRITICALLY *Thorndike's law of effect says that behaviors with favorable consequences will occur more frequently and behaviors with less favorable consequences will occur less frequently.*

Reinforcement

THINKING CRITICALLY *What are the different kinds of reinforcement and how effective are they?*

Reinforcement procedures strengthen responses by making them more likely to occur again. There are two ways to reinforce a behavior: positive reinforcement and negative reinforcement (see Figure 20.2).

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.

POSITIVE REINFORCEMENT

Behavior is followed by a desirable event or state.

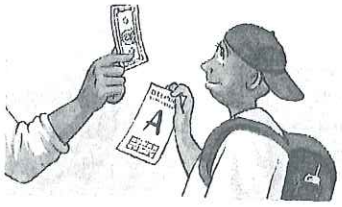

$10 for an A makes it more likely a student will earn more As.

NEGATIVE REINFORCEMENT

Behavior ends an undesirable event or state.

Taking aspirin relieves headaches and makes it more likely that aspirin will be taken in the future.

Figure 20.2 Reinforcement Strengthens Behavior Positive and negative reinforcement work differently, but both types of reinforcement make a behavior more likely to happen again.

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. The words *positive* and *negative* in this context do not mean "good" and "bad." Here, *positive* simply means that something desirable is *presented.* And *negative* means that something undesirable is *removed.* For example, Roshni's headache is undesirable. If taking aspirin provides relief from the headache, Roshni's behavior of taking aspirin has been negatively reinforced. Negative reinforcement, like all reinforcement, *strengthens* a behavior. So, Roshni becomes more likely to take aspirin to escape a headache in the future.

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.

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 candy. The end of the tantrum negatively reinforced the mother's candy-buying behavior, which is likely to recur in the future.
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 (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 more long term. It's no surprise, then, that drugs that produce the most immediate reinforcement, like nicotine and cocaine, are the most addictive (Marlatt, 1991). 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 for that behavior is delayed by 30 seconds or more. But humans 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 aren't issued until the end of the pay period 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 (Mischel & others, 1989).

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 20.3).

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 6-month-old baby, he will probably put it in his mouth to see if it tastes good. (He's checking for a primary reinforcer!) When the baby discovers the money doesn't taste good, he 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 value—so much value that you would work hard to get more.

Figure 20.3 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 basic rewards.

PRIMARY REINFORCEMENT

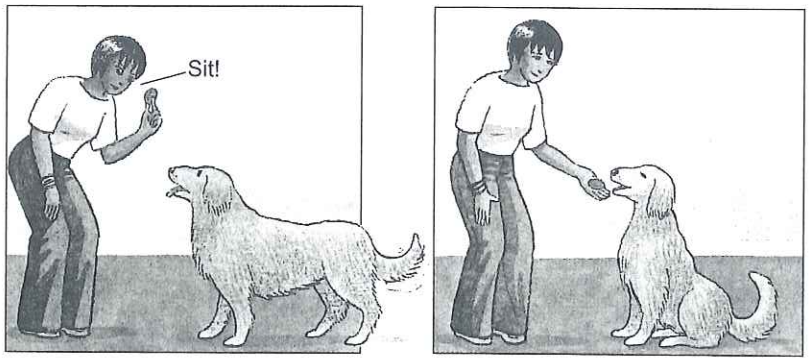

Food is a primary reinforcer for a dog.

