

Module 27

Operant Conditioning

Module Learning Objectives

- 27-1** Describe operant conditioning, and explain how operant behavior is reinforced and shaped.
- 27-2** Discuss the differences between positive and negative reinforcement, and identify the basic types of reinforcers.
- 27-3** Explain how the different reinforcement schedules affect behavior.
- 27-4** Discuss how punishment and negative reinforcement differ, and explain how punishment affects behavior.
- 27-5** Describe the controversy over Skinner's views of human behavior.

Henrik Sorenson



Operant Conditioning

- 27-1** What is operant conditioning, and how is operant behavior reinforced and shaped?

It's one thing to classically condition a dog to salivate at the sound of a tone, or a child to fear moving cars. To teach an elephant to walk on its hind legs or a child to say *please*, we turn to operant conditioning.

Classical conditioning and operant conditioning are both forms of associative learning, yet their difference is straightforward:

- *Classical conditioning* forms associations between stimuli (a CS and the US it signals). It also involves *respondent behavior*—actions that are automatic responses to a stimulus (such as salivating in response to meat powder and later in response to a tone).
- In **operant conditioning**, organisms associate their own actions with consequences. Actions followed by reinforcers increase; those followed by punishers often decrease. Behavior that *operates* on the environment to *produce* rewarding or punishing stimuli is called *operant behavior*.

Skinner's Experiments

B. F. Skinner (1904–1990) was a college English major and an aspiring writer who, seeking a new direction, entered psychology graduate school. He went on to become modern behaviorism's most influential and controversial figure. Skinner's work elaborated on what psychologist Edward L. Thorndike (1874–1949) called the **law of effect**: Rewarded

AP® Exam Tip


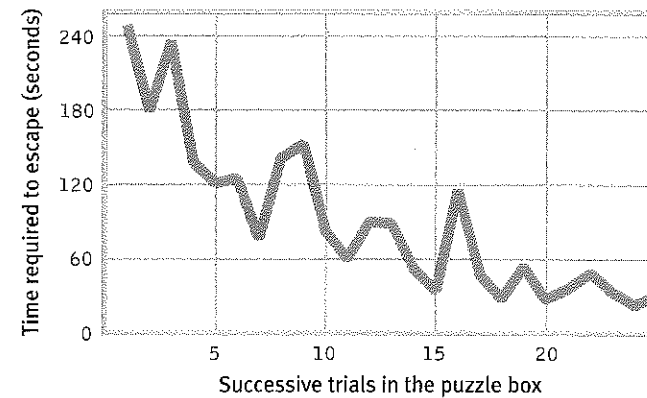
Don't be fooled by the fact that classical conditioning is presented before operant conditioning. Classical conditioning was understood before operant conditioning, but operant conditioning has a larger impact on our day-to-day lives.

operant conditioning a type of learning in which behavior is strengthened if followed by a reinforcer or diminished if followed by a punisher.

law of effect Thorndike's principle that behaviors followed by favorable consequences become more likely, and that behaviors followed by unfavorable consequences become less likely.

Figure 27.1


Cat in a puzzle box
Thorndike used a fish reward to entice cats to find their way out of a puzzle box (right) through a series of maneuvers. The cats' performance tended to improve with successive trials (left), illustrating Thorndike's law of effect. (Adapted from Thorndike, 1898.)



behavior is likely to recur (**FIGURE 27.1**). Using Thorndike's law of effect as a starting point, Skinner developed a behavioral technology that revealed principles of *behavior control*. These principles also enabled him to teach pigeons such unpigeon-like behaviors as walking in a figure 8, playing Ping-Pong, and keeping a missile on course by pecking at a screen target.

Figure 27.2

A Skinner box Inside the box, the rat presses a bar for a food reward. Outside, a measuring device (not shown here) records the animal's accumulated responses.




