6

CHAPTER SIX

Learning

ADAM AND TERI'S DAUGHTER, CARLY, celebrated her first birthday on September 11, 2001. The family happened to be living in Boston at the time, and they awoke that morning full of anticipation for a fun-filled day of birthday celebration.

What they got instead was a phone call from a friend in Texas, urging them to turn on the local news. Like many Americans, Adam and Teri watched with sadness and horror as terrorist attacks in New York, Pennsylvania, and the nation's capital took place before their eyes. American Airlines Flight 11, which crashed into the North Tower of the World Trade Center, had originated from Boston that morning, heightening the sense of uncertainty and anxiety that already had begun to define the day. Adam and Teri watched in shock as United Airlines Flight 175 crashed into the South Tower on live television. As the news reports filtered in throughout the day, each more disturbing than the last, the couple could scarcely avert their eyes from the television, and they ended up having CNN on all day long.

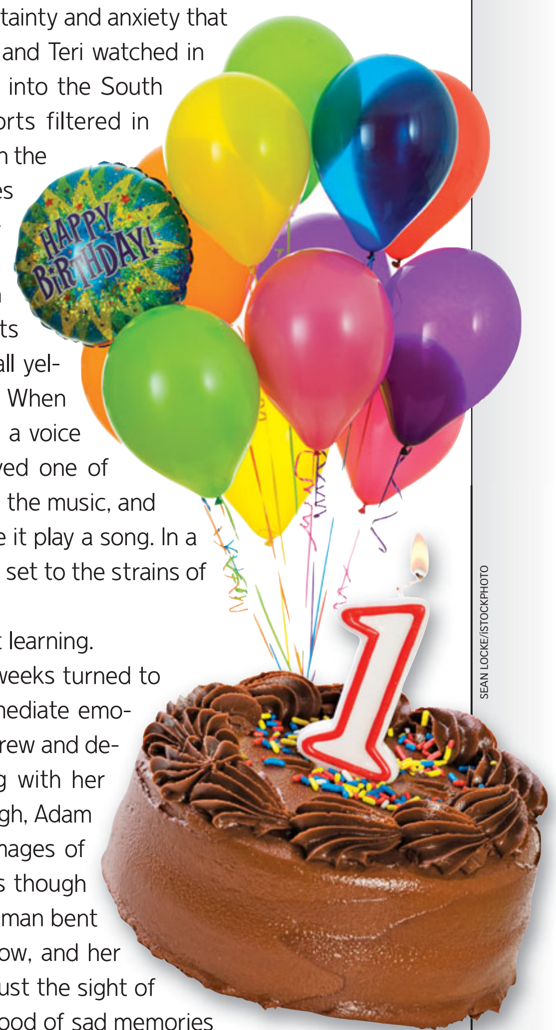

SEAN LOCKE/ISTOCKPHOTO

Yet through it all, young Carly played with her presents, blissfully unaware of the events unfolding on the TV screen. One gift, a small yellow soccer goal, turned out to be a favorite. When the ball hit the back of the net, it triggered a voice that yelled, "Goooooaaaallll!" and then played one of several songs at random. Carly loved to hear the music, and she would repeatedly whack the toy to make it play a song. In a surreal scene, fire, turmoil, and carnage were set to the strains of "John Jacob Jingleheimer Schmidt."

And that's what makes this a story about learning.

Quite a curious thing happened. As the weeks turned to months and 2001 turned to 2002, the immediate emotional impact of 9/11 faded for Adam. Carly grew and developed, and she continued to love playing with her soccer goal. Each time it played a song, though, Adam felt a chill run through his body and saw images of burning buildings in his mind's eye. It was as though John Jacob Jingleheimer Schmidt was a madman bent on bedeviling his life. Carly is much older now, and her baby toys have been put up on a shelf. But just the sight of that little yellow goal can still bring back a flood of sad memories

AL SWINKELS

AP PHOTO/CARMEN TAYLOR

• As baby Carly played with her new soccer goal, television images showed the horrifying events of September 11, 2001. Carly's parents, Adam and Teri, learned an association that lasted for years between the baby's toy and the 9/11 events.

and call up a welter of unpleasant emotions for her parents.

What's at work here is a type of learning based on association. Adam and Teri came to associate a unique historical tragedy and a child's toy, and as a result, either of the two stimuli produced certain mental and emotional reactions. The fear and sadness that were triggered by watching the events of 9/11 came to be triggered by an innocuous plaything, and it was an effect that lasted for years. In this chapter, we'll consider this type of learning as well as other ways that knowledge is acquired and stored. ■

Defining Learning: Experience That Causes a Permanent Change

Learning is shorthand for a collection of different techniques, procedures, and outcomes that produce changes in an organism's behavior. Learning psychologists have identified and studied as many as 40 different kinds of learning. However, there is a basic principle at the core of all of them. **Learning** involves *some experience that results in a relatively permanent change in the state of the learner.* This definition emphasizes several key ideas: Learning is based on experience; learning produces changes in the organism; and these changes are relatively permanent. Think back to Adam and Teri's experiences on September 11, 2001—seeing the horrors of 9/11 unfold on their TV screen and hearing Carly's toy changed their response to what had been a harmless child's toy. Furthermore, the association they learned lasted for years.

Learning can be conscious and deliberate or unconscious. For example, memorizing the names of all the U.S. presidents is a conscious and deliberate activity, with an explicit awareness of the learning process as it is taking place. In comparison, the kind of learning that associated Carly's toy with images of horror is much more implicit. Adam and Teri certainly weren't aware of or consciously focused on learning as it was taking place. Some other forms of learning start out explicitly but become more implicit over time. When you first learned to drive a car, for example, you probably devoted a lot of attention to the many movements and sequences that needed to be carried out simultaneously ("Step lightly on the accelerator while you push the turn indicator, and look in the rearview mirror while you turn the steering wheel"). That complex interplay of motions is now probably quite effortless and automatic for you. Explicit learning has become implicit over time.

● How are learning and memory linked?

These distinctions in learning might remind you of similar distinctions in memory and for good reason. In Chapter 5, you read about the differences between *implicit* and *explicit* memories as well as *procedural, semantic,* and *episodic* memories. Do different forms of learning mirror different types of memory? It's not that simple, but it is true that learning and memory are inextricably linked. Learning produces memories; and conversely, the existence of memories implies that knowledge was acquired, that experience was registered and recorded in the brain, or that learning has taken place.

The Case of Habituation

If you've ever lived under the flight path of your local airport, near railroad tracks, or by a busy highway, you probably noticed the loud noises when you first moved in. You probably also noticed that after a while, the roar wasn't quite so deafening anymore and

learning Some experience that results in a relatively permanent change in the state of the learner.

habituation A general process in which repeated or prolonged exposure to a stimulus results in a gradual reduction in responding.

that eventually you ignored the sounds of the planes, trains, or automobiles in your vicinity.

Habituation is *a general process in which repeated or prolonged exposure to a stimulus results in a gradual reduction in responding*. For example, a car that backfires unexpectedly as you walk by will produce a startle response: You'll jump back; your eyes will widen; your muscles will tense; and your body will experience an increase in sweating, blood pressure, and alertness. If another car were to backfire a block later, you might show another startle response, but it would be less dramatic and subside more quickly. If a third backfire should occur, you would likely not respond at all. You would have become *habituated* to the sound of a car backfiring.

Why won't the noise from a highway near your home keep you awake at night?

Habituation is a simple form of learning. An experience results in a change in the state of the learner: In the preceding example, you begin by reacting one way to a stimulus and, with experience, your reactions change. However, this kind of change usually isn't permanent. In most cases of habituation, a person will exhibit the original reaction if enough time has gone by. Thus, when you return home from a 2-week vacation, the roar of the jets passing over your home will probably sound just as loud as ever.

Learning and Behaviorism

As you'll recall from Chapter 1, a sizable chunk of psychology's history was devoted to a single dominant viewpoint. Behaviorism, with its insistence on measuring only observable, quantifiable behavior and its dismissal of mental activity as irrelevant and unknowable, was the major outlook of most psychologists working from the 1930s through the 1950s. This was also the period during which most of the fundamental work on learning theory took place. Most behaviorists argued that the "permanent change in experience" that resulted from learning could be demonstrated equally well in almost any organism: rats, dogs, pigeons, mice, pigs, or humans. From this perspective, behaviorists viewed learning as a purely behavioral, eminently observable activity that did not necessitate any mental activity.

As you'll see shortly, in many ways the behaviorists were right. Much of what we know about how organisms learn comes directly from the behaviorists' observations of behaviors. However, the behaviorists also overstated their case. Some important cognitive considerations—that is, elements of mental activity—need to be addressed in order to understand the learning process.

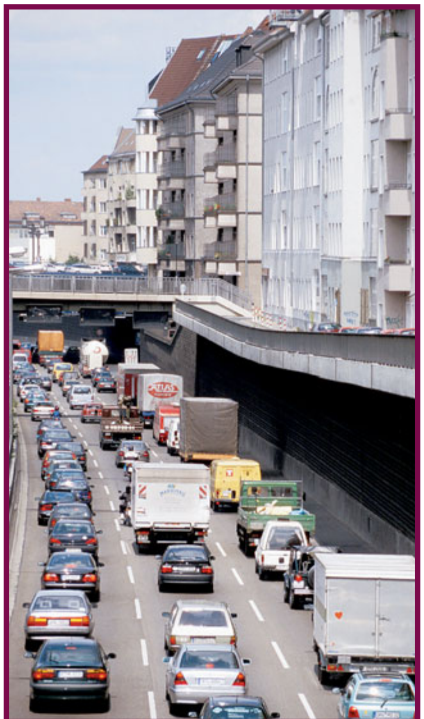

MICHAEL KLINEC/ALAMY

Living near a busy highway can be unpleasant. Most people who live near major highways become habituated to the sound of traffic.

summary quiz [6.1]

1. ________ is defined as an experience that results in a relatively permanent change in an organism's behavior.
   a. Behaviorism c. Habituation
   b. Learning d. Acquisition

2. Andi lives near the flight path of a large airport. At first, she was keenly aware of the loud roar of jets on the runway, but eventually, she no longer noticed the sound. This illustrates
   a. habituation. c. operant learning.
   b. classical conditioning. d. biological preparedness.

3. Most behaviorists in the mid-20th century argued that
   a. elements of cognitive activity are involved in human learning.
   b. habituation is not necessary for learning to occur.
   c. all learning is an observable activity.
   d. nonhuman animals learn in a fundamentally different way than humans.

Classical Conditioning: One Thing Leads to Another

You'll recall from Chapter 1 that the early behaviorists were greatly influenced by the work of Russian physiologist Ivan Pavlov, who had revealed the mechanics of one form of learning, which came to be called classical conditioning. **Classical conditioning** occurs *when a neutral stimulus evokes a response after being paired with a stimulus that naturally evokes a response*. In his classic experiments, Pavlov showed that dogs learned to salivate to neutral stimuli such as a bell or a tone after that stimulus had been associated with another stimulus that naturally evokes salivation, such as food.

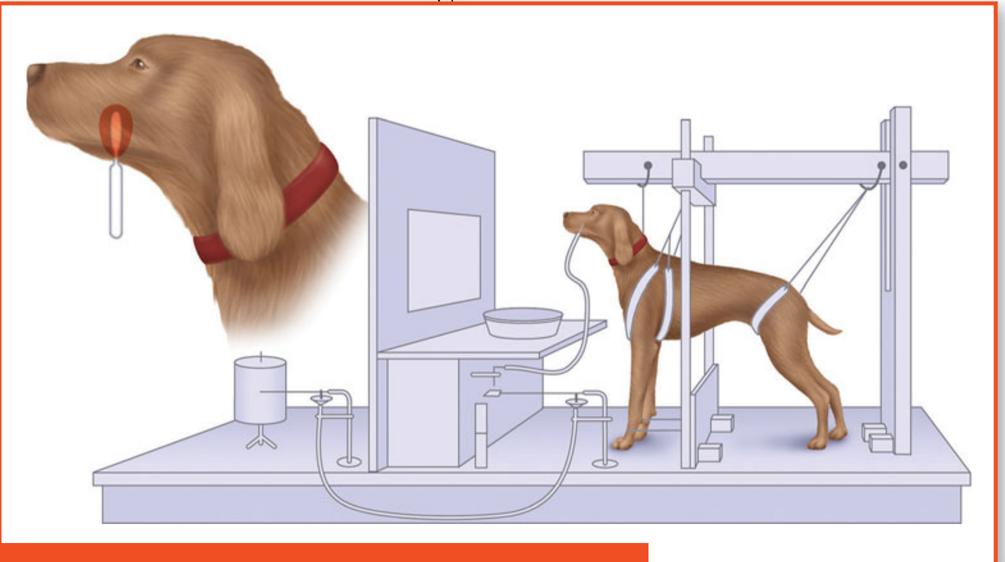

FIGURE 6.1
Pavlov's Apparatus for Studying Classical Conditioning
Pavlov presented auditory stimuli to the animals using a bell or a tuning fork. Visual stimuli could be presented on the screen. The inset shows a close-up of the tube inserted in the dog's salivary gland for collecting saliva.

Pavlov's Experiments on Classical Conditioning

Pavlov's basic experimental setup involved cradling dogs in a harness to administer various kinds foods and to measure the salivary response to each, as shown in FIGURE 6.1. He noticed that dogs that previously had been in the experiment began to produce a kind of "anticipatory" salivary response as soon as they were put in the harness, before any food was presented. Pavlov and his colleagues regarded these responses as annoyances at first because they interfered with collecting naturally occurring salivary secretions. In reality, the dogs were exhibiting classical conditioning.

When the dogs were initially presented with a plate of food, they began to salivate. No surprise here—placing food in front of most animals will launch the salivary process. Pavlov called the presentation of food an **unconditioned stimulus (US)**, or *something that reliably produces a naturally occurring reaction in an organism*. He called the dogs' salivation an **unconditioned response (UR)**, or *a reflexive reaction that is reliably elicited by an unconditioned stimulus*.

Pavlov soon discovered that he could make the dogs salivate to stimuli that don't usually make animals salivate. In various experiments, Pavlov paired the presentation of food with the sound of a buzzer, the ticking of a metronome, the humming of a tuning fork, or the flash of a light (Pavlov, 1927). Sure enough, he found that the dogs salivated to the sound of a buzzer, the ticking of a metronome, the humming of a tuning fork, or the flash of a light, each of which had become a **conditioned stimulus (CS)**, or *a stimulus that is initially neutral and produces no reliable response in an organism* (see FIGURE 6.2 on page 165). When dogs hear the sound of a buzzer in the wild, they're not known to salivate. However, when the buzzer (CS) is paired over time with the food (US), the animal will learn to associate food with the sound, and eventually the CS is sufficient to produce a response, or salivation. This response resembles the UR, but Pavlov called it the **conditioned response (CR)**, or *a reaction that resembles an unconditioned response but is produced by a conditioned stimulus*. As you can imagine, a range of stimuli might be used as a CS, and as we noted earlier, several different stimuli became the CS in Pavlov's experiment.