SECONDARY REINFORCEMENT

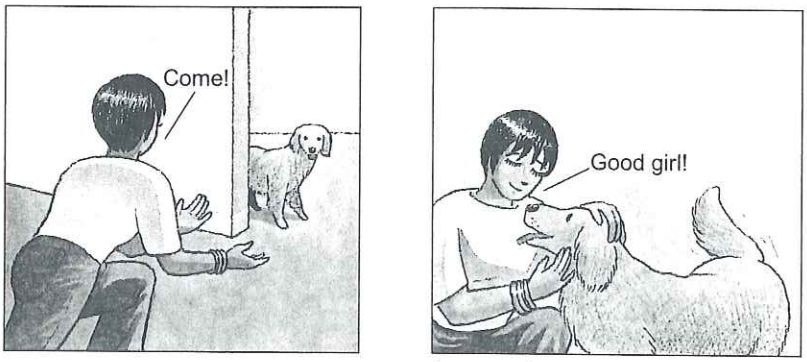

An owner's words can become secondary reinforcement when they're associated with petting and approval.

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 her a B+ 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 affect on their behavior.

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.

THINKING CRITICALLY SUMMARY *All reinforcements increase the likelihood of behaviors. Positive reinforcements follow a desired behavior with a desirable event or state; negative reinforcement follows a desired behavior with the removal of an undesirable event or state. Reinforcers that are delivered soon after the behavior (immediate) are more effective than reinforcers delivered long after the behavior occurred (delayed). Finally, reinforcers that are "naturally" rewarding (like food) are called primary reinforcers, and reinforcers that we learn are desirable are called secondary reinforcers.*

Punishment

THINKING CRITICALLY *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 20.4):

- *The behavior leads to an undesirable event.* 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.

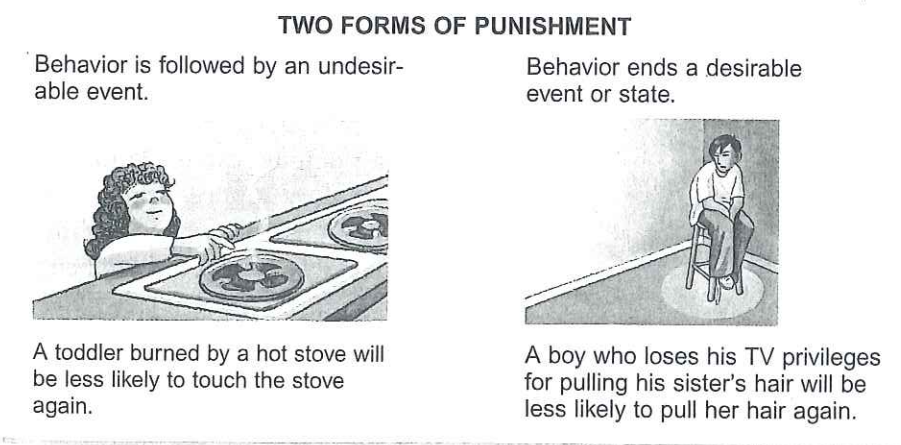

Figure 20.4 Punishment Weakens Behavior Here are two forms of punishment. In one, the punished behavior is followed by an undesirable event. In the other, the punished behavior is followed by the loss of a desirable event or state. Although the types of punishment differ, each decreases the likelihood that the behavior will happen again.

- *The behavior ends a desirable state or event.* 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.

Problems with Punishment

Many learning experts oppose the use of punishment to control behavior. 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 just purchase a radar detector rather than drive moderately.

Punishment can also lead to fear, anxiety, and lower self-esteem. Frequently punished children or animals 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 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) (Straus & Gelles, 1980).

So, despite all these problems, why is punishment used so often? One reason seems to be that when a punished individual is obedient or contrite, even for just a few minutes, this consequence positively reinforces the behavior it followed—which in this case is punishment. This desirable consequence in turn makes the 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.