For his pioneering studies, Skinner designed an **operant chamber**, popularly known as a *Skinner box* (**FIGURE 27.2**). The box has a bar (a lever) that an animal presses—or a key (a disc) the animal pecks—to release a reward of food or water. It also has a device that records these responses. This design creates a stage on which rats and others animals act out Skinner's concept of **reinforcement**: any event that strengthens (increases the frequency of) a preceding response. What is reinforcing depends on the animal and the conditions. For people, it may be praise, attention, or a paycheck. For hungry and thirsty rats, food and water work well. Skinner's experiments have done far more than teach us how to pull habits out of a rat. They have explored the precise conditions that foster efficient and enduring learning.

Shaping Behavior

Imagine that you wanted to condition a hungry rat to press a bar. Like Skinner, you could tease out this action with **shaping**, gradually guiding the rat's actions toward the desired behavior. First, you would watch how the animal naturally behaves, so that you could build on its existing behaviors. You might give the rat a bit of food each time it approaches the bar. Once the rat is approaching regularly, you would give the food only when it moves close to the bar, then closer still. Finally, you would require it to touch the bar to get food. With this method of *successive approximations*, you reward responses that are ever-closer to the final desired behavior, and you ignore all other responses. By making rewards contingent on desired behaviors, researchers and animal trainers gradually shape complex behaviors.

Shaping can also help us understand what nonverbal organisms perceive. Can a dog distinguish red and green? Can a baby hear the difference between lower- and higher-pitched tones? If we can shape them to respond to one stimulus and not to another, then we know they can perceive the difference. Such experiments have even shown that some animals can form concepts. When experimenters reinforced pigeons for pecking after



Reinforcers vary with circumstances What is reinforcing (a heat lamp) to one animal (a cold meerkat) may not be to another (an overheated child). What is reinforcing in one situation (a cold snap at the Taronga Zoo in Sydney) may not be in another (a sweltering summer day).

seeing a human face, but not after seeing other images, the pigeon's behavior showed that it could recognize human faces (Herrnstein & Loveland, 1964). In this experiment, the human face was a **discriminative stimulus**. Like a green traffic light, discriminative stimuli signal that a response will be reinforced. After being trained to discriminate among classes of events or objects—flowers, people, cars, chairs—pigeons can usually identify the category in which a new pictured object belongs (Bhatt et al., 1988; Wasserman, 1993). They have even been trained to discriminate between the music of Bach and Stravinsky (Porter & Neuringer, 1984).

In everyday life, we continually reinforce and shape others' behavior, said Skinner, though we may not mean to do so. Isaac's whining, for example, annoys his dad, but look how he typically responds:

Isaac: Could you take me to the mall?

Father: (*Ignores Isaac and stays focused on his phone*)

Isaac: Dad, I need to go to the mall.

Father: (*distracted*) Uh, yeah, just a minute.

Isaac: DAAAD! The mall!!

Father: Show some manners! Okay, where are my keys. . .

Isaac's whining is reinforced, because he gets something desirable—his dad's attention. Dad's response is reinforced because it gets rid of something aversive—Isaac's whining.

Or consider a teacher who pastes gold stars on a wall chart beside the names of children scoring 100 percent on spelling tests. As everyone can then see, some children consistently do perfect work. The others, who take the same test and may have worked harder than the academic all-stars, get no rewards. The teacher would be better advised to apply the principles of operant conditioning—to reinforce all spellers for gradual improvements (successive approximations toward perfect spelling of words they find challenging).

Types of Reinforcers

How do positive and negative reinforcement differ, and what are the basic types of reinforcers?

Up to now, we've mainly been discussing **positive reinforcement**, which strengthens a response by *presenting* a typically pleasurable stimulus after a response. But, as we saw in the whining Isaac story, there are *two* basic kinds of reinforcement (**TABLE 27.1** on the next page).


discriminative stimulus

in operant conditioning, a stimulus that elicits a response after association with reinforcement (in contrast to related stimuli not associated with reinforcement).

positive reinforcement

increasing behaviors by presenting positive reinforcers. A positive reinforcer is any stimulus that, when *presented* after a response, strengthens the response.

Shaping a dog to play the piano Using a method of successive approximations, with a food reward for each small step—hopping up on the piano bench, putting her paws on the keys, actually making sounds—this dog was taught to "play" the piano, and now does so frequently!



operant chamber in operant conditioning research, a chamber (also known as a *Skinner box*) containing a bar or key that an animal can manipulate to obtain a food or water reinforcer; attached devices record the animal's rate of bar pressing or key pecking.

reinforcement in operant conditioning, any event that *strengthens* the behavior it follows.

shaping an operant conditioning procedure in which reinforcers guide behavior toward closer and closer approximations of the desired behavior.