● Why do some dogs seem to know when it's dinner time?

Let's apply these four basic elements of the classical conditioning process—the US, UR, CS, and CR—to a real-world example. Consider your own dog (or cat). You probably

classical conditioning When a neutral stimulus evokes a response after being paired with a stimulus that naturally evokes a response.

unconditioned stimulus (US) Something that reliably produces a naturally occurring reaction in an organism.

unconditioned response (UR) A reflexive reaction that is reliably elicited by an unconditioned stimulus.

conditioned stimulus (CS) A stimulus that is initially neutral and produces no reliable response in an organism.

conditioned response (CR) A reaction that resembles an unconditioned response but is produced by a conditioned stimulus.

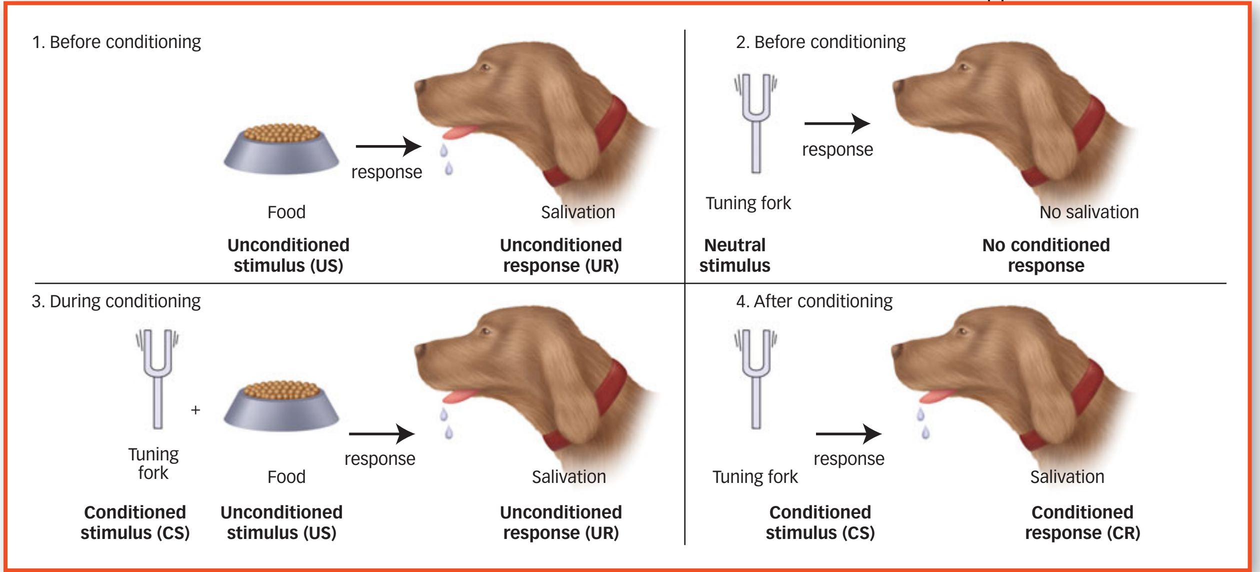

FIGURE 6.2 **The Elements of Classical Conditioning** *In classical conditioning, a previously neutral stimulus (e.g., the sound of a tuning fork) is paired with an unconditioned stimulus (e.g., the presentation of food). After several trials associating the two, the conditioned stimulus (the sound) alone can produce a conditioned response*

think you have the only dog that can tell time because she always knows when dinner's coming and gets ready to eat. Sorry to burst your bubble, but your dog is no clock-watching wonder hound. Instead, the presentation of food (the US) has become associated with a complex CS—your getting up, moving into the kitchen, opening the cabinet, working the can opener—such that the CS alone signals to your dog that food is on the way and therefore initiates the CR of her getting ready to eat.

As another example, think back to Adam and Teri's experiences on September 11, 2001. As they watched the World Trade Center collapsing on television, they felt sadness, fear, and anxiety. The images of devastation and horror were the US, and the negative feelings those images caused were the UR. However, Carly's soccer goal acted as the CS. The toy—and especially the songs it played—was an initially neutral stimulus that was associated with the US that day. As the horrific images flashed across the screen, "John Jacob Jingleheimer Schmidt" provided an endless sound track. Eventually the CS all by itself—the music played by the toy—was sufficient to produce the CR: feelings of sadness, fear, and anxiety.

When Pavlov's findings first appeared in the scientific and popular literature (Pavlov, 1923a, 1923b), they produced a flurry of excitement because psychologists now had demonstrable evidence of how conditioning produced learned behaviors. This was the kind of psychology that Watson and the behaviorists were proposing: An organism experiences events or stimuli that are observable and measurable, and changes in that organism can be directly observed and measured. Dogs learned to salivate to the sound of a buzzer, and there was no need to resort to explanations about why it had happened, what the dog wanted, or how the animal thought about the situation. In other words, there was no need to consider the mind in this classical conditioning paradigm, which appealed to Watson and the behaviorists. Pavlov also appreciated the significance of his discovery and embarked on a systematic investigation of the mechanisms of classical conditioning. (The Real World box on the next page shows how Pavlov's ideas help explain how drug overdoses occur.)

"I THINK MOM'S USING THE CAN OPENER."

[THE REAL WORLD]

Understanding Drug Overdoses

All too often, police are confronted with a perplexing problem: the sudden death of addicts from a drug overdose. These deaths are puzzling for at least three reasons: The victims are often experienced drug users, the dose taken is usually not larger than what they usually take, and the deaths tend to occur in unusual settings. Experienced drug users are just that: experienced! You'd think that if a heroin addict or crack cocaine user were ingesting a typical amount of a substance he or she had used many times before, the chances of an overdose would be *lower* than usual.

Classical conditioning provides some insight into how these deaths occur. First, when classical conditioning takes place, the CS is more than a simple bell or tone: It also includes the overall *context* within which the conditioning takes place. Indeed, Pavlov's dogs often began to salivate even as they approached the experimental apparatus. Second, many CRs are compensatory reactions to the US. In some of Pavlov's early experiments, he used a very mild acid solution as the US because it produces large amounts of saliva that dilute the acid in the dog's mouth. When that salivary response is eventually conditioned to the sound of a tone, in a way it represents the remnants of the body's natural reaction to the presentation of the US.

These two finer points of classical conditioning help explain what happens when someone takes a drug such as heroin (Siegel, 1984). When the drug is injected, the entire setting (the drug paraphernalia, the room, the lighting, the addict's usual companions) functions as the CS, and the addict's brain reacts to the heroin by secreting neurotransmitters that counteract its effects. Over time, this protective physiological response becomes part of the CR, and like all CRs, it occurs in the presence of the CS but prior to the actual administration of the drug. These compensatory physiological reactions are also what make drug abusers take increasingly larger doses to achieve the same effect; ultimately, these reactions produce *drug tolerance*, discussed in Chapter 8.

Based on these principles of classical conditioning, taking drugs in a new environment can be fatal for a longtime drug user. If an addict injects the usual dose in a setting that is sufficiently novel or where heroin has never been taken before, the CS is now altered. What's more, the physiological compensatory CR either does not occur or is substantially decreased. As a result, the addict's usual dose becomes an overdose, and death often results. This effect has also been shown experimentally: Rats that have had extensive experience with morphine in one setting were much more likely to survive dose increases in that same setting than in a novel one (Siegel, 1976).

The basic principles of classical conditioning help explain this real-world tragedy of drug overdose. Intuitively, addicts may stick with the crack houses, opium dens, or "shooting galleries" with which they're familiar for just this reason.

AP PHOTO/CHRIS GARDNER

Although opium dens and crack houses may be considered blight, it is often safer for addicts to use drugs there. The environment becomes part of the addict's CS, so ironically, busting crack houses may contribute to more deaths from drug overdose when addicts are pushed to use drugs in new situations.

acquisition The phase of classical conditioning when the CS and the US are presented together.

extinction The gradual elimination of a learned response that occurs when the US is no longer presented.

spontaneous recovery The tendency of a learned behavior to recover from extinction after a rest period.

generalization A process in which the CR is observed even though the CS is slightly different from the original one used during acquisition.

The Basic Principles of Classical Conditioning

Classical conditioning requires some period of association between the CS and US. This period is called **acquisition**, or *the phase of classical conditioning when the CS and the US are presented together.* During the initial phase of classical conditioning, typically there is a gradual increase in learning: It starts low, rises rapidly, and then slowly tapers off, as shown on the left side of **FIGURE 6.3** (on page 167). Pavlov's dogs gradually increased their amount of salivation over several trials of pairing a tone with the presentation of food, and similarly, your dog eventually learned to associate your kitchen preparations with the subsequent appearance of food. After learning has been established, the CS by itself will reliably elicit the CR.

After Pavlov and his colleagues had explored the process of acquisition extensively, they turned to the next logical question: What would happen if they continued to present

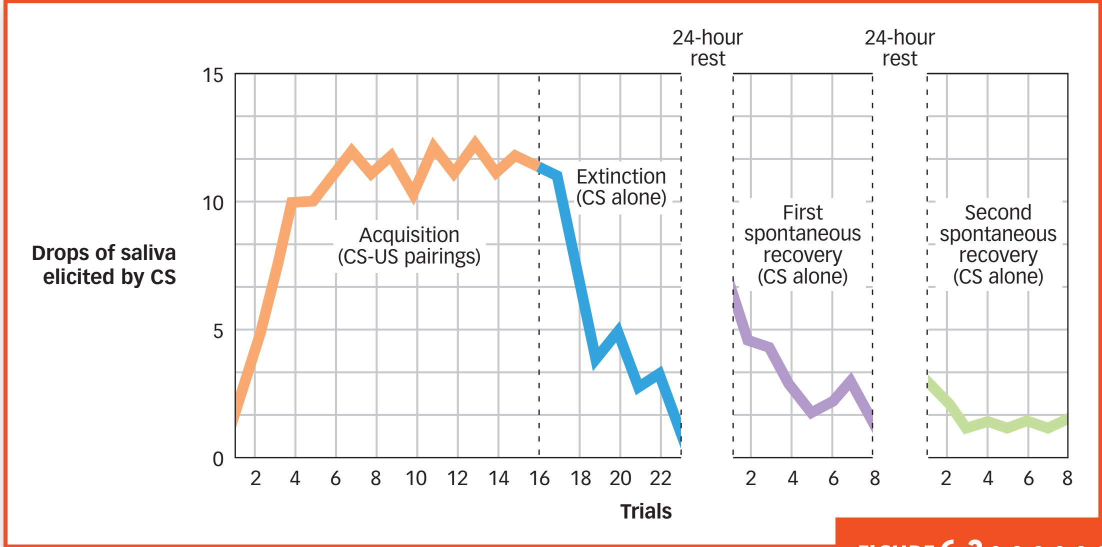

FIGURE 6.3 **Acquisition, Extinction, and Spontaneous Recovery** *In classical conditioning, the CS is originally neutral and produces no specific response. After several trials pairing the CS with the US, the CS alone comes to elicit the salivary response (the CR). Learning tends to take place fairly rapidly and then levels off as stable responding develops. In extinction, the CR diminishes quickly until it no longer occurs. A rest period, however, is typically followed by spontaneous recovery of the CR. In fact, a well-learned CR may show spontaneous recovery after more than one rest.*

the CS (buzzer) but stopped presenting the US (food)? The result is just as you might imagine: As shown on the right side of the first panel in FIGURE 6.3 (above), behavior declines abruptly and continues to drop until eventually the dog ceases to salivate to the sound of the buzzer. This process is called **extinction**, *the gradual elimination of a learned response that occurs when the CS is presented but no longer paired with the US.* Similarly, if you make noises in the kitchen without subsequently presenting a meaty plate of Alpo, eventually your dog will stop salivating or even getting aroused every time you walk into the kitchen.

Having established that he could produce learning through conditioning and then extinguish it, Pavlov wondered if this elimination of conditioned behavior was permanent. Is a single session of extinction sufficient to knock out the CR completely, or is there some residual change in the dog's behavior so that the CR might reappear?

● How does conditioned behavior change when the unconditioned stimulus is removed?

To explore this question, Pavlov extinguished the classically conditioned salivation response and then allowed the dogs to have a short rest period. When they were brought back to the lab and presented with the CS again, they displayed **spontaneous recovery**, *the tendency of a learned behavior to recover from extinction after a rest period.* This phenomenon is shown in the middle panel in FIGURE 6.3. Notice that this recovery takes place even though there have not been any additional associations between the CS and US. Some spontaneous recovery of the conditioned response even takes place in what is essentially a second extinction session after another period of rest (see the right-hand panel in FIGURE 6.3). Clearly, extinction had not completely wiped out the learning that had been acquired. The ability of the CS to elicit the CR was weakened, but it was not eliminated. In fact, if the CS-US pairings are introduced again, the animal will show rapid conditioning, much more rapid than during the initial acquisition phase. This effect is known as *savings*, since it suggests that some underlying neural changes that occurred during the initial learning are "saved" no matter how many extinction trials are conducted, and it is a good illustration of the permanence of some kinds of learning.

Another important principle governing classical conditioning is **generalization**, in which *the CR is observed even though the CS is slightly different from the original one used during acquisition.* Suppose you decide to break down and buy a new can opener, replacing the crummy one that you've had for years. Let's say the new one makes a slightly different sound. Do you think your dog will be stumped,

RON KIMBALL/KIMBALLSTOCK

unable to anticipate the presentation of her food? Will a whole new round of conditioning need to be established with this modified CS?

Probably not. It wouldn't be very adaptive for an organism if each little change in the CS-US pairing required an extensive regimen of new learning. Rather, your dog's conditioning to the sound of the old can opener will probably generalize to the sound of the new one. As you might expect, the more the new stimulus changes, the less conditioned responding is observed. If you replaced a hand-held can opener with an electric can opener, your dog would probably show a much weaker conditioned response (Pearce, 1987; Rescorla, 2006).

Some generalization studies used a 1,000-hertz (Hz) tone as the CS during the acquisition phase. The test stimuli used were tones of higher or lower pitches. As you might expect, an animal gives the maximum response to the original stimulus of 1,000 Hz, with a systematic drop-off as the pitch of the replacement stimulus is farther away from the original tone of 1,000 Hz regardless of whether the tone was higher or lower. Interestingly, when the stimulus is one of the octaves of the original stimulus (octaves in music are tones that are direct multiples of each other), either 500 Hz or 2,000 Hz, there is a slight increase in responding. In these cases, the rate of responding is lower than that of the original CS but higher than it is in other cases of dissimilar tones. The animals clearly show that they detect octaves just like we do, and in this case, responding has generalized to those octaves (see **FIGURE 6.4**, below).