Bob Daemmrich/Stock, Boston

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

ıere is a role for punishment in learning, but it is a limited one. sure punishment can effectively control certain behaviors, espe- if the punisher's goal is to protect a child from a dangerous situa- For example, if a toddler has developed the bad habit of running he street, a harsh reprimand or swat on the behind may be appro- . A young child needs to develop some fear and avoidance of the . But punishment is generally most effective when used least. Have ver had a class where the teacher was constantly punishing the stu- by losing her temper and yelling at the class? How effective was acher's behavior after it occurred several times? Compare that sce- with a class's reaction to angry behavior on the part of a teacher arely "loses it." If your memories are similar to mine, you'll see that ss often punishment happens, the more effective it is.

r all the reasons we've discussed here, most psychologists recom- *reinforcing an incompatible behavior* as an effective alternative to hment. Rather than punishing a child for lying, for example, par- ıight consider reinforcing the child for telling the truth. This will ıse the amount of truth telling, and because a child cannot tell the and lie at the same time, the amount of lying must decrease. The philosophy here is to "catch the child being good" and reinforce dingly. This approach will lead to a more gradual change of behav- ıt the change will be more permanent than the temporary suppres- ıf behavior that follows punishment. Reinforcement also has two benefits. First, it tends to lead to *approach* behaviors that draw e together (one reason kids are often so eager to see their mostly rcing grandparents) rather than feelings of fear, anxiety, and low teem. Second, children who model positive reinforcement are fun to be with than children who model aggressive behaviors.

KING CRITICALLY SUMMARY *Punishment is anything that ıses the likelihood of a behavior. It occurs when a behavior is followed undesirable event or when a behavior is followed by the removal of a ble state or event. Many learning experts point out that trying to control ors using punishments can cause undesirable side effects. Punishment ecrease the behavior, but it does not affect the internal desire to engage behavior. Also, punishment can be stressful, and children are likely to physically aggressive forms of punishment. Punishment can effectively dangerous behaviors (like a child running into the street) if it is used gly and immediately after the dangerous behavior occurs.*

nforcement Procedures

KING CRITICALLY *How can you use operant conditioning to a new behavior or make an operantly conditioned behavior stop?*

hat you know a little bit about how reinforcement and punish- work, we can begin to explore some other procedures that make

Marc Romanelli/The Image Bank

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.

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.

To train a pigeon to turn in clockwise circles, for example, you would start by providing a food reward every time the pigeon turns 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 is to gradually extend how far the pigeon must turn its head before you give it a food reward. After a series of increasingly longer turns, the pigeon will finally turn all the way around. By breaking the circling behavior into a succession of gradual steps, you can easily shape 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 the seat, ran beside you until you were reasonably well balanced, and then let go. You probably managed to roll several feet on your own before falling, at which point you were rewarded with a hearty, "Good job!" It really wasn't that good a job—you only made it a few feet—but for a first attempt it deserved reinforcement. Gradually, as your riding skills improved, your trainer made you ride farther and farther before giving you a compliment. You were being shaped!

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

discrimination
Ability to distinguish between two similar signals or stimuli.

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

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 *extinction.* **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.

Jeff Greenberg/Stock, Boston

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

Life would be fairly chaotic if we made the same response to all stimuli that were similar. For safety reasons, students and teachers need to learn to *discriminate* between class bells and fire alarms. We do this 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 immediately 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.

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

THINKING CRITICALLY SUMMARY *To teach a new behavior using operant conditioning, you would shape the behavior. This means that you would reward behaviors as they become more similar to the one you want to teach. Discrimination is a learned response of behaving differently toward similar signals. Finally, operantly conditioned behaviors stop when a consequence no longer follows the behavior.*

Signals Can Be Powerful! The dog is occasionally reinforced with a bowl full of food after hearing the can opener. That's enough to send him running every time he hears it!

Dennis the Menace

"I think Mom's using the can opener."

Schedules of Reinforcement

THINKING CRITICALLY *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 do. 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 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 accustomed 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

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 (Skinner, 1961). Two of the partial reinforcement schedules, called *interval schedules,* focus on the time that elapses between reinforcements. The other two partial

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

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

reinforcement schedules, called *ratio schedules,* focus on the number of responses before reinforcement occurs. Take a closer look at each of these four partial reinforcement schedules.