Table 27.1 Ways to Increase Behavior

Operant Conditioning Term	Description	Examples
Positive reinforcement	Add a desirable stimulus	Pet a dog that comes when you call it; pay the person who paints your house
Negative reinforcement	Remove an aversive stimulus	Take painkillers to end pain; fasten seat belt to end loud beeping

Negative reinforcement *strengthens* a response by *reducing or removing* something negative. Isaac's whining was *positively* reinforced, because Isaac got something desirable—his father's attention. His dad's response to the whining (taking Isaac to the mall) was negatively reinforced, because it ended an aversive event—Isaac's whining. Similarly, taking aspirin may relieve your headache, and pushing the snooze button will silence your annoying alarm. These welcome results provide negative reinforcement and increase the odds that you will repeat these behaviors. For drug addicts, the negative reinforcement of ending withdrawal pangs can be a compelling reason to resume using (Baker et al., 2004). Note that *negative reinforcement is not punishment*. (Some friendly advice: Repeat the last five words in your mind.) Rather, negative reinforcement *removes* a punishing (aversive) event. Think of negative reinforcement as something that provides relief—from that whining teenager, bad headache, or annoying alarm.

Sometimes negative and positive reinforcement coincide. Imagine a worried student who, after goofing off and getting a bad test grade, studies harder for the next test. This increased effort may be *negatively* reinforced by reduced anxiety, and *positively* reinforced by a better grade. Whether it works by reducing something aversive, or by giving something desirable, *reinforcement is any consequence that strengthens behavior*.

HI AND LOIS



negative reinforcement increasing behaviors by stopping or reducing negative stimuli. A negative reinforcer is any stimulus that, when removed after a response, strengthens the response. (Note: Negative reinforcement is not punishment.)

primary reinforcer an innately reinforcing stimulus, such as one that satisfies a biological need.

conditioned reinforcer a stimulus that gains its reinforcing power through its association with a primary reinforcer; also known as a *secondary reinforcer*.

PRIMARY AND CONDITIONED REINFORCERS

Getting food when hungry or having a painful headache go away is innately satisfying. These **primary reinforcers** are unlearned. **Conditioned reinforcers**, also called *secondary reinforcers*, get their power through learned association with primary reinforcers. If a rat in a Skinner box learns that a light reliably signals a food delivery, the rat will work to turn on the light. The light has become a conditioned reinforcer. Our lives are filled with conditioned reinforcers—money, good grades, a pleasant tone of voice—each of which has been linked with more basic rewards.

IMMEDIATE AND DELAYED REINFORCERS

Let's return to the imaginary shaping experiment in which you were conditioning a rat to press a bar. Before performing this "wanted" behavior, the hungry rat will engage in a sequence of "unwanted" behaviors—scratching, sniffing, and moving around. If you present

food immediately after any one of these behaviors, the rat will likely repeat that rewarded behavior. But what if the rat presses the bar while you are distracted, and you delay giving the reinforcer? If the delay lasts longer than about 30 seconds, the rat will not learn to press the bar. You will have reinforced other incidental behaviors—more sniffing and moving—that intervened after the bar press.

Unlike rats, humans do respond to delayed reinforcers: the paycheck at the end of the week, the good grade at the end of the term, the trophy at the end of the season. Indeed, to function effectively we must learn to delay gratification. In laboratory testing, some 4-year-olds show this ability. In choosing a candy, they prefer having a big one tomorrow to munching on a small one right now. Learning to control our impulses in order to achieve more valued rewards is a big step toward maturity (Logue, 1998a,b). No wonder children who make such choices have tended to become socially competent and high-achieving adults (Mischel et al., 1989).