● How can changing can openers affect a conditioned dog's response?

When an organism generalizes to a new stimulus, two things are happening. First, by responding to the new stimulus used during generalization testing, the organism demonstrates that it recognizes the similarity between the original CS and the new stimulus. Second, by displaying *diminished* responding to that new stimulus, it also tells us that it notices a difference between the two stimuli. In the second case, the organism shows **discrimination**, or *the capacity to distinguish between similar but distinct stimuli.*

Here's a true story about a talented golden retriever named Splash. Splash was very well trained to perform a number of behaviors when his name was called, as in "Go, Splash," to fetch a ball. The sound of his name was the CS, and running after a target was the US. Repeated attempts to trick him, by yelling, "Go, Splat!" or, "Go, Crash!" or even, "Go, Spla!" resulted in predictable outcomes. Splash would start to move, but then hesitate, showing that he discriminated between the appropriate stimulus ("Splash!") and the substituted ones ("Splat!").

Conceptually, generalization and discrimination are two sides of the same coin. The more organisms show one, the less they show the other, and training can modify the balance between the two.

TIMOTHY A. CLARY/AFP/GETTY IMAGES

● *Some people desire money as an end in itself to the extent that they will put others' welfare at risk in order to accumulate vast amounts of cash. Disgraced financier Bernard Madoff was sentenced to 150 years in prison after being convicted for defrauding hundreds of investors out of billions of dollars in order to create a vast fortune for himself. Individuals such as Madoff may be showing the effects of second-order conditioning.*

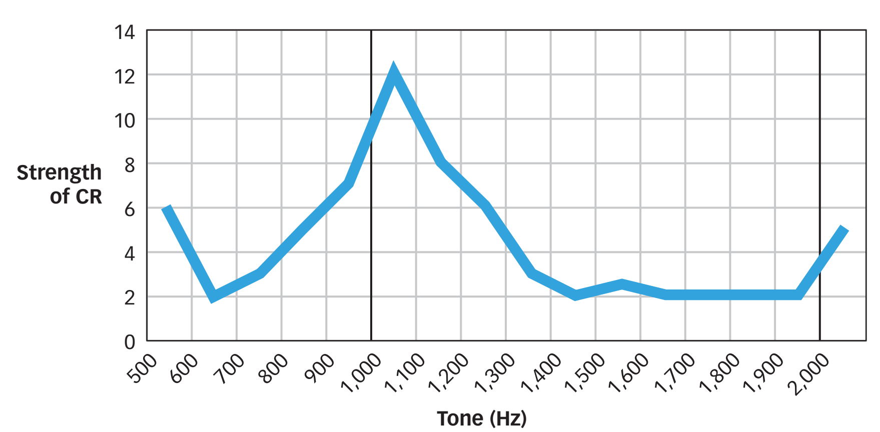

FIGURE 6.4 **Stimulus Generalization** *In this experiment, an animal was conditioned using a 1,000-Hz tone (the CS) and tested with a variety of tones of higher and lower pitches. As the pitches move farther away from the original CS, the strength of the CR drops off systematically. However, when the tone is an octave of the original (i.e., either 500 or 2,000 Hz), there is an increase in the CR.*

Conditioned Emotional Responses: The Case of Little Albert

WATSON & RAYNOR, 1920

John Watson and Rosalie Rayner show Little Albert an unusual bunny mask. Why doesn't the mere presence of these experimenters serve as a conditioned stimulus in itself?

Before you conclude that classical conditioning is merely a sophisticated way to train your dog, let's revisit the larger principles of Pavlov's work. Classical conditioning demonstrates that durable, substantial changes in behavior can be achieved simply by setting up the proper conditions. With the skillful association of a naturally occurring US with an appropriate CS, an organism can learn to perform a variety of behaviors, often after relatively few acquisition trials. There is no reference to an organism's *wanting* to learn the behavior, *willingness* to do it, *thinking* about the situation, or *reasoning* through the available options. We don't need to consider internal and cognitive explanations to demonstrate the effects of classical conditioning: The stimuli, the eliciting circumstances, and the resulting behavior are there to be observed by one and all.

It was this kind of simplicity that appealed to behaviorists. In fact, Watson and his followers thought that it was possible to develop general explanations of pretty much *any* behavior of *any* organism based on classical conditioning principles.

As a step in that direction, Watson embarked on a controversial study with his research assistant Rosalie Rayner (Watson & Rayner, 1920). To support his contention that even complex behaviors were the result of conditioning, Watson enlisted the assistance of 9-month-old "Little Albert." Watson presented Little Albert with a variety of stimuli: a white rat, a dog, a rabbit, various masks, and a burning newspaper. Albert's reactions in most cases were curiosity or indifference, and he showed no fear of any of the items. Then Watson unexpectedly struck a large steel bar with a hammer, producing a loud noise. Predictably, this caused Albert to cry, tremble, and be generally displeased.

Watson and Rayner then led Little Albert through the acquisition phase of classical conditioning. Albert was presented with a white rat. As soon as he reached out to touch it, the steel bar was struck. This pairing occurred again and again over several trials. Eventually, the sight of the rat alone caused Albert to recoil in terror, crying and clamoring to get away from it. In this situation, a US (the loud sound) was paired with a CS (the presence of the rat) such that the CS all by itself was sufficient to produce the CR (a fearful reaction). Little Albert also showed stimulus generalization. The sight of a white rabbit, a seal-fur coat, and a Santa Claus mask produced the same kinds of fear reactions in the infant.

Why did Albert fear the rat?

What was Watson's goal in all this? First, he wanted to show that a relatively complex reaction could be conditioned using Pavlovian techniques. Second, he wanted to show that emotional responses such as fear and anxiety could be produced by classical conditioning and therefore need not be the product of deeper unconscious processes or early life experiences as Freud and his followers had argued (see Chapter 1). Instead, Watson proposed that fears could be learned, just like any other behavior. Third, Watson wanted to confirm that conditioning could be applied to humans as well as to other animals. This study was controversial in its cavalier treatment of a young child, especially given that Watson and Rayner did not follow up with Albert or his mother during the ensuing years (Harris, 1979). Modern ethical guidelines that govern the treatment of research participants make sure that this kind of study could not be conducted today. At the time, however, it was consistent with a behaviorist view of psychology. As Watson (1930) summarized his position several years later:

> Give me a dozen healthy infants, well-formed, and my own specified world to bring them up in and I'll guarantee to take any one at random and train him to become any type of specialist I might select—doctor, lawyer, artist, merchant-chief and, yes, even beggar-man and thief, regardless of his talents, penchants, tendencies, abilities, vocations, and race of his ancestors. (p. 104)

discrimination The capacity to distinguish between similar but distinct stimuli.

Watson was promoting a staunch view that learning and the environment were responsible for determining behavior, more so than genetics or personality. He intended his statements to be extreme in order to shake up the young discipline of psychology and highlight the importance of acquired experiences in shaping behavior.

A Deeper Understanding of Classical Conditioning

As a form of learning, classical conditioning could be reliably produced, it had a simple set of principles, and it had applications to real-life situations. In short, classical conditioning offered a good deal of utility for psychologists who sought to understand the mechanisms underlying learning, and it continues to do so today.

Like a lot of strong starters, though, classical conditioning has been subjected to deeper scrutiny in order to understand exactly how, when, and why it works. Let's examine three areas that give us a closer look at the mechanisms of classical conditioning.

The Neural Elements of Classical Conditioning

Pavlov saw his research as providing insights into how the brain works. After all, he was trained in medicine, not psychology, and was a bit surprised when psychologists became excited by his findings. Recent research has clarified some of what Pavlov hoped to understand about conditioning and the brain.

For example, fear conditioning has been extensively studied in part because the brain substrates are particularly evident. In Chapter 3, you saw that the amygdala plays an important role in the experience of emotion, including fear and anxiety. So, it should come as no surprise that the amygdala, particularly an area known as the *central nucleus,* is also critical for emotional conditioning.

Consider a rat who is conditioned to a series of CS-US pairings where the CS is a tone and the US is a mild electric shock. When rats experience sudden painful stimuli in nature, they show a defensive reaction, known as *freezing,* where they crouch down and sit motionless. In addition, their autonomic nervous systems go to work: Heart rate and blood pressure increase, and various hormones associated with stress are released. When fear conditioning takes place, these two components—one behavioral and one physiological—occur, except that now they are elicited by the CS.

The central nucleus of the amygdala plays a role in producing both of these outcomes through two distinct connections with other parts of the brain. If connections linking the amygdala to the midbrain are disrupted, the rat does not exhibit the behavioral freezing response. If the connections between the amygdala and the hypothalamus are severed, the autonomic responses associated with fear cease (LeDoux et al., 1988). Hence, the action of the amygdala is an essential element in fear conditioning, and its links with other areas of the brain are responsible for producing specific features of conditioning. The amygdala is involved in fear conditioning in people as well as rats and other animals (Phelps & LeDoux, 2005).

● What is the role of the amygdala in fear conditioning?

The Cognitive Elements of Classical Conditioning

Pavlov's work was a behaviorist's dream come true. In this view, conditioning is something that *happens to* a dog, a rat, or a person, apart from what the organism thinks about the conditioning situation. However, eventually someone was bound to ask an important question: Why didn't Pavlov's dogs salivate to Pavlov? After all, he was instrumental in the arrival of the CS. If Pavlov delivered the food to the dogs, why didn't they form an association with him? Indeed, if Watson was present whenever the unpleasant US was sounded, why didn't Little Albert come to fear *him?*

Maybe classical conditioning isn't such an unthinking, mechanical process as behaviorists originally had assumed (Rescorla, 1966, 1988). Somehow, Pavlov's dogs were sensitive to the fact that Pavlov was not a *reliable* indicator of the arrival of food. Pavlov was linked with the arrival of food, but he was also linked with other activities that had nothing to do with food, including checking on the apparatus, bringing the dog from

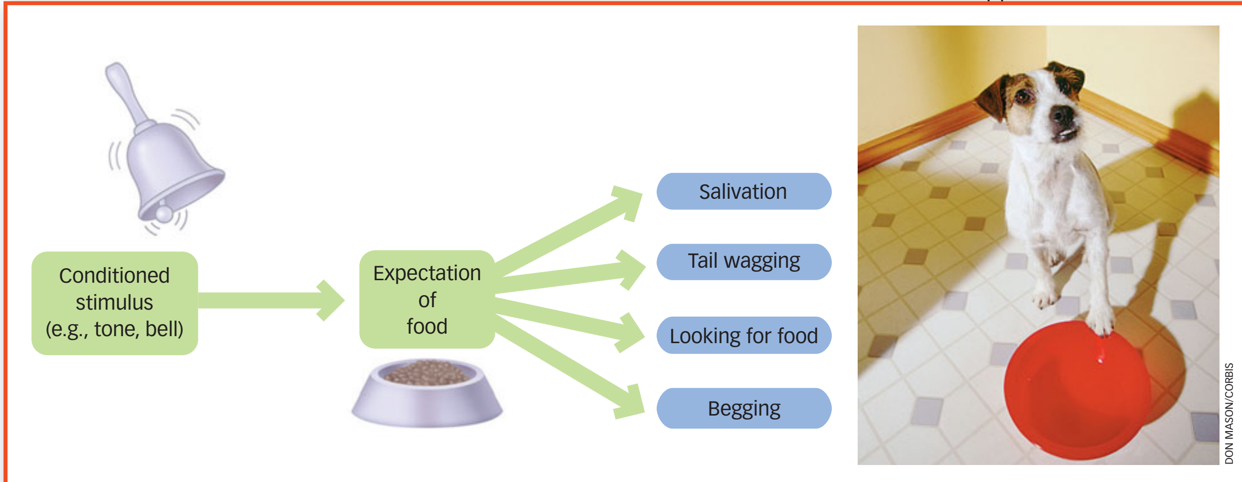

FIGURE 6.5 **Expectation in Classical Conditioning** *In the Rescorla-Wagner model of classical conditioning, a CS serves to set up an expectation. The expectation in turn leads to an array of behaviors associated with the presence of the CS.*

the kennel to the laboratory, and standing around and talking with his assistants. These observations suggest that perhaps cognitive components are involved in classical conditioning after all.

Robert Rescorla and Allan Wagner (1972) were the first to theorize that classical conditioning only occurs when an animal has learned to set up an *expectation* (see **FIGURE 6.5**). The sound of a tone, because of its systematic pairing with food, served to set up this cognitive state for the laboratory dogs; Pavlov, because of the lack of any reliable link with food, did not. Rescorla and Wagner predicted that conditioning would be easier when the CS was an *unfamiliar* event than when it was familiar. The reason is that familiar events, being familiar, already have expectations associated with them, making new conditioning difficult. For example, Adam didn't recoil in horror every time he saw his daughter Carly, even though she was present during the acquisition phase of 9/11. The familiarity of Carly in multiple contexts made her, thankfully, a poor CS for Adam's fear conditioning. In short, classical conditioning might appear to be a primitive and unthinking process, but it is actually quite sophisticated and incorporates a significant cognitive element.

How does familiarity with the stimulus hinder new conditioning?

The Evolutionary Elements of Classical Conditioning

In addition to this cognitive component, evolutionary mechanisms also play an important role in classical conditioning. As you learned in Chapter 1, evolution and natural selection go hand in hand with adaptiveness: Behaviors that are adaptive allow an organism to survive and thrive in its environment. In the case of classical conditioning, psychologists began to appreciate how this type of learning could have adaptive value. In fact, there are good reasons why animals may have evolved to condition to novel stimuli more easily than to familiar stimuli. An example comes from conditioning of food aversions and food preferences.

You may think food preference is a matter of personal taste, but in fact food aversions can be classically conditioned. A psychology professor was once on a job interview in Southern California, and his hosts took him to lunch at a Middle Eastern restaurant. Suffering from a case of bad hummus, he was up all night long. Needless to say, he was in pretty rough shape the following day, and he didn't get the job offer.

This colleague developed a lifelong aversion to hummus. Why would one bad incident taint food preferences in such a lasting way? On the face of it, this looks like a case of classical conditioning. The hummus was the CS, its apparent toxicity was the US, and the resulting gastric distress was the UR. The UR (the nausea) became linked to the once-neutral CS (the hummus) and became a CR (an aversion to hummus). However, this case has several unusual aspects.

Under certain conditions, people may develop food aversions. This serving of hummus looks inviting and probably tastes delicious, but at least one psychologist avoids it like the plague.