Fixed-Interval Schedule

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

A **fixed-interval schedule** is a partial reinforcement schedule that 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 fixed, and there is no way the rat can earn reinforcement during 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 interval, when there is no way to earn a reinforcement. Toward the end of the fixed 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 Figure 20.5.

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? Most 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.

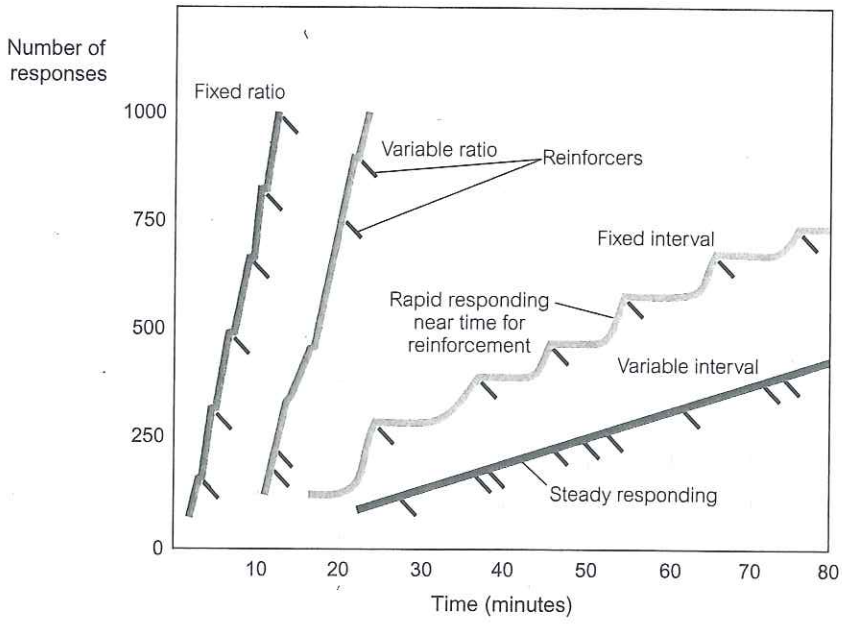

Figure 20.5 The Cumulative Record Skinner used graphs like this one to track the responses of pigeons and rats. The horizontal axis of the graph measures time, and the vertical axis measures total responses. The small slash marks indicate points at which the animal received a reinforcement (such as a 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. Notice that the ratio schedules produce a faster response rate (a steeper line) than the interval schedules. (Adapted from Skinner, 1961.)

Variable-Interval Schedule

A **variable-interval schedule** is a partial reinforcement schedule that 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" 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 reward.

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 risk. Not so on a variable-interval schedule, where any given interval might be short. Instead, the rat learns to respond at a moderate, steady rate, as Figure 20.5 shows. Fast responses don't result in extra rewards, so speed is not important with a variable-interval schedule.

Pop quizzes use a variable-interval schedule. When a quiz can occur at any time, you have 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.

Fixed-Ratio Schedule

A **fixed-ratio schedule** is a partial reinforcement schedule that 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 speed: 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 Figure 20.5, it will take a short break. After "catching its breath," the rat will run through the next set of responses as rapidly as possible.

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

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

Movie rental stores that run "rent 10, get 1 free" specials are using a fixed-ratio schedule of reinforcement. You may have found yourself renting 10 DVDs quickly to qualify for your reinforcement—the free DVD. After watching the free DVD, you may not rent again for a while. (You're probably pretty sick of movies after watching them several nights in a row.) However, after "catching your breath" for a few days, you and your friends may start renting again to earn the next free DVD.

Variable-Ratio Schedule

A **variable-ratio schedule** is a partial reinforcement schedule that 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 receiving a reinforcement. After all, the next response *could* be the response 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 (1953) found that pigeons sometimes pecked 150,000 times without a reward after having been on a high variable-ratio schedule.

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?

Bob Daemmrich/Stock, Boston

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.