To our detriment, small but immediate consequences (the enjoyment of late-night videos or texting, for example) are sometimes more alluring than big but delayed consequences (feeling alert tomorrow). For many teens, the immediate gratification of risky, unprotected sex in passionate moments prevails over the delayed gratifications of safe sex or saved sex. And for many people, the immediate rewards of today's gas-guzzling vehicles, air travel, and air conditioning prevail over the bigger future consequences of global climate change, rising seas, and extreme weather.

Reinforcement Schedules

27-3 How do different reinforcement schedules affect behavior?


In most of our examples, the desired response has been reinforced every time it occurs. But **reinforcement schedules** vary. With **continuous reinforcement**, learning occurs rapidly, which makes this the best choice for mastering a behavior. But extinction also occurs rapidly. When reinforcement stops—when we stop delivering food after the rat presses the bar—the behavior soon stops. If a normally dependable candy machine fails to deliver a chocolate bar twice in a row, we stop putting money into it (although a week later we may exhibit spontaneous recovery by trying again).

Real life rarely provides continuous reinforcement. Salespeople do not make a sale with every pitch. But they persist because their efforts are occasionally rewarded. This persistence is typical with **partial (intermittent) reinforcement** schedules, in which responses are sometimes reinforced, sometimes not. Learning is slower to appear, but *resistance to extinction* is greater than with continuous reinforcement. Imagine a pigeon that has learned to peck a key to obtain food. If you gradually phase out the food delivery until it occurs only rarely, in no predictable pattern, the pigeon may peck 150,000 times without a reward (Skinner, 1953). Gambling machines and lottery tickets reward gamblers in much the same way—occasionally and unpredictably. And like pigeons, slot players keep trying, time and time again. With intermittent reinforcement, hope springs eternal.

Lesson for child caregivers: Partial reinforcement also works with children. *Occasionally* giving in to children's tantrums for the sake of peace and quiet intermittently reinforces the tantrums. This is the very best procedure for making a behavior persist.

Skinner (1961) and his collaborators compared four schedules of partial reinforcement. Some are rigidly fixed, some unpredictably variable.

Fixed-ratio schedules reinforce behavior after a set number of responses. Coffee shops may reward us with a free drink after every 10 purchased. In the laboratory, rats may be reinforced on a fixed ratio of, say, one food pellet for every 30 responses. Once conditioned, animals will pause only briefly after a reinforcer before returning to a high rate of responding (**FIGURE 27.3** on the next page).



"Oh, not bad. The light comes on, I press the bar, they write me a check. How about you?"

reinforcement schedule a pattern that defines how often a desired response will be reinforced.

continuous reinforcement reinforcing the desired response every time it occurs.


partial (intermittent) reinforcement reinforcing a response only part of the time; results in slower acquisition of a response but much greater resistance to extinction than does continuous reinforcement.

fixed-ratio schedule in operant conditioning, a reinforcement schedule that reinforces a response only after a specified number of responses.

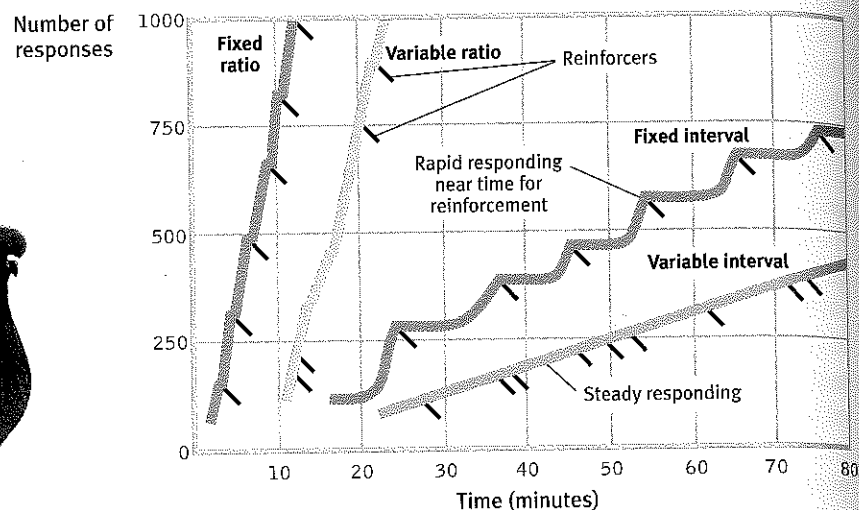
"The charm of fishing is that it is the pursuit of what is elusive but attainable, a perpetual series of occasions for hope." -SCOTTISH AUTHOR JOHN BUCHAN (1875-1940)