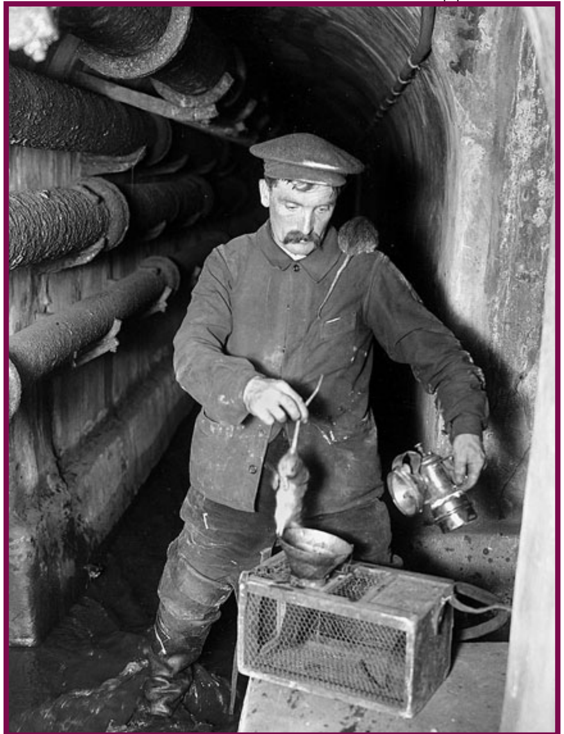

BOYER/ROGER VIOLLET/GETTY IMAGES

Rats can be difficult to poison because of learned taste aversions, which are an evolutionarily adaptive element of classical conditioning. Here a worker tries his best in the sewers of France.

For starters, all of the psychologist's hosts also ate the hummus, yet none of them reported feeling ill. It's not clear, then, what the US was; it couldn't have been anything that was actually in the food. What's more, the time between the hummus and the distress was several hours; usually a response follows a stimulus fairly quickly. Most baffling, this aversion was cemented with a single acquisition trial. Usually it takes several pairings of a CS and US to establish learning.

These peculiarities are not so peculiar from an evolutionary perspective. What seems like a mindbug is actually the manifestation of an adaptive process. Any species that forages or consumes a variety of foods needs to develop a mechanism by which it can learn to avoid any food that once made it ill. To have adaptive value, this mechanism should have several properties.

First, there should be rapid learning that occurs in perhaps one or two trials. If learning takes more trials than this, the animal could die from eating a toxic substance. Second, conditioning should be able to take place over very long intervals, perhaps up to several hours. Toxic substances often don't cause illness immediately, so the organism would need to form an association between food and the illness over a longer term. Third, the organism should develop the aversion to the smell or taste of the food rather than its ingestion. It's more adaptive to reject a potentially toxic substance based on smell alone than it is to ingest it. Finally, learned aversions should occur more often with novel foods than familiar ones. It is not adaptive for an animal to develop an aversion to everything it has eaten on the particular day it got sick. Our psychologist friend didn't develop an aversion to the Coke he drank with lunch or the scrambled eggs he had for breakfast that day, only to the unfamiliar hummus.

John Garcia and his colleagues illustrated the adaptiveness of classical conditioning in a series of studies with rats (Garcia & Koelling, 1966). They paired a variety of CSs with a US, such as injection of a toxic substance, that caused nausea and vomiting hours later. If the CS was water laced with a harmless but distinctly flavored novel substance (such as strawberry), the rats developed a strong aversion to the smell and taste of strawberries. But if the CS was a familiar food that the animal had eaten before, the aversion was much less likely to develop.

This research had an interesting application. It led to the development of a technique for dealing with an unanticipated side effect of radiation and chemotherapy: Cancer patients who experience nausea from their treatments often develop aversions to foods they ate before the therapy. Broberg and Bernstein (1987) reasoned that, if the findings with rats generalized to humans, a simple technique should minimize the negative consequences of this effect. They gave their patients an unusual food (coconut or root beer–flavored candy) at the end of the last meal before undergoing treatment. Sure enough, the conditioned food aversions that the patients developed were overwhelmingly for one of the unusual flavors and not for any of the other foods in the meal. As a result, patients were spared developing aversions to more common foods that they are more likely to eat. Understanding the basis of mindbugs can have practical as well as theoretical value.

● **How have cancer patients' discomfort been eased by our understanding of food aversions?**

Studies such as these suggest that evolution has provided each species with a kind of **biological preparedness**, *a propensity for learning particular kinds of associations over others,* so that some behaviors are relatively easy to condition in some species but not others. For example, the taste and smell stimuli that produce food aversions in rats do not work with most species of birds. Birds depend primarily on visual cues for finding food and are relatively insensitive to taste and smell. However, it is relatively easy to produce a food aversion in birds using an unfamiliar visual stimulus as the CS, such as a brightly colored food (Wilcoxon, Dragoin, & Kral, 1971). Indeed, most researchers agree that conditioning works best with stimuli that are biologically relevant to the organism (Domjan, 2005).

biological preparedness A propensity for learning particular kinds of associations over others.

Culture & Community

Is It Possible That Humans Have an Innate Ability to Understand Geometry?

In a study (Dhaene et al. 2006) of the Munduruku, an isolated indigenous tribe located in the Amazon, Munduruku children and adults were compared to American children and adults on their basic comprehension of geometric shapes. In each test, the participants identified which figure among the series of six images presented to them did *not* belong in the group. Each series tested basic geometric concepts like parallels, shapes, distance, and symmetry.

All participants performed well above the level of chance, and only American adults showed a significant advantage. Before this study, it was largely believed that people must "learn" geometry through cultural interventions like maps, mathematical tools, or the terms used in geometry. In contrast, this study provides evidence that core knowledge of geometry is a universal intuition of the human mind.

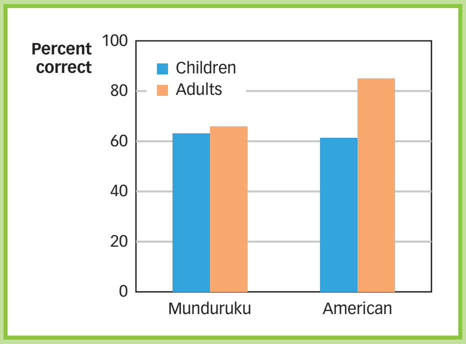

summary quiz [6.2]

4. When you make noise with your can opener in the kitchen, but no longer follow the noise with a bowl of dog food, your dog will stop running into the kitchen at the sound of the can opener. Your clever dog is displaying

a. acquisition.
b. generalization.
c. discrimination.
d. extinction.

5. Every time Little Albert reached out to touch the cute white rat, the experimenter startled Albert with a loud noise. Soon Little Albert feared not only the rat but also a white rabbit, a fur coat, and a Santa Claus mask. This phenomenon is called

a. extinction.
b. spontaneous recovery.
c. generalization.
d. discrimination.

6. The classical conditioning of fear involves which area of the brain?

a. amygdala
b. hypothalamus
c. hippocampus
d. temporal cortex

7. Jeff loved chili dogs, but one day he ate one and soon after became violently ill. After that, the mere sight or smell of a chili dog made Jeff nauseous. Jeff's nausea is an example of a(n)

a. conditioned stimulus.
b. conditioned response.
c. unconditioned stimulus.
d. unconditioned response.

operant conditioning A type of learning in which the consequences of an organism's behavior determine whether it will be repeated in the future.

law of effect The principle that behaviors that are followed by a "satisfying state of affairs" tend to be repeated and those that produce an "unpleasant state of affairs" are less likely to be repeated.

operant behavior Behavior that an organism produces that has some impact on the environment.

reinforcer Any stimulus or event that functions to increase the likelihood of the behavior that led to it.

punisher Any stimulus or event that functions to decrease the likelihood of the behavior that led to it.

Operant Conditioning: Reinforcements from the Environment

The learned behaviors you've seen so far share a common feature: They all occurred beyond the voluntary control of the organism. Most animals don't voluntarily salivate or feel spasms of anxiety; rather, these animals exhibit these responses involuntarily during the conditioning process. In fact, these reflexlike behaviors make up only a small portion of our behavioral repertoires. The remainder are behaviors that we voluntarily perform, behaviors that modify and change the environment around us. The study of classical conditioning is the study of behaviors that are *reactive*. We turn now to a different form of learning: **operant conditioning**, *a type of learning in which the consequences of an organism's behavior determine whether it will be repeated in the future*. The study of operant conditioning is the exploration of behaviors that are *active*.

The Early Days: The Law of Effect

The study of how active behavior affects the environment began at about the same time as classical conditioning. In fact, Edward L. Thorndike (1874–1949) first examined active behaviors back in the 1890s, before Pavlov published his findings. Thorndike's research focused on *instrumental behaviors*—that is, behavior that required an organism to *do* something, solve a problem, or otherwise manipulate elements of its environment (Thorndike, 1898). For example, Thorndike completed several experiments using a puzzle box, which was a wooden crate with a door that would open when a concealed lever was moved in the right way (see **FIGURE 6.6**). A hungry cat placed in a puzzle box would try various behaviors to get out—scratching at the door, meowing loudly, sniffing the inside of the box, putting its paw through the openings—but only one behavior opened the door and led to food: tripping the lever in just the right way. After this happened, Thorndike placed the cat back in the box for another round. Don't get the wrong idea. Thorndike probably really liked cats. Far from teasing them, he was after an important behavioral principle.

Fairly quickly, the cats became quite skilled at triggering the lever for their release. At first, the cats enacted any number of likely (but ultimately ineffective) behaviors, but only one behavior led to freedom and food. Over time, the ineffective behaviors became less and less frequent, and the one instrumental behavior (going right for the latch) became more frequent (see **FIGURE 6.7**, below).

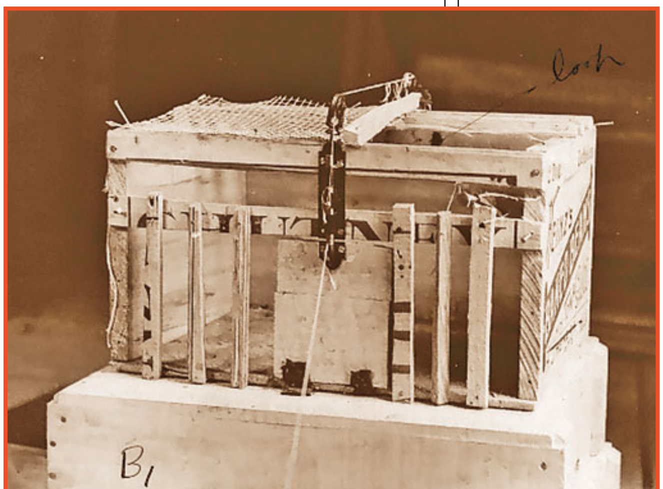

YALE UNIVERSITY LIBRARY

FIGURE 6.6
Thorndike's Puzzle Box *In Thorndike's original experiments, food was placed just outside the door of the puzzle box, where the cat could see it. If the cat triggered the appropriate lever, it would open the door and let the cat out.*

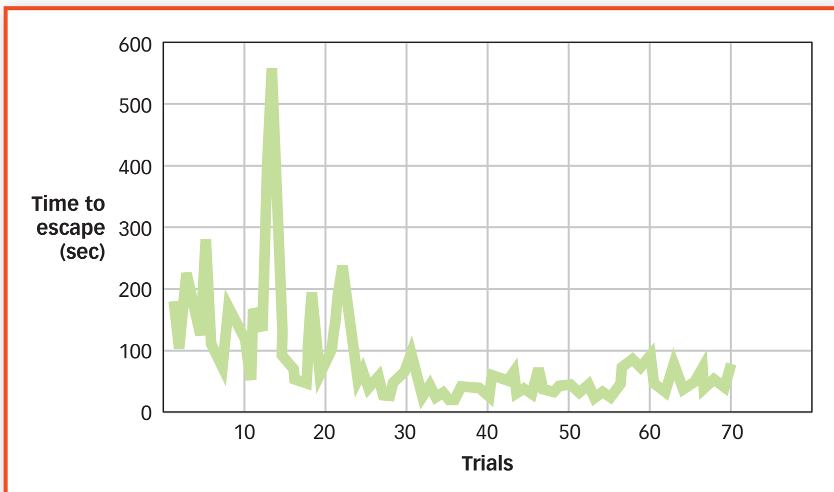

FIGURE 6.7
The Law of Effect *Thorndike's cats displayed trial-and-error behavior when trying to escape from the puzzle box. They made lots of irrelevant movements and actions until, over time, they discovered the solution. Once they figured out what behavior was instrumental in opening the latch, they stopped all other ineffective behaviors and escaped from the box faster and faster.*

From these observations, Thorndike developed the **law of effect,** which states that *behaviors that are followed by a "satisfying state of affairs" tend to be repeated, and those that produce an "unpleasant state of affairs" are less likely to be repeated.*

The circumstances that Thorndike used to study learning were very different from those in studies of classical conditioning. Remember that in classical conditioning experiments, the US occurred on every training trial no matter what the animal did. Pavlov delivered food to the dog whether it salivated or not. But in Thorndike's work, the behavior of the animal determined what happened next. If the behavior was "correct" (i.e., the latch was triggered), the animal was rewarded with food. Incorrect behaviors produced no results and the animal was stuck in the box until it performed the correct behavior. Although different from classical conditioning, Thorndike's work resonated with most behaviorists at the time: It was still observable, quantifiable, and free from explanations involving the mind (Galef, 1998).

● **What is the relationship between behavior and reward?**

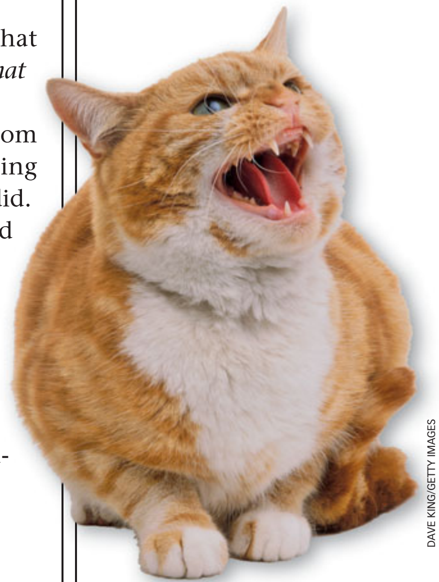

DAVE KING/GETTY IMAGES

Reinforcement, Punishment, and the Development of Operant Conditioning

Several decades after Thorndike's work, B. F. Skinner (1904–1990) coined the term ***operant behavior*** to refer to *behavior that an organism produces that has some impact on the environment.* In Skinner's system, all of these emitted behaviors "operated" on the environment in some manner, and the environment responded by providing events that either strengthened those behaviors (i.e., they *reinforced* them) or made them less likely to occur (i.e., they *punished* them).