THINKING CRITICALLY SUMMARY *Continuous reinforcement establishes a strong association between the behavior and its consequence, so it is effective when teaching something new. Behaviors learned through continuous reinforcement become extinct quickly when the reinforcement is removed, however. Partial reinforcement takes longer but is more resistant to extinction. The four partial reinforcement schedules are fixed interval, variable interval, fixed ratio, and variable ratio.*

New Understandings of Operant Conditioning

THINKING CRITICALLY *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.

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

cognitive map
Mental representation of a place.

overjustification effect
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.

The Role of Cognition

With **latent learning,** the learning occurs but is not apparent until th learner has an incentive to demonstrate it. Latent learning is probab the best example of how our thinking—not just whether we are rei forced—affects our learning. In a now-classic experiment, researche demonstrated how a **cognitive map,** or mental representation of place, can influence learning (Tolman & Honzik, 1930). In this exper ment, researchers trained one group of rats to find its way through maze by putting a food reward in a box at the end of the maze. As th number of trials increased, the rats in this group completed the maz faster and faster (see Figure 20.6). The researchers also placed a secon group of rats in the maze, but they did not reward them with food fc finding the end of the maze. Rats in this second group wandere through the maze, exploring it, but their times did not improve durin this first phase of the experiment. In the second phase of the experimen the researchers put a food reward in the box at the maze's end for th second group of rats. Now the performance of the second group of rat rapidly improved until it matched that of the first group. The rats in th second group had developed a cognitive map as they wandered throug the maze in the first phase of the experiment. They had learned, but th learning occurred cognitively before it was expressed behaviorally.

The overjustification effect provides more support for cognition role in operant conditioning. The **overjustification effect** is the resul of promising a reward for doing what one already likes to do. Th reward may replace the person's original motivation so that the behav ior stops if the reward is eliminated. Rewarding an already enjoyabl behavior overjustifies it and may actually *decrease* the frequency of tha behavior. This is the direct opposite of the effect Skinner's principle