Figure 27.3

Intermittent reinforcement schedules
Skinner's laboratory pigeons produced these response patterns to each of four reinforcement schedules. (Reinforcers are indicated by diagonal marks.) For people, as for pigeons, reinforcement linked to number of responses (a *ratio schedule*) produces a higher response rate than reinforcement linked to amount of time elapsed (an *interval schedule*). But the predictability of the reward also matters. An unpredictable (*variable*) schedule produces more consistent responding than does a predictable (*fixed*) schedule. Adapted from "Teaching Machines" by B. F. Skinner. Copyright © 1961, Scientific American, Inc. All Rights Reserved.



Vitaly Titov & Maria Sidelnikova/Shutterstock



Variable-ratio schedules provide reinforcers after a seemingly unpredictable number of responses. This is what slot-machine players and fly-casting anglers experience—unpredictable reinforcement—and what makes gambling and fly fishing so hard to extinguish even when both are getting nothing for something. Because reinforcers increase as the number of responses increases, variable-ratio schedules produce high rates of responding.

Fixed-interval schedules reinforce the first response after a fixed time period. Animals on this type of schedule tend to respond more frequently as the anticipated time for reward draws near. People check more frequently for the mail as the delivery time approaches. A hungry child jiggles the Jell-O more often to see if it has set. Pigeons peck keys more rapidly as the time for reinforcement draws nearer. This produces a choppy stop-start pattern rather than a steady rate of response (see Figure 27.3).

Variable-interval schedules reinforce the first response after *varying* time intervals. Like the longed-for responses that finally reward persistence in rechecking e-mail or Facebook, variable-interval schedules tend to produce slow, steady responding. This makes sense, because there is no knowing when the waiting will be over (TABLE 27.2).

In general, response rates are higher when reinforcement is linked to the number of responses (a ratio schedule) rather than to time (an interval schedule). But responding is more consistent when reinforcement is unpredictable (a variable schedule) than when it is predictable (a fixed schedule). Animal behaviors differ, yet Skinner (1956) contended

Table 27.2 Schedules of Reinforcement

	Fixed	Variable
<i>Ratio</i>	Every so many: reinforcement after every <i>n</i> th behavior, such as buy 10 coffees, get 1 free, or pay per product unit produced	After an unpredictable number: reinforcement after a random number of behaviors, as when playing slot machines or fly casting
<i>Interval</i>	Every so often: reinforcement for behavior after a fixed time, such as Tuesday discount prices	Unpredictably often: reinforcement for behavior after a random amount of time, as in checking for a Facebook response

AP® Exam Tip

The word "interval" in schedules of reinforcement means that an interval of time must pass before reinforcement. There is nothing the learner can do to shorten the interval. The word "ratio" refers to the ratio of responses to reinforcements. If the learner responds with greater frequency, there will be more reinforcements.

variable-ratio schedule
in operant conditioning, a reinforcement schedule that reinforces a response after an unpredictable number of responses.

fixed-interval schedule
in operant conditioning, a reinforcement schedule that reinforces a response only after a specified time has elapsed.

variable-interval schedule
in operant conditioning, a reinforcement schedule that reinforces a response at unpredictable time intervals.

that the reinforcement principles of operant conditioning are universal. It matters little, he said, what response, what reinforcer, or what species you use. The effect of a given reinforcement schedule is pretty much the same: "Pigeon, rat, monkey, which is which? It doesn't matter. . . Behavior shows astonishingly similar properties."

Punishment

27.4 How does punishment differ from negative reinforcement, and how does punishment affect behavior?

Reinforcement increases a behavior; **punishment** does the opposite. A *punisher* is any consequence that *decreases* the frequency of a preceding behavior (TABLE 27.3). Swift and sure punishers can powerfully restrain unwanted behavior. The rat that is shocked after touching a forbidden object and the child who is burned by touching a hot stove will learn not to repeat those behaviors. A dog that has learned to come running at the sound of an electric can opener will stop coming if its owner runs the machine to attract the dog and banish it to the basement.