In order to study operant behavior scientifically, Skinner developed a variation on Thorndike's puzzle box. The *operant chamber,* or *Skinner box,* as it is commonly called (FIGURE 6.8), allows a researcher to study the behavior of small organisms in a controlled environment.

Skinner's approach to the study of learning focused on *reinforcement* and *punishment.* These terms, which have commonsense connotations, turned out to be rather difficult to define. For example, some people love roller coasters, whereas others find them horrifying; the chance to go on one will be a reinforcement for one group but a punishment for another. Dogs can be trained with praise and a good belly rub—procedures that are nearly useless for most cats. Skinner settled on a "neutral" definition that would characterize each term by its effect on behavior. Therefore, a **reinforcer** is *any stimulus or event that functions to increase the likelihood of the behavior that led to it,* whereas a **punisher** is *any stimulus or event that functions to decrease the likelihood of the behavior that led to it.*

WALTER DAWN/PHOTO RESEARCHERS

FIGURE 6.8 Skinner Box *In a typical Skinner box, or operant conditioning chamber, a rat, pigeon, or other suitably sized animal is placed in this environment and observed during learning trials that use operant conditioning principles.*

Whether a particular stimulus acts as a reinforcer or punisher depends in part on whether it increases or decreases the likelihood of a behavior. Presenting food is usually reinforcing, producing an increase in the behavior that led to it; removing food is often punishing, leading to a decrease in the behavior. Turning on an electric shock is typically punishing (the behavior that led to it); turning it off is rewarding (and increases the behavior that led to it).

To keep these possibilities distinct, Skinner used the term *positive* for situations in which a stimulus was presented and *negative* for situations in which it was removed. Consequently, there is *positive reinforcement* (where something desirable is presented) and *negative reinforcement* (where something undesirable is removed), as well as *positive punishment* (where something unpleasant is administered) and *negative punishment* (where something desirable is removed). Here the words *positive* and *negative* mean, respectively, something that is *added* or something that is *taken away.* As you can see from TABLE 6.1 (on page 176), positive and negative reinforcement increase the likelihood of the behavior and positive and negative punishment decrease the likelihood of the behavior.

TABLE 6.1

Reinforcement and Punishment

	Increases the Likelihood of Behavior	Decreases the Likelihood of Behavior
Stimulus is presented	Positive reinforcement	Positive punishment
Stimulus is removed	Negative reinforcement	Negative punishment

COURTESY DAIRY QUEEN

• *Negative reinforcement involves the removal of something undesirable from the environment. When Daddy stops the car, he gets a reward: His little monster stops screaming. However, from the child's perspective, this is positive reinforcement. The child's tantrum results in something positive added to the environment: stopping for a snack.*

These distinctions can be confusing at first; after all, "negative reinforcement" and "punishment" both sound like they should be "bad" and produce the same type of behavior. There are a couple of ways to keep track of these distinctions. First, remember that *positive* and *negative* simply mean *presentation* or *removal,* and the terms don't necessarily mean "good" or "bad" as they do in everyday speech. Negative reinforcement, for example, involves something pleasant; it's the *removal* of something unpleasant, like a shock, and the absence of a shock is indeed pleasant.

Second, bear in mind that reinforcement is generally more effective than punishment in promoting learning. There are many reasons (Gershoff, 2002), but one reason is this: Punishment signals that an unacceptable behavior has occurred, but it doesn't specify what should be done instead. Spanking a young child for starting to run into a busy street certainly stops the behavior—which, in this case, is probably a good idea. But it doesn't promote any kind of learning about the desired behavior. Reinforcers and punishers often gain their functions from basic biological mechanisms. Food, comfort, shelter, and warmth are examples of *primary reinforcers* because they help satisfy biological needs. However, the vast majority of reinforcers or punishers in our daily lives have little to do with biology. Handshakes, verbal approval, an encouraging grin, a bronze trophy, or money all serve powerful reinforcing functions, yet none of them taste very good or help keep you warm at night. The point is, we learn to perform a lot of behaviors based on reinforcements that have little or nothing to do with biological satisfaction.

● **Why is reinforcement more constructive than punishment in learning desired behavior?**

These *secondary reinforcers* derive their effectiveness from their associations with primary reinforcers through classical conditioning. For example, money starts out as a neutral CS that, through its association with primary USs, such as acquiring food or shelter, takes on a conditioned emotional element. Flashing lights, originally a neutral CS, acquire powerful negative elements through association with a speeding ticket and a fine. Under normal circumstances, as long as the CS-US link is maintained, the secondary reinforcers and punishers can be used to modify and control behavior. If the links are broken (i.e., if an extinction procedure is introduced), secondary reinforcers typically lose these functions. Money that is no longer backed by a solvent government quickly loses its reinforcing capacity and becomes worth no more than the paper it is printed on.

But as long as behaviors are linked with reinforcement, those behaviors should continue to occur, right? Actually, no. Sometimes, the presentation of rewards can cause the opposite effect: a decrease in performing the behavior. An example of such a mindbug is **overjustification effect,** *when external rewards can undermine the intrinsic satisfaction of performing a behavior.* In one study nursery school children were given colored pens and paper and were asked to draw whatever they wanted (Lepper & Greene, 1978). For a young child, the pleasures of drawing and creative expression are rewarding all by themselves. Some children, though, received a "Good Player Award" for their efforts at artwork, whereas other children did not. As you may have guessed, the Good Players spent more time at the task than the other children. As you may not have guessed, when the

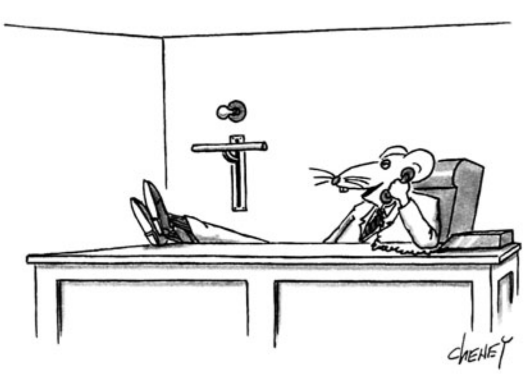

"Oh, not bad. The light comes on, I press the bar, they write me a check. How about you?"

THE NEW YORKER COLLECTION 1993 TOM CHENEY FROM CARTOONBANK.COM.

experimenters stopped handing out the Good Player certificates to the first group, the amount of time the children spent drawing dropped significantly below that of the group that never received any external reinforcements.

This was a case of *over*justification, or too much reinforcement. The children who received the extrinsic reinforcement of the certificate came to view their task as one that gets rewards. The children who didn't receive the extrinsic reinforcement continued to perform the task for its own sake. When the extrinsic rewards were later removed, children in the first group found little reason to continue engaging in the task. Other researchers have found that when people are paid for tasks such as writing poetry, drawing, or finding solutions to economic and business problems, they tend to produce *less* creative solutions when monetary rewards are offered (Amabile, 1996). You can weigh in on these issues in the Where Do You Stand? box at the end of this chapter (on page 193).

Can rewards backfire?

AL SWINKELS

Drawing pictures is fun. Drawing pictures for external rewards might, oddly enough, make drawing pictures seem like much less fun.

The Basic Principles of Operant Conditioning

After establishing how reinforcement and punishment produced learned behavior, Skinner and other scientists began to expand the parameters of operant conditioning. They started by investigating some phenomena that were well known in classical conditioning, such as discrimination, generalization, and extinction.

For example, if a pigeon is reinforced for pecking a key whenever a particular tone is sounded but never reinforced if the tone is absent, that tone will quickly become a *discriminative stimulus,* or a stimulus that is associated with reinforcement for key pecking in that situation. Pigeons, reinforced under these conditions, will quickly learn to engage in vigorous key pressing whenever the tone sounds but cease if it is turned off. The tone sets the occasion, or context, for the pigeon to emit the response.

You similarly modify your behavior based on what context you're in. We all take off our clothes at least once a day, but usually not in public. We scream at rock concerts but not in libraries. We say, "Please pass the gravy," at the dinner table but not in a classroom. Although these observations may seem like nothing more than common sense, Thorndike was the first to recognize the underlying message: Learning takes place *in contexts,* not in the free range of any plausible situation. As Skinner rephrased it later, most behavior is under *stimulus control,* which develops when a particular response only occurs when the appropriate stimulus is present.

Stimulus control, perhaps not surprisingly, shows both discrimination and generalization effects similar to those we saw with classical conditioning. To demonstrate this, researchers used either a painting by the French Impressionist Claude Monet or one of Pablo Picasso's paintings from his Cubist period for the discriminative stimulus (Watanabe, Sakamoto, & Wakita, 1995). Some participants were reinforced only if they responded when the Monet painting was present; others were reinforced for responding to the Picasso. Later, the participants were tested on new paintings, and they discriminated appropriately: Those trained with the Monet painting responded when other paintings by Monet were presented, and those trained with a Picasso painting reacted when other paintings by Picasso were shown. If these results don't seem particularly startling to you, it might help to know that the research participants were pigeons who were trained to key-peck to these various works of art. Stimulus control, and its ability to foster stimulus discrimination and stimulus generalization, is effective even if the stimulus has no meaning to the respondent.

As in classical conditioning, operant behavior undergoes extinction when the reinforcements stop. Pigeons cease pecking at a key if food is no longer presented following the behavior. You wouldn't put more money into a vending machine if it failed to give you its promised candy bar or soda. Warm smiles that are greeted with scowls and

overjustification effect Circumstances when external rewards can undermine the intrinsic satisfaction of performing a behavior.

In research on stimulus control, participants trained with Picasso paintings, such as the one on the left, responded to other paintings by Picasso or even to paintings by other Cubists. Participants trained with Monet paintings, such as the one on the right, responded to other paintings by Monet or by other French Impressionists. Interestingly, the participants in this study were pigeons.

shaping Learning that results from the reinforcement of successive approximations to a final desired behavior.

frowns will quickly disappear. On the surface, extinction of operant behavior looks like that of classical conditioning: The response rate drops off fairly rapidly, and, if a rest period is provided, spontaneous recovery is typically seen.

One important difference between classical and operant conditioning is that, in classical conditioning, responses are usually hardwired—behaviors that the animal already displays, such as salivation or fear. Classical conditioning only changes the conditions in which these behaviors are produced. By contrast, operant conditioning can produce brand-new behaviors. Have you ever been to AquaLand and wondered how the dolphins learn to jump up in the air, twist around, splash back down, do a somersault, and then jump through a hoop, all in one smooth motion? These behaviors are the result of **shaping**, or *learning that results from the reinforcement of successive approximations to a final desired behavior.* The outcomes of one set of behaviors shape the next set of behaviors

How can operant conditioning produce complex behaviors?

To illustrate the effects of shaping, Skinner noted that if you put a rat in a Skinner box and wait for it to press the bar, you could end up waiting a very long time; bar

B. F. Skinner shaping a dog named Agnes. In the span of 20 minutes, Skinner was able to use reinforcement of successive approximations to shape Agnes's behavior. The result was a pretty neat trick: to wander in, stand on hind legs, and jump.

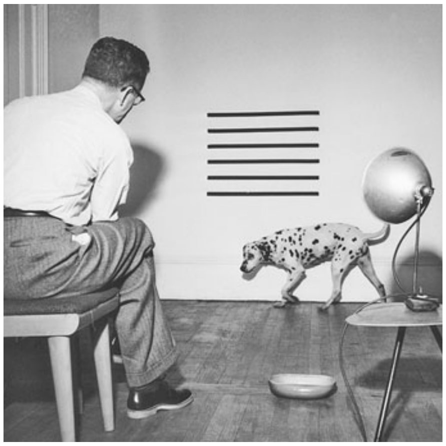

1 minute

4 minutes

8 minutes

LIBRARY OF CONGRESS/LOOK MAGAZINE PHOTOGRAPHIC COLLECTION pressing just isn't very high in a rat's natural hierarchy of responses. However, it is relatively easy to "shape" bar pressing. Watch the rat closely: If it turns in the direction of the bar, deliver a food reward. This will reinforce turning toward the bar, making such a movement more likely. Now wait for the rat to take a step toward the bar before delivering food; this will reinforce moving toward the bar. After the rat walks closer to the bar, wait until it touches the bar before presenting the food. Notice that none of these behaviors is the final desired behavior—reliably pressing the bar. Rather, each behavior is a *successive approximation* to the final product, or a behavior that gets incrementally closer to the overall desired behavior. In the dolphin example—and indeed, in many instances of animal training in which relatively simple animals seem to perform astoundingly complex behaviors—each smaller behavior is reinforced until the overall sequence of behavior gets performed reliably.

In shaping, complicated behaviors can be established by pairing responses with reinforcement. But what if reinforcement occurs, regardless of what the organism is doing? To find out, Skinner (1947) put several pigeons in Skinner boxes, set the food dispenser to deliver food every 15 seconds, and left the birds to their own devices. Later he returned and found the birds engaging in odd, idiosyncratic behaviors, such as pecking aimlessly in a corner or turning in circles. He referred to these behaviors as "superstitious" and offered a behaviorist analysis of their occurrence. The pigeons, he argued, were simply repeating behaviors that had been accidentally reinforced. A pigeon that just happened to have pecked randomly in the corner when the food showed up had connected the delivery of food to that behavior. Because this pecking behavior was "reinforced" by the delivery of food, the pigeon was likely to repeat it. Now pecking in the corner was more likely to occur, and it was more likely to be reinforced 15 seconds later when the food appeared again.

● **How would a behaviorist explain superstitions?**

For each pigeon, the behavior reinforced would be whatever the pigeon happened to be doing when the food was first delivered. Skinner's pigeons acted as though there was a causal relationship between their behaviors and the appearance of food when it was merely an accidental correlation. Superstitious behavior is not limited to pigeons, of course. Baseball players who enjoy several home runs on a day when they happened to have not showered are likely to continue that tradition, laboring under the belief that the accidental correlation between poor personal hygiene and a good day at bat is somehow causal. This "stench causes home runs" hypothesis is just one of many examples of human superstitions (Gilbert et al., 2000; Radford & Radford, 1949).

Why does Tiger Woods always wear his Sunday red shirt for the final round of a golf tournament? Some people think that he is engaging in superstitious behavior, but it's more than just that: Tiger feels more aggressive when wearing a red shirt, which helps him perform down the stretch.