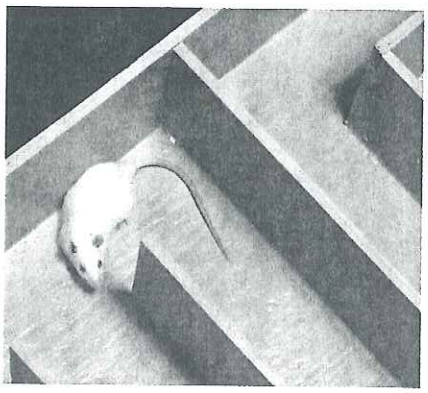

Will and Deni McIntyre/Photo Researchers

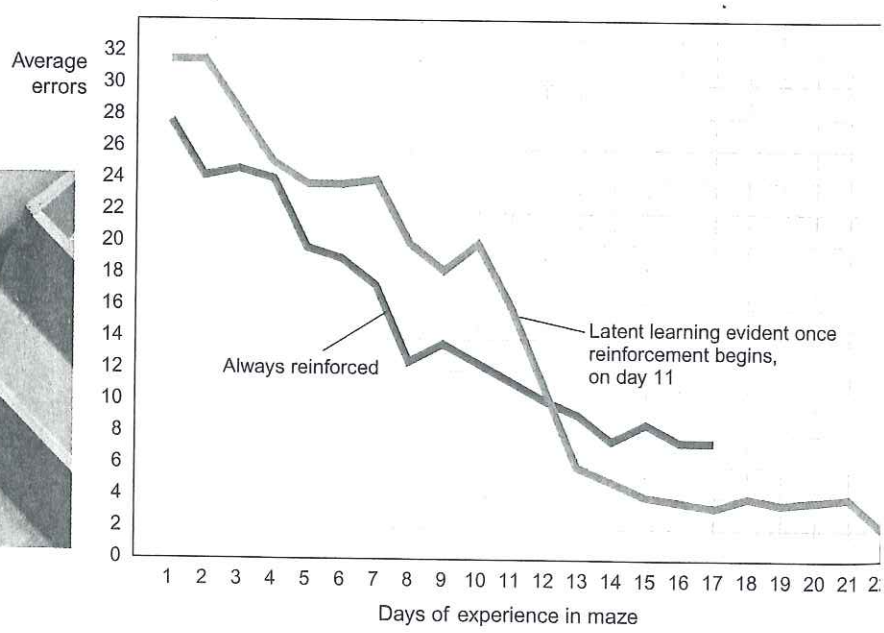

Figure 20.6 Learning Without Reward If this rat is allowed to wander through the maze on several occasions, it will develop a cognitive map of the maze—an example of latent learning. But the rat will not demonstrate that learning until researchers add some positive reinforcement for showing the knowledge of the maze paths. (From Tolman & Honzik, 1930.)

l predict—an increase in the rewarded behavior. Unfortunately, ıstification sometimes happens in school. Activities like reading, ı should be (and are) naturally reinforcing for most young chil- can be overjustified if the school provides lots of special rewards, ıs gold stars, grades, and special parties for meeting reading goals. langer is that these rewards may begin to overwhelm the child's al motivation to read and become the primary means of maintain- ıe behavior. When the rewards stop, the behavior stops as well. at, I don't get a free pizza for reading 10 books?")

ren grades can lead to overjustification. One experiment divided graders into two groups. The researchers told one group to read a ge and informed students in the group that they would be graded ow much they'd learned. They told the other group that they should the passage but they wouldn't be graded, simply questioned to find hat they remembered. There was little difference in how much the roups remembered, but the second group thought the passage was interesting (Grolnick & Ryan, 1987). The point is that there seems a link between external, environmental rewards and internal, cog- factors. If we ignore cognition, we won't get the whole picture t how operant conditioning works.

Role of Biology

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

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

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.

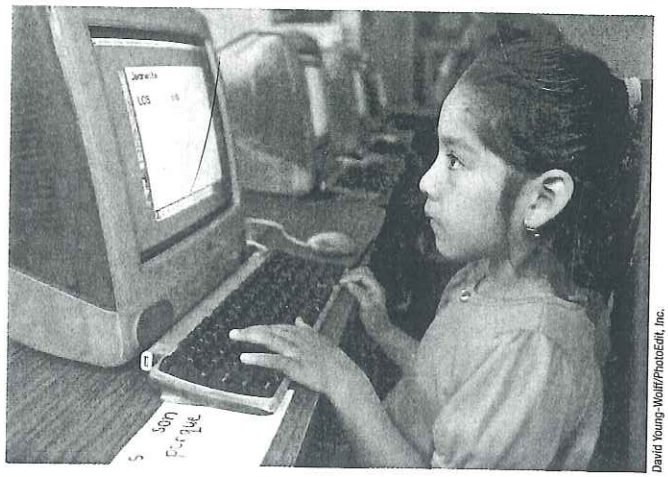

David Young-Wolff/PhotoEdit, Inc.

Individualized Learning Many computerized 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.

Pigeons, for example, easily learn to flap their wings to avoid 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 (Foree & LoLordo, 1973). 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.

THINKING CRITICALLY *Both cognition and biology influence operant conditioning in ways that Skinner never imagined (or admitted). Some examples of learning are clearly influenced by how we think: Latent learning is learning that does not become apparent until the learner has an incentive to demonstrate it. Cognitive maps, or mental representation of places, also affect learning. The overjustification effect happens when a reward for doing what one already likes to do replaces the person's original motivation so that the behavior stops if the reward is eliminated. Other examples of learning demonstrate the influence of biology: Pigeons can more easily be conditioned to perform behaviors that match biological tendencies. They will learn to peck for food, but it is hard to teach them to flap their wings for food.*