Table 27.3 Ways to Decrease Behavior

Type of Punisher	Description	Examples
<i>Positive punishment</i>	Administer an aversive stimulus	Spray water on a barking dog; give a traffic ticket for speeding
<i>Negative punishment</i>	Withdraw a rewarding stimulus	Take away a teen's driving privileges; revoke a library card for nonpayment of fines

Criminal behavior, much of it impulsive, is also influenced more by swift and sure punishers than by the threat of severe sentences (Darley & Alter, 2011). Thus, when Arizona introduced an exceptionally harsh sentence for first-time drunk drivers, the drunk-driving rate changed very little. But when Kansas City police started patrolling a high crime area to increase the sureness and swiftness of punishment, that city's crime rate dropped dramatically.

How should we interpret the punishment studies in relation to parenting practices? Many psychologists and supporters of nonviolent parenting note four major drawbacks of physical punishment (Gershoff, 2002; Marshall, 2002).

1. *Punished behavior is suppressed, not forgotten.* This temporary state may (negatively) reinforce parents' punishing behavior. The child swears, the parent swats, the parent hears no more swearing and feels the punishment successfully stopped the behavior. No wonder spanking is a hit with so many U.S. parents of 3- and 4-year-olds—more than 9 in 10 of whom acknowledged spanking their children (Kazdin & Benjet, 2003).
2. *Punishment teaches discrimination among situations.* In operant conditioning, *discrimination* occurs when an organism learns that certain responses, but not others, will be reinforced. Did the punishment effectively end the child's swearing? Or did the child simply learn that it's not okay to swear around the house, though okay elsewhere?
3. *Punishment can teach fear.* In operant conditioning, *generalization* occurs when an organism's response to similar stimuli is also reinforced. A punished child may associate fear not only with the undesirable behavior but also with the person who delivered the punishment or the place it occurred. Thus, children may learn to fear a punishing teacher and try to avoid school, or may become more anxious (Gershoff et al., 2010). For such reasons, most European countries and most U.S. states now ban

Try This


Telemarketers are reinforced by which schedule? People checking the oven to see if the cookies are done are on which schedule? Airline frequent-flyer programs that offer a free flight after every 25,000 miles of travel use which reinforcement schedule?

ANSWERS: Telemarketers are reinforced on a variable-ratio schedule (after varying numbers of rings). Cookie checkers are reinforced on a fixed-interval schedule. Frequent-flyer programs use a fixed-ratio schedule.

AP® Exam Tip

Remember that *any kind of reinforcement* (positive, negative, primary, conditioned, immediate, delayed, continuous, or partial) encourages the behavior. *Any kind of punishment* discourages the behavior. Positive and negative do not refer to values—it's not that positive reinforcement (or punishment) is the good kind and negative is the bad. Think of positive and negative mathematically; a stimulus is added with positive reinforcement (or punishment) and a stimulus is subtracted with negative reinforcement (or punishment).

punishment an event that tends to decrease the behavior that it follows.



Children see, children do? Children who often experience physical punishment tend to display more aggression.

- hitting children in schools and child-care institutions (www.stophitting.com). Thirty-three countries, including those in Scandinavia, further outlaw hitting by parents, providing children the same legal protection given to spouses.
4. *Physical punishment may increase aggression by modeling aggression as a way to cope with problems.* Studies find that spanked children are at increased risk for aggression (and depression and low self-esteem). We know, for example, that many aggressive delinquents and abusive parents come from abusive families (Straus & Gelles, 1980; Straus et al., 1997).

Some researchers note a problem. Well, *yes*, they say, physically punished children may be more aggressive, for the same reason that people who have undergone psychotherapy are more likely to suffer depression—because they had preexisting problems that triggered the treatments (Larzelere, 2000, 2004). Which is the chicken and which is the egg? Correlations don't hand us an answer.