JAMIE SQUIRE/GETTY IMAGES

12 minutes | 16 minutes | 20 minutes

fixed interval schedule (FI) An operant conditioning principle in which reinforcements are presented at fixed time periods, provided that the appropriate response is made.

variable interval schedule (VI) An operant conditioning principle in which behavior is reinforced based on an average time that has expired since the last reinforcement.

fixed ratio schedule (FR) An operant conditioning principle in which reinforcement is delivered after a specific number of responses have been made.

variable ratio schedule (VR) An operant conditioning principle in which the delivery of reinforcement is based on a particular average number of responses.

intermittent reinforcement An operant conditioning principle in which only some of the responses made are followed by reinforcement.

Schedules of Reinforcement

The law of effect states that behaviors that are reinforced tend to occur more often. But how much more often? Partly, this depends on how often reinforcement is received.

Skinner was intrigued by this fact, and he explored dozens of what came to be known as *schedules of reinforcement* (Ferster & Skinner, 1957) (see **FIGURE 6.9** on page 181). The two most important are *interval schedules,* based on the time intervals between reinforcements, and *ratio schedules,* based on the ratio of responses to reinforcements.

Under a **fixed interval schedule (FI)**, *reinforcements are presented at fixed time periods, provided that the appropriate response is made*. For example, on a 2-minute fixed interval schedule, a response will be reinforced, but only after 2 minutes have expired since the last reinforcement. Pigeons in Skinner boxes produce predictable patterns of behavior under these schedules. They show little responding right after the presentation of reinforcement, but as the next time interval draws to a close, they show a burst of responding. If this pattern seems odd to you, consider how often undergraduates behave exactly like this. They do relatively little work until just before the upcoming exam, then engage in a burst of reading and studying—and then probably take a little time off after the exam before they start seriously preparing for the next test.

Under a **variable interval schedule (VI)**, a *behavior is reinforced based on an average time that has expired since the last reinforcement*. For example, on a 2-minute variable interval schedule, responses will be reinforced every 2 minutes *on average* but not after each 2-minute period. Variable interval schedules typically produce steady, consistent responding because the time until the next reinforcement is less predictable. For example, a radio station might advertise that they give away concert tickets every hour, which is true, but the DJs are likely to say, "Sometime this hour, I'll be giving away a pair of tickets to see the Arctic Monkeys in concert!" which is also true. The reinforcement—getting the tickets—might average out to once an hour across the span of the broadcasting day, but the presentation of the reinforcement is variable: It might come early in the 10 o'clock hour, later in the 11 o'clock hour, and so on. The result is to keep listeners tuned in steadily throughout the day, rather than just tuning in every hour on the hour for a chance to win those tickets.

How does a radio station use scheduled reinforcements to keep you listening?

© RICHARD HUTCHINGS/PHOTOEDIT

Radio station promotions and giveaways often follow a variable interval schedule of reinforcement.

Both fixed interval schedules and variable interval schedules tend to produce slow, methodical responding because the reinforcements follow a time scale that is independent of how many responses occur. It doesn't matter if a rat on a fixed interval schedule presses a bar 1 time during a 2-minute period or 100 times: The reinforcing food pellet won't drop out of the shoot until 2 minutes have elapsed, regardless of the number of responses.

Under a **fixed ratio schedule (FR)**, *reinforcement is delivered after a specific number of responses have been made*. One schedule might present reinforcement after every fourth response; a different schedule might present reinforcement after every 20 responses. The special case of presenting reinforcement after *each* response is called *continuous reinforcement*. There are many situations in which people find themselves being reinforced on a fixed ratio schedule. Book clubs often give you a "freebie" after a set number of regular purchases, pieceworkers get paid after making a fixed number of products, and some credit card companies return to their customers a percentage of the amount charged.

Under a **variable ratio schedule (VR)**, *the delivery of reinforcement is based on a particular average number of responses*. For example, under a 10-response variable ratio schedule, reinforcement follows every 10th response—on average. Slot machines in modern casinos pay off on variable ratio schedules that are determined by the random number generator that controls the play of the machines. A casino might advertise that they pay off on "every 100 pulls on average," which could be true. However, one player might hit a jackpot after 3 pulls on a slot machine, whereas another player might not hit until after 80 pulls. The ratio of responses to reinforcements is variable, which probably helps casinos stay in business.

IMAGINECHINA

These pieceworkers in a textile factory get paid following a fixed ratio schedule: They receive payment after some set number of shirts have been sewn.

All ratio schedules encourage high and consistent rates of responding because the number of rewards received is directly related to the number of responses made. Unlike a rat following a fixed interval schedule, where food is delivered at a specified time regardless of the number of responses, rats following a ratio schedule should respond quickly and often. Not surprisingly, variable ratio schedules produce slightly higher rates of responding than fixed ratio schedules primarily because there's always the possibility of a reward after the very next response—even if a reinforcement was just obtained.

How do ratio schedules work to keep you spending your money?

All of these schedules of reinforcement provide **intermittent reinforcement**, meaning that *only some of the responses made are followed by reinforcement.* They all produce behavior that is much more resistant to extinction than a continuous reinforcement schedule. One way to think about this effect is to recognize that the more irregular and intermittent a schedule is, the more difficult it becomes for an organism to detect when it has actually been placed on extinction.

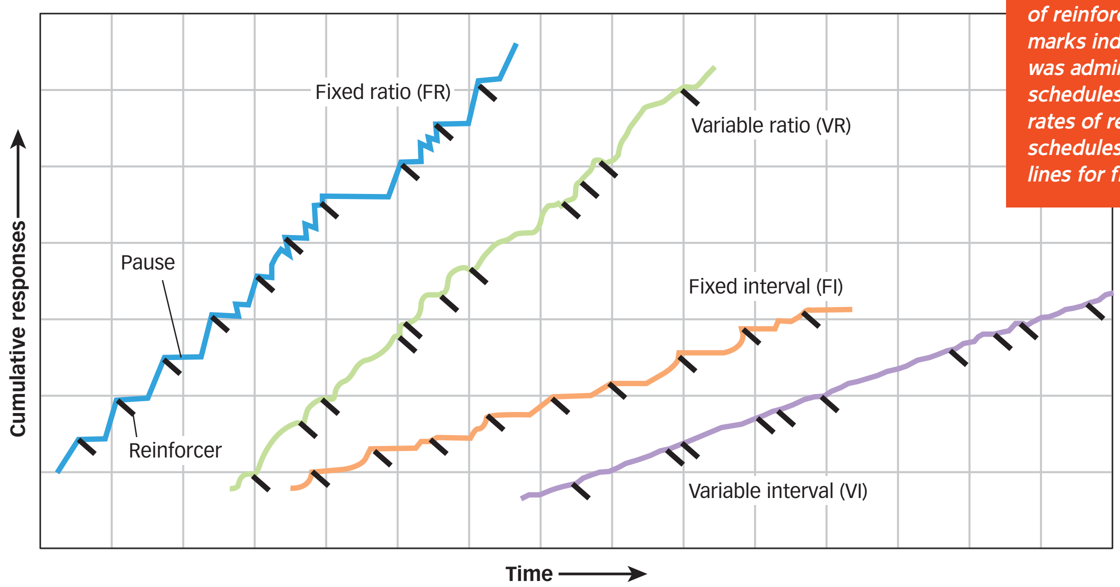

FIGURE 6.9 **Reinforcement Schedules** *Different schedules of reinforcement produce different rates of responding. These lines represent the amount of responding that occurs under each type of reinforcement. The black slash marks indicate when reinforcement was administered. Notice that ratio schedules tend to produce higher rates of responding than do interval schedules, as shown by the steeper lines for fixed ratio and variable ratio.*

OTTMAR BIERWAGEN/SPECTRUM PHOTOFILE

Slot machines in casinos pay out following a variable ratio schedule. This helps explain why some gamblers feel incredibly lucky, whereas others (like this chap) can't believe they can play a machine for so long without winning a thing.

For example, if you've just put a dollar into a soda machine that, unbeknownst to you, is broken, no soda comes out. Because you're used to getting your sodas on a continuous reinforcement schedule—one dollar produces one soda—this abrupt change in the environment is easily noticed, and you are unlikely to put additional money into the machine. In other words, you'd quickly show extinction. However, if you've put your dollar into a slot machine that, unbeknownst to you, is broken, do you stop after one or two plays? Almost certainly not. If you're a regular slot player, you're used to going for many plays in a row without winning anything, so it's difficult to tell that anything is out of the ordinary. Under conditions of intermittent reinforcement, all organisms will show considerable resistance to extinction and continue for many trials before they stop responding.

This relationship between intermittent reinforcement schedules and the robustness of the behavior they produce is called the **intermittent reinforcement effect**, *the fact that operant behaviors that are maintained under intermittent reinforcement schedules resist extinction better than those maintained under continuous reinforcement.* In one extreme case, Skinner gradually extended a variable ratio schedule until he managed to get a pigeon to make an astonishing 10,000 pecks at an illuminated key for one food reinforcer! Behavior maintained under a schedule like this is virtually immune to extinction.

A Deeper Understanding of Operant Conditioning

Like classical conditioning, operant conditioning also quickly proved powerful. It's difficult to argue this fact when a rat learns to perform relatively complex behaviors after only 20 minutes of practice, prompted by little more than the skillful presentation of rat chow. The results are evident: "Learning" in its most fundamental sense is a change in behavior brought about by experience. This observation was enough for the early behaviorists, who didn't include the mind in the analysis of an organism's actions. Skinner was satisfied to observe an organism perform the behavior; he didn't look for a deeper explanation of mental processes (Skinner, 1950). However, some research on operant conditioning digs deeper into the underlying mechanisms that produce the familiar outcomes of reinforcement. Let's examine three elements that expand our view of operant conditioning beyond strict behaviorism: the neural, cognitive, and evolutionary elements of operant conditioning.

The Neural Elements of Operant Conditioning

Soon after psychologists came to appreciate the range and variety of things that could function as reinforcers, they began looking for underlying brain mechanisms that might account for these effects. The first hint of how specific brain structures might contribute to the process of reinforcement came from the discovery of what came to be called *pleasure centers*. James Olds and his associates inserted tiny electrodes into different parts of a rat's brain and allowed the animal to control electric stimulation of its own brain by pressing a bar. They discovered that some brain areas, particularly those in the limbic system (see Chapter 3), produced what appeared to be intensely positive experiences: The rats would press the bar repeatedly to stimulate these structures. The researchers observed that these rats would ignore food, water, and other life-sustaining necessities for hours on end simply to receive stimulation directly in the brain. They then called these parts of the brain "pleasure centers" (Olds, 1956) (see FIGURE 6.10, on page 183).

● **Where are the brain's "pleasure centers"?**

In the years since these early studies, researchers have identified a number of structures and pathways in the brain that deliver rewards through stimulation (Wise, 1989, 2005). The neurons in the *medial forebrain bundle,* a pathway that meanders its way from the midbrain through the *hypothalamus* into the *nucleus accumbens,* are the most susceptible to stimulation that produces pleasure. This is not surprising as psychologists have identified this bundle of cells as crucial to behaviors that clearly involve pleasure, such as eating, drinking, and engaging in sexual activity. Second, the neurons along

intermittent reinforcement effect The fact that operant behaviors that are maintained under intermittent reinforcement schedules resist extinction better than those maintained under continuous reinforcement.

latent learning A condition in which something is learned but it is not manifested as a behavioral change until sometime in the future.

this pathway are *dopaminergic;* that is, they secrete the neurotransmitter *dopamine*. Remember from Chapter 3 that higher levels of dopamine in the brain are usually associated with positive emotions.

Researchers have found good support for this "reward center." First, as you've just seen, rats will work to stimulate this pathway at the expense of other basic needs (Olds & Fobes, 1981). However, if drugs that block the action of dopamine are administered to the rats, they cease stimulating the pleasure centers (Stellar, Kelley, & Corbett, 1983). Second, drugs such as cocaine, amphetamine, and opiates activate these pathways and centers (Moghaddam & Bunney, 1989), but dopamine-blocking drugs dramatically diminish their reinforcing effects (White & Milner, 1992). Third, fMRI studies (see Chapter 3) show increased activity in the nucleus accumbens in heterosexual men looking at pictures of attractive women (Aharon et al., 2001) and in individuals who believe they are about to receive money (Knutson et al., 2001). Finally, rats given primary reinforcers such as food or water or that are allowed to engage in sexual activity show increased dopamine secretion in the nucleus accumbens—but only if the rats are hungry, thirsty, or sexually aroused (Damsma et al., 1992). After all, food tastes a lot better when we are hungry, and sexual activity is more pleasurable when we are aroused. These biological structures underlying rewards and reinforcements probably evolved to ensure that species engaged in activities that helped survival and reproduction.

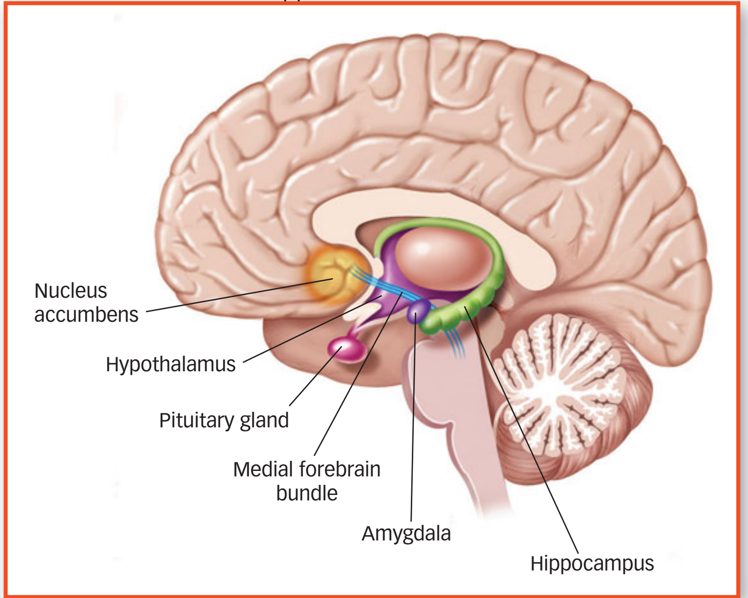

FIGURE 6.10 **Pleasure Centers in the Brain** *The nucleus accumbens, medial forebrain bundle, and hypothalamus are all major pleasure centers in the brain.*

The Cognitive Elements of Operant Conditioning

In addition to studying the brain substrates of operant conditioning, other researchers began to question Skinner's strictly behaviorist interpretation of learning and to suggest that cognition might play a role. Edward Chace Tolman (1886–1959) was the strongest early advocate of a cognitive approach to operant learning. Tolman argued that there was more to learning than just knowing the circumstances in the environment (the properties of the stimulus) and being able to observe a particular outcome (the reinforced response). Instead, Tolman focused less on the stimulus-response connection and more on what happens in the organism's mind when faced with the stimulus.

One phenomenon that suggested that simple stimulus-response interpretations of operant learning were inadequate was **latent learning**, in which *something is learned but it is not manifested as a behavioral change until sometime in the future*. Latent learning can easily be established in rats and occurs without any obvious reinforcement, a finding that posed a direct challenge to the then-dominant behaviorist position that all learning required some form of reinforcement (Tolman & Honzik, 1930b). Tolman gave three groups of rats access to a complex maze every day for over 2 weeks. The control group never received any reinforcement for navigating the maze. They were simply allowed to run around until they reached the goal box at the end of the maze. In FIGURE 6.11 (on page 184) you can see that over the 2 weeks of the study, this group (in green) got a little better at finding their way through the maze but not by much. A second group of rats received regular reinforcements; when they reached the goal box, they found a small food reward there. Not surprisingly, these rats showed clear learning, as can be seen in blue in FIGURE 6.11. A third group was treated exactly like the control group for the first 10 days and then rewarded for the last 7 days. This group's behavior (in orange) was quite striking. For the first 10 days, they behaved like the rats in the control group. However, during the final 7 days, they behaved a lot like the rats in the second group that had been reinforced every day. Clearly, the rats in this third group had learned a lot about the maze and the location of the goal box during those first 10 days even though they had not received any reinforcements for their behavior. In other words, they showed evidence of latent learning.

Edward Chace Tolman advocated a cognitive approach to operant learning and provided evidence that in maze learning experiments, rats develop a mental picture of the maze, which he called a cognitive map.

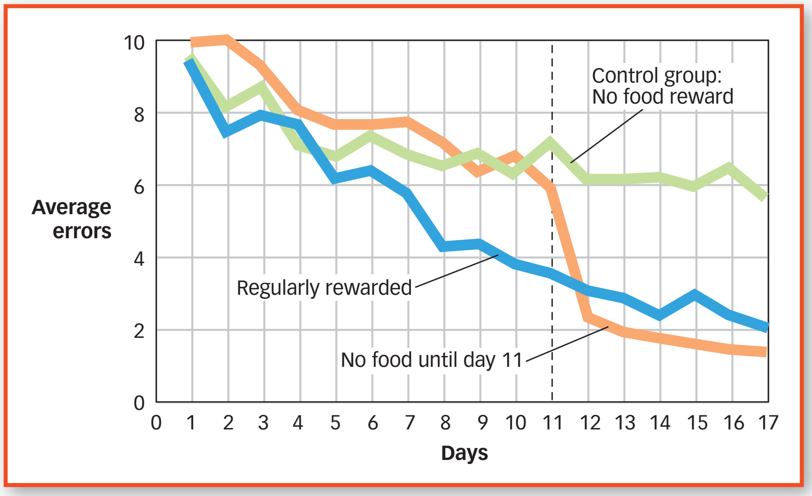

FIGURE 6.11 **Latent Learning** *Rats in a control group that never received any reinforcement (in green) improved at finding their way through the maze over 17 days but not by much. Rats that received regular reinforcements (in blue) showed fairly clear learning; their error rate decreased steadily over time. Rats in the latent learning group (in orange) were treated exactly like the control group rats for the first 10 days and then like the regularly rewarded group for the last 7 days. Their dramatic improvement on day 12 shows that these rats had learned a lot about the maze and the location of the goal box even though they had never received reinforcements. Notice also that on the last 7 days, these latent learners actually seem to make fewer errors than their regularly rewarded counterparts.*

These results suggested to Tolman that beyond simply learning "start here, end here," his rats had developed a sophisticated mental picture of the maze. Tolman called this a **cognitive map**, or *a mental representation of the physical features of the environment*. One simple experiment provided support for Tolman's theories and wreaked havoc with the noncognitive explanations offered by staunch behaviorists. Tolman trained a group of rats in the maze shown in **FIGURE 6.12a**. As you can see, rats run down a straightaway, take a left, a right, a long right, and then end up in the goal box at the end of the maze.

● What are "cognitive maps," and why are they a challenge to behaviorism?

After they had mastered the maze, Tolman changed things around a bit and put them in the maze shown in **FIGURE 6.12b**. The goal box was still in the same place relative to the start box. However, many alternative paths now spoked off the main platform, and the main straightaway that the rats had learned to use was blocked. Most behaviorists would predict that the rats in this situation—running down a familiar path only to find it blocked—would show stimulus generalization and pick the next closest path, such as

FIGURE 6.12 **Cognitive Maps** *(a) Rats trained to run from a start box to a goal box in the maze on the left mastered the task quite readily. When these rats were then placed in the maze on the right (b), in which the main straightaway had been blocked, they did something unusual. Rather than simply backtrack and try the next closest runway (i.e., those labeled 8 or 9 in the figure), which would be predicted by stimulus generalization, the rats typically chose runway 5, which led most directly to where the goal box had been during their training. The rats had formed a cognitive map of their environment and so knew where they needed to end up, spatially, compared to where they began.*

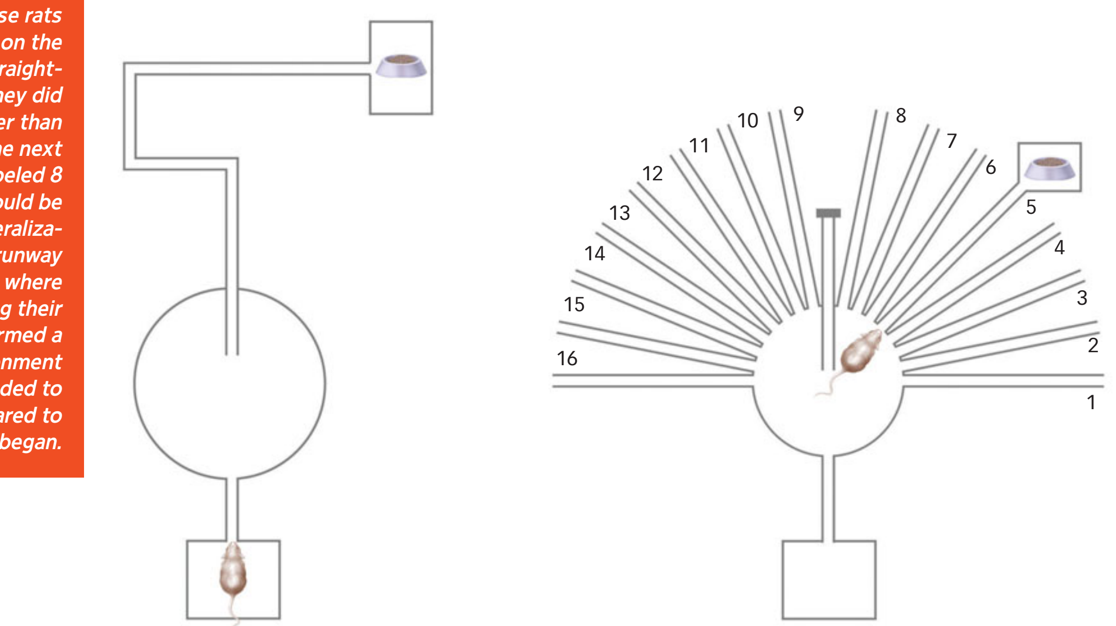

(a) Apparatus used in preliminary training

(b) Apparatus used in test trial

one immediately adjacent to the straightaway. This was not what Tolman observed. When faced with the blocked path, the rats instead ran all the way down the path that led directly to the goal box. The rats had formed a sophisticated cognitive map of their environment and behaved in a way that suggested they were successfully following that map after the conditions had changed. Latent learning and cognitive maps suggest that operant conditioning involves much more than an animal responding to a stimulus. Tolman's experiments strongly suggest that there is a cognitive component, even in rats, to operant learning.

The Evolutionary Elements of Operant Conditioning

As you'll recall, classical conditioning has an adaptive value that has been fine-tuned by evolution. Not surprisingly, we can also view operant conditioning from an evolutionary perspective. This viewpoint grew out of a set of curious observations from the early days of conditioning experiments. Several behaviorists were using simple T mazes like the one shown in **FIGURE 6.13.** If a rat found food in one arm of the maze on the first trial of the day, it typically ran down the *other* arm on the very next trial. A staunch behaviorist wouldn't expect the rats to behave this way. After all, the rats in these experiments were hungry, and they had just been reinforced for turning in a particular direction. According to operant conditioning, this should *increase* the likelihood of turning in that same direction, not reduce it. With additional trials the rats eventually learned to go to the arm with the food, but they had to learn to overcome this initial tendency to go "the wrong way." How can we explain this mindbug?

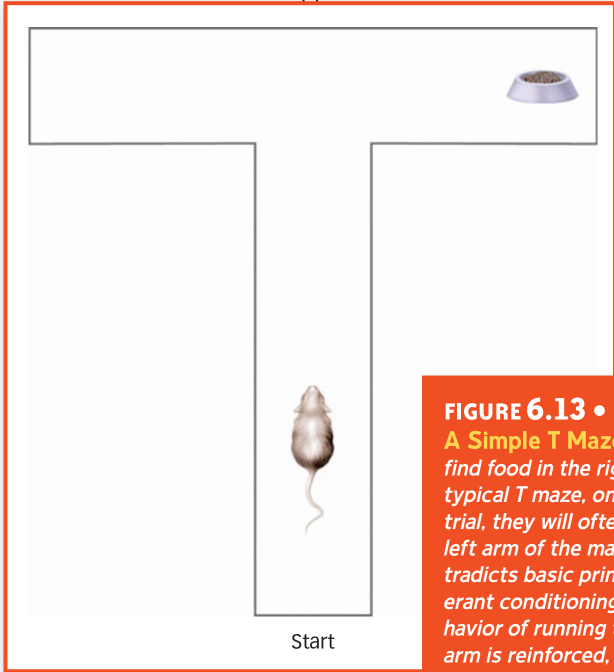

FIGURE 6.13 **A Simple T Maze** *When rats find food in the right arm of a typical T maze, on the next trial, they will often run to the left arm of the maze. This contradicts basic principles of operant conditioning: If the behavior of running to the right arm is reinforced, it should be more likely to occur again in the future. However, this behavior is perfectly consistent with a rat's evolutionary preparedness. Like most foraging animals, rats explore their environments in search of food and seldom return to where food has already been found. Quite sensibly, if food has already been found in the right arm of the T maze, the rat will search the left arm next to see if more food is there.*

What was puzzling from a behaviorist perspective makes sense when viewed from an evolutionary perspective. Rats are foragers, and like all foraging species, they have evolved a highly adaptive strategy for survival. They move around in their environment looking for food. If they find it somewhere, they eat it (or store it) and then go look somewhere else for more. If they do not find food, they forage in another part of the environment. So, if the rat just found food in the *right* arm of a T maze, the obvious place to look next time is the *left* arm. The rat knows that there isn't any more food in the right arm because it just ate the food it found there! Indeed, foraging animals such as rats have well-developed spatial representations that allow them to search their environment efficiently. So, in this case it's not the rat who is the victim of a mindbug—it's the behaviorist theorist!

● What explains a rat's behavior in a T-maze?

Two of Skinner's former students, Keller Breland and Marian Breland, were among the first researchers to discover that it wasn't just rats in T mazes that presented a problem for behaviorists (Breland & Breland, 1961). For example, the Brelands, who made a career out of training animals for commercials and movies, often used pigs because pigs are surprisingly good at learning all sorts of tricks. However, they discovered that it was extremely difficult to teach a pig the simple task of dropping coins in a box. Instead of depositing the coins, the pigs persisted in rooting with them as if they were digging them up in soil, tossing them in the air with their snouts and pushing them around. The Brelands tried to train raccoons at the same task, with different but equally dismal results. The raccoons spent their time rubbing the coins between their paws instead of dropping them in the box.

Having learned the association between the coins and food, the animals began to treat the coins as stand-ins for food. Pigs are biologically predisposed to root out their food, and raccoons have evolved to clean their food by rubbing it with their paws is exactly what each species of animal did with the coins.

cognitive map A mental representation of the physical features of the environment.

The misbehavior of organisms: Pigs are biologically predisposed to root out their food, just as raccoons are predisposed to wash their food. Trying to train either species to behave differently can prove to be an exercise in futility.

The Brelands' work shows that each species, including humans, is biologically predisposed to learn some things more readily than others and to respond to stimuli in ways that are consistent with its evolutionary history (Gallistel, 2000). Such adaptive behaviors, however, evolved over extraordinarily long periods and in particular environmental contexts. If those circumstances change, some of the behavioral mechanisms that support learning can lead an organism astray. Raccoons that associated coins with food failed to follow the simple route to obtaining food by dropping the coins in the box; "nature" took over, and they wasted time rubbing the coins together. The point is that although much of every organism's behavior results from predispositions sharpened by evolutionary mechanisms, these mechanisms sometimes can have ironic consequences.

summary quiz [6.3]

8. Which of the following is an example of a secondary reinforcer?
a. food
b. shelter
c. warmth
d. money

9. Some college students do relatively little work until just before the upcoming exam, when they then engage in a burst of studying. After the exam, they take some time off before starting the cycle again. These students are operating under which schedule of reinforcement?
a. fixed interval
b. variable interval
c. fixed ratio
d. variable ratio

10. The neurons in the __________ are crucial to behaviors associated with pleasure, such as eating, drinking, and sexual activity.
a. pituitary gland
b. medial forebrain bundle
c. hippocampus
d. parietal lobe

11. __________ is a condition in which something is learned but not manifested in a behavioral change until sometime in the future.
a. Latent learning
b. Observational learning
c. Implicit learning
d. Successive approximation

Observational Learning: Look at Me

The guiding principle of operant conditioning is that reinforcement determines future behavior. That tenet fit well with behaviorism's insistence on observable action and the behaviorists' reluctance to consider what was going on in the mind. As we've already seen, however, cognition helps explain why operant conditioning doesn't always happen as behaviorists would expect. The next section looks at learning by keeping one's eyes and ears open to the surrounding environment and further chips away at strict behaviorist doctrine.

STANLEY B. BURNS, MD & THE BURNS ARCHIVE N.Y./PHOTO RESEARCHERS

Observational learning plays an important role in surgical training, as illustrated by the medical students observing famed German surgeon Vincenz Czerny (beard and white gown) perform stomach surgery in 1901 at a San Francisco hospital.

Learning without Direct Experience

Four-year-old Rodney and his 2-year-old sister Margie had always been told to keep away from the stove, and that's good advice for any child. Being a mischievous imp, however, Rodney decided one day to heat up a burner and place his hand over it . . . until the singeing of his flesh led him to recoil, shrieking in pain. Rodney was more scared than hurt, really—and no one hearing this story doubts that he learned something important that day. But little Margie, who stood by watching these events unfold, *also* learned the same lesson. Rodney's story is a behaviorist's textbook example: The administration of punishment led to a learned change in his behavior. But how can we explain Margie's learning? She received neither punishment nor reinforcement—indeed, she didn't even have direct experience with the wicked appliance—yet it's arguable that she's just as likely to keep her hands away from stoves in the future as Rodney is.