If one adjusts for preexisting antisocial behavior, then an occasional single swat or two to misbehaving 2- to 6-year-olds looks more effective (Baumrind et al., 2002; Larzelere & Kuhn, 2005). That is especially so if two other conditions are met:

1. The swat is used only as a backup when milder disciplinary tactics, such as a time-out (removing them from reinforcing surroundings), fail.
2. The swat is combined with a generous dose of reasoning and reinforcing.

Other researchers remain unconvinced. After controlling for prior misbehavior, they report that more frequent spankings of young children predict future aggressiveness (Grogan-Kaylor, 2004; Taylor et al., 2010).

Parents of delinquent youths are often unaware of how to achieve desirable behaviors without screaming at or hitting their children (Patterson et al., 1982). Training programs can help transform dire threats ("Apologize right now or I'm taking that cell phone away!") into positive incentives ("You're welcome to have your phone back when you apologize."). Stop and think about it. Aren't many threats of punishment just as forceful, and perhaps more effective, when rephrased positively? Thus, "If you don't get your homework done, I'm not giving you money for a movie!" would better be phrased as . . .

In classrooms, too, teachers can give feedback on papers by saying, "No, but try this . . ." and "Yes, that's it!" Such responses reduce unwanted behavior while reinforcing more desirable alternatives. Remember: *Punishment tells you what not to do; reinforcement tells you what to do.*

What punishment often teaches, said Skinner, is how to avoid it. Most psychologists now favor an emphasis on reinforcement.

Skinner's Legacy

27-5 Why did Skinner's ideas provoke controversy?

B. F. Skinner stirred a hornet's nest with his outspoken beliefs. He repeatedly insisted that external influences (not internal thoughts and feelings) shape behavior. And he urged people to use operant principles to influence others' behavior at school, work, and home. Knowing that behavior is shaped by its results, he said we should use rewards to evoke more desirable behavior.

Skinner's critics objected, saying that he dehumanized people by neglecting their personal freedom and by seeking to control their actions. Skinner's reply: External consequences already haphazardly control people's behavior. Why not administer those consequences toward human betterment? Wouldn't reinforcers be more humane than the punishments used in homes, schools, and prisons? And if it is humbling to think that our history has shaped us, doesn't this very idea also give us hope that we can shape our future?

Before You Move On

▶ ASK YOURSELF

Does your social media behavior (such as checking for new messages) make sense now that you've learned about the different kinds of reinforcement schedules?

▶ TEST YOURSELF

Fill in the three blanks below with one of the following terms: negative reinforcement (NR), positive punishment (PP), and negative punishment (NP). The first answer, positive reinforcement (PR) is provided for you.

Type of Stimulus	Give It	Take It Away
Desired (for example, a teen's use of the car):	1. PR	2.
Undesired/aversive (for example, an insult):	3.	4.

Answers to the Test Yourself questions can be found in Appendix E at the end of the book.

Module 27 Review

27-1 What is operant conditioning, and how is operant behavior reinforced and shaped?

- In *operant conditioning*, behaviors followed by reinforcers increase; those followed by punishers often decrease.
- Expanding on Edward Thorndike's *law of effect*, B. F. Skinner and others found that the behavior of rats or pigeons placed in an *operant chamber* (Skinner box) can be *shaped* by using reinforcers to guide closer and closer approximations of the desired behavior.

27-2 How do positive and negative reinforcement differ, and what are the basic types of reinforcers?

- *Reinforcement* is any consequence that strengthens behavior. *Positive reinforcement* adds a desirable stimulus to increase the frequency of a behavior. *Negative reinforcement* removes an aversive stimulus to increase the frequency of a behavior.
- *Primary reinforcers* (such as receiving food when hungry or having nausea end during an illness) are innately satisfying—no learning is required.
- *Conditioned (or secondary) reinforcers* (such as cash) are satisfying because we have learned to associate them with more basic rewards (such as the food or medicine we buy with them).
- Immediate reinforcers (such as a purchased treat) offer immediate payback; delayed reinforcers (such as a weekly paycheck) require the ability to delay gratification.

B. F. Skinner "I am sometimes asked, 'Do you think of yourself as you think of the organisms you study?' The answer is yes. So far as I know, my behavior at any given moment has been nothing more than the product of my genetic endowment, my personal history, and the current setting" (1983).