Margie's is a case of **observational learning**, in which *learning takes place by watching the actions of others*. Observational learning challenges behaviorism's reinforcement-based explanations of classical and operant conditioning, but there is no doubt that this type of learning produces changes in behavior. In all societies, appropriate social behavior is passed on from generation to generation largely through observation (Bandura, 1965), not only through deliberate training of the young but also through young people observing the patterns of behaviors of their elders. Tasks such as using chopsticks or learning to operate a TV's remote control are more easily acquired if we watch these activities being carried out before we try ourselves. Even complex motor tasks, such as performing surgery, are learned in part through extensive observation and imitation of models. And anyone who is about to undergo surgery is grateful for observational learning. Imagine if surgeons had to learn by trial-and-error or by Skinner's technique for shaping of successive approximations!

● Why might a younger sibling appear to learn faster than a first-born?

Observational Learning in Humans

In a series of studies that have become landmarks in psychology, Albert Bandura and his colleagues investigated the parameters of observational learning (Bandura, Ross, & Ross, 1961). The researchers escorted individual preschoolers into a play area, stocked with a number of desirable toys that 4-year-olds typically like: stickers, ink stamps, crayons. An adult *model,* someone whose behavior might serve as a guide for others, was then led into the room and seated in the opposite corner, where there were several toys including a Bobo doll, which is a large inflatable plastic toy with a weighted bottom that allows it to bounce back upright when knocked down. The adult played quietly for a bit but then started aggressing toward the Bobo doll, knocking it down, jumping on it, hitting it with the mallet, kicking it around the room, and yelling "Pow!" and "Kick him!" When the children who observed these actions were later allowed to play with a

observational learning A condition in which learning takes place by watching the actions of others.

ALEX SEGRE/ALAMY

Video games have become a must-have device in many households. Research on observational learning suggests that seeing violent images—in video games, on television, or in movies—can increase the likelihood of enacting violent behavior.

variety of toys, including a Bobo doll, they were more than twice as likely to interact with it in an aggressive manner as a group of children who hadn't observed the aggressive model.

As **FIGURE 6.14** shows, the degree of imitation that the children showed was startling. In fact, the adult model purposely used novel behaviors such as hitting the doll with a mallet or throwing it up in the air so that the researchers could distinguish aggressive acts that were clearly the result of observational learning. The children in these studies also showed that they were sensitive to the consequences of the actions they observed. When they saw the adult models being punished for behaving aggressively, the children showed considerably less aggression. When the children observed a model being rewarded and praised for aggressive behavior, they displayed an increase in aggression (Bandura, Ross, & Ross, 1963).

What did the Bobo doll experiment show about children and aggressive behavior?

The observational learning seen in Bandura's studies has implications for social learning, cultural transmission of norms and values, and psychotherapy, as well as moral and ethical issues (Bandura, 1977, 1994). For example, a recent review of the literature on the effects of viewing violence on subsequent behavior concluded that viewing media violence has both immediate and long-term effects in increasing the likelihood of aggressive and violent behavior among youth (Anderson et al., 2003). This conclusion

FIGURE 6.14

Beating Up Bobo *Children who were exposed to an adult model who behaved aggressively toward a Bobo doll were likely to behave aggressively themselves. This behavior occurred in the absence of any direct reinforcement. Observational learning was responsible for producing the children's behaviors.*

ALBERT BANDURA, DEPT. OF PSYCHOLOGY, STAMFORD UNIVERSITY speaks volumes about the impact of violence and aggression as presented on TV, in movies, and in video games on our society, but it is hardly surprising in light of Bandura's pioneering research more than 40 years earlier.

Observational Learning in Animals

Humans aren't the only creatures capable of learning through observing. A wide variety of species learns by observing. In one study, for example, pigeons watched other pigeons get reinforced for either pecking at the feeder or stepping on a bar. When placed in the box later, the pigeons tended to use whatever technique they had observed other pigeons using earlier (Zentall, Sutton, & Sherburne, 1996). In another series of studies, researchers showed that laboratory-raised rhesus monkeys that had never seen a snake would develop a fear of snakes simply by observing the fear reactions of other monkeys (Cook & Mineka, 1990; Mineka & Cook, 1988). These results also support our earlier discussion of how each species has evolved particular biological predispositions for specific behaviors. Virtually every rhesus monkey raised in the wild has a fear of snakes, which strongly suggests that such a fear is one of this species' predispositions.

Observational learning may involve a neural component as well. *Mirror neurons* are a type of cell found in the brains of primates (including humans). Mirror neurons fire when an animal performs an action, such as when a monkey reaches for a food item. More importantly, however, mirror neurons also fire when an animal watches someone *else* perform the same specific task (Rizzolatti & Craighero, 2004). Although this "someone else" is usually a fellow member of the same species, some research suggests that mirror neurons in monkeys also fire when they observe humans performing an action (Fogassi et al., 2005). For example, monkeys' mirror neurons fired when they observed humans grasping for a piece of food, either to eat it or to place it in a container.

● **What do mirror neurons do?**

Mirror neurons, then, may play a critical role in the imitation of behavior as well as the prediction of future behavior (Rizzolatti, 2004). If the neurons fire when another organism is seen performing an action, it could indicate an awareness of intentionality, or that the animal is anticipating a likely course of future actions. Both of these elements—rote imitation of well-understood behaviors and an awareness of how behavior is likely to unfold—contribute to observational learning.

summary quiz [6.4]

12. After watching her 4-year-old brother Anthony burn his hand on a hot stove, 2-year-old Isabel refused to even go near the stove. Her behavior is best explained by the concept of

a. negative reinforcement.
b. positive reinforcement.
c. observational learning.
d. punishment.

13. Which is true of observational learning?

a. Although humans learn by observing others, nonhuman animals seem to lack this capability.
b. If a child sees an adult being punished for engaging in a certain behavior, the child is less likely to imitate the behavior.
c. Humans learn complex behaviors more readily by trial and error than by observation.
d. Viewing media violence does not affect the likelihood of aggressive behavior among youth.

14. Mirror neurons in the brain fire
- a. when an individual performs an action, but not when the individual watches someone else perform that action.
- b. when an individual watches someone perform an action, but not when the individual performs the action.
- c. when an individual watches someone get punished, but not when the individual watches someone get rewarded.
- d. when an individual either performs an action or watches someone else perform that action.

Implicit Learning: Under the Wires

So far, we have covered a lot of what is known about learning with only the briefest consideration of *awareness* in the learning process. You may remember we distinguished between explicit learning and implicit learning at the beginning of the chapter. People often know that they are learning, are aware of what they're learning, and can describe what they know about a topic. If you have learned something concrete, such as doing arithmetic or typing on a computer keyboard, you know that you know it and you know *what* it is you know.

But did Pavlov's dogs *know* that they had been conditioned to salivate to a bell? Did Adam and Teri in our opening vignette understand that they had learned to associate their child's toy with an emotional event? Were Bandura's young research participants aware that the adult model was affecting their behavior? It certainly makes sense to ask whether these basic learning processes in humans require an awareness on the part of the learner.

For starters, it's safe to assume that people are sensitive to the patterns of events that occur in the world around them. Most people don't stumble through life thoroughly unaware of what's going on. But people usually are attuned to linguistic, social, emotional, or sensorimotor events in the world around them so much so that they gradually build up internal representations of those patterns that were acquired without explicit awareness. This process is often called **implicit learning**, or *learning that takes place largely independent of awareness of both the process and the products of information acquisition.* As an example, although children are often given explicit rules of social conduct ("Don't chew with your mouth open"), they learn how to behave in a civilized way through experience. They're probably not aware of when or how they learned a particular course of action and may not even be able to state the general principle underlying their behavior. Yet most kids have learned not to eat with their feet, to listen when they are spoken to, and not to kick the dog. Implicit learning is knowledge that sneaks in "under the wires."

● How can you learn something without being aware of it?

Ways to Study Implicit Learning

Early studies of implicit learning showed research participants 15- or 20-letter strings and asked them to memorize them. The letter strings, which at first glance look like nonsense syllables, were actually formed using a complex set of rules called an *artificial grammar*. Take a look at the letter strings shown in **FIGURE 6.15** (on page 191). The ones on the left are "correct" and follow the rules of the artificial grammar; the ones on the right all violated the rules. The differences are pretty subtle, and if you haven't been through the learning phase of the experiment, both sets look a lot alike. Participants were not told anything about the rules, but with experience, they gradually developed

implicit learning Learning that takes place largely independent of awareness of both the process and the products of information acquisition.

Grammatical Strings	Nongrammatical Strings
VXJJ	VXTJJ
XXVT	XVTVVJ
VJTVXJ	VJTTVTV
VJTVTV	VJTXXVJ
XXXXVX	XXXVTJJ

FIGURE 6.15 **Artificial Grammar and Implicit Learning** *These are examples of letter strings formed by an artificial grammar. Research participants are exposed to the rules of the grammar and are later tested on new letter strings. Participants show reliable accuracy at distinguishing the valid, grammatical strings from the invalid, nongrammatical strings even though they usually can't explicitly state the rule they are following when making such judgments. Using an artificial grammar is one way of studying implicit learning (Reber, 1996).*

a vague, intuitive sense of the "correctness" of particular letter groupings. These letter groups became familiar to the participants, and they processed these letter groupings more rapidly and efficiently than the "incorrect" letter groupings (Reber, 1967, 1996). Research participants were then asked to classify new letter strings based on whether they follow the rules of the grammar. People turn out to be quite good at this task (usually they get between 60% and 70% correct), but they are unable to provide much in the way of explicit awareness of the rules and regularities that they are using. The experience is like when you come across a sentence with a grammatical error—you are immediately aware that something is wrong, and you can certainly make the sentence grammatical. But unless you are a trained linguist, you'll probably find it difficult to articulate which rules of English grammar were violated or which rules you used to repair the sentence.

Why are tasks learned implicitly difficult to explain to others?

Other studies of implicit learning have used a *serial reaction time* task (Nissen & Bullemer, 1987). Here research participants are presented with five small boxes on a computer screen. Each box lights up briefly; when it does, the person is asked to press the button that is just underneath that box as quickly as possible. As with the artificial grammar task, the sequence of lights appears to be random, but in fact it follows a pattern. Research participants eventually get faster with practice as they learn to anticipate which box is most likely to light up next. But, if asked, they are generally unaware that there is a pattern to the lights.

Implicit learning has some characteristics that distinguish it from explicit learning. For example, when asked to carry out implicit tasks, people differ relatively little from one another, but on explicit tasks, such as conscious problem solving, they show large individual-to-individual differences (Reber, Walkenfeld, & Hernstadt, 1991). Implicit learning also seems to be unrelated to IQ: People with high scores on standard intelligence tests are no better at implicit learning tasks, on average, than those whose scores are more modest (Reber & Allen, 2000). Implicit learning changes little across the life span. Eight-month-old infants can develop implicit learning of complex auditory patterns as well as college students (Saffran, Aslin, & Newport, 1996). At the other end of the life span, implicit learning abilities decline more slowly in old age than explicit learning abilities (Howard & Howard, 1997).

Implicit learning is remarkably resistant to various disorders that are known to affect explicit learning. A group of patients suffering from various psychoses were so severely impaired that they could not solve simple problems that college students had little difficulty with. Yet these patients were able to solve an artificial grammar learning task about as well as college students (Abrams & Reber, 1988). Other studies have found that profoundly amnesic patients not only show normal implicit memories but also display virtually normal implicit learning of artificial grammar (Knowlton, Ramus, & Squire, 1992). In fact, these patients made accurate judgments about novel letter strings even though they had essentially no explicit memory of having been in the learning phase of the experiment!

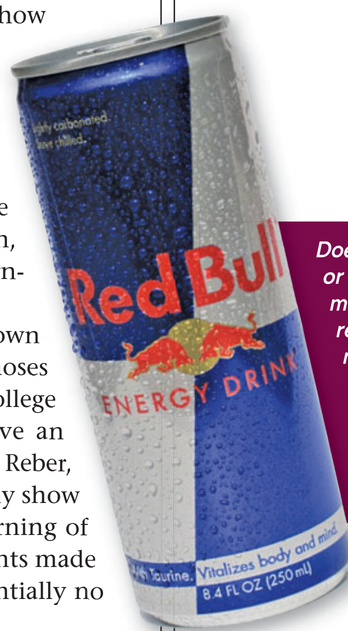

Does studying all night help or hurt your chances on the morning's exam? The latest research shows that information is consolidated by the brain during sleep, strengthening learning and making you better prepared to take the test.

Implicit and Explicit Learning Use Distinct Neural Pathways

The fact that patients suffering from psychoses or amnesia show implicit learning strongly suggests that the brain structures that underlie implicit leaning are distinct from those that underlie explicit learning. What's more, it appears that distinct regions of the brain may be activated depending on how people approach a task.

For example, in one study, participants saw a series of dot patterns, each of which looked like an array of stars in the night sky (Reber et al., 2003). Actually, all the stimuli were constructed to conform to an underlying prototypical dot pattern. The dots, however, varied so much that it was virtually impossible for a viewer to guess that they all had this common structure. Before the experiment began, half of the participants were told about the existence of the prototype; in other words, they were given instructions that encouraged explicit processing. The others were given standard implicit learning instructions: They were told nothing other than to attend to the dot patterns.

What technology shows that implicit and explicit learning are associated with separate structures of the brain?

The participants were then scanned as they made decisions about new dot patterns, attempting to categorize them into those that conformed to the prototype and those that did not. Interestingly, both groups performed equally well on this task, correctly classifying about 65% of the new dot patterns. However, the brain scans revealed that the two groups were making these decisions using very different parts of their brains (see FIGURE 6.16). Participants who were given the explicit instructions showed *increased* brain activity in the prefrontal cortex, parietal cortex, hippocampus, and a variety of other areas known to be associated with the processing of explicit memories. Those given the implicit instructions showed *decreased* brain activation primarily in the occipital region, which is involved in visual processing. This finding suggests that participants recruited distinct brain structures in different ways depending on whether they were approaching the task using explicit or implicit learning.

FIGURE 6.16
Implicit and Explicit Learning Activate Different Brain Areas
Research participants were scanned with fMRI while engaged in either implicit or explicit learning about the categorization of dot patterns. The occipital region (in blue) showed decreased brain activity after implicit learning. The areas in yellow, orange, and red showed increased brain activity during explicit learning, including the left temporal lobe (far left), right frontal lobe (second from left and second from right), and parietal lobe (second from right and far right) (Reber et al., 2003).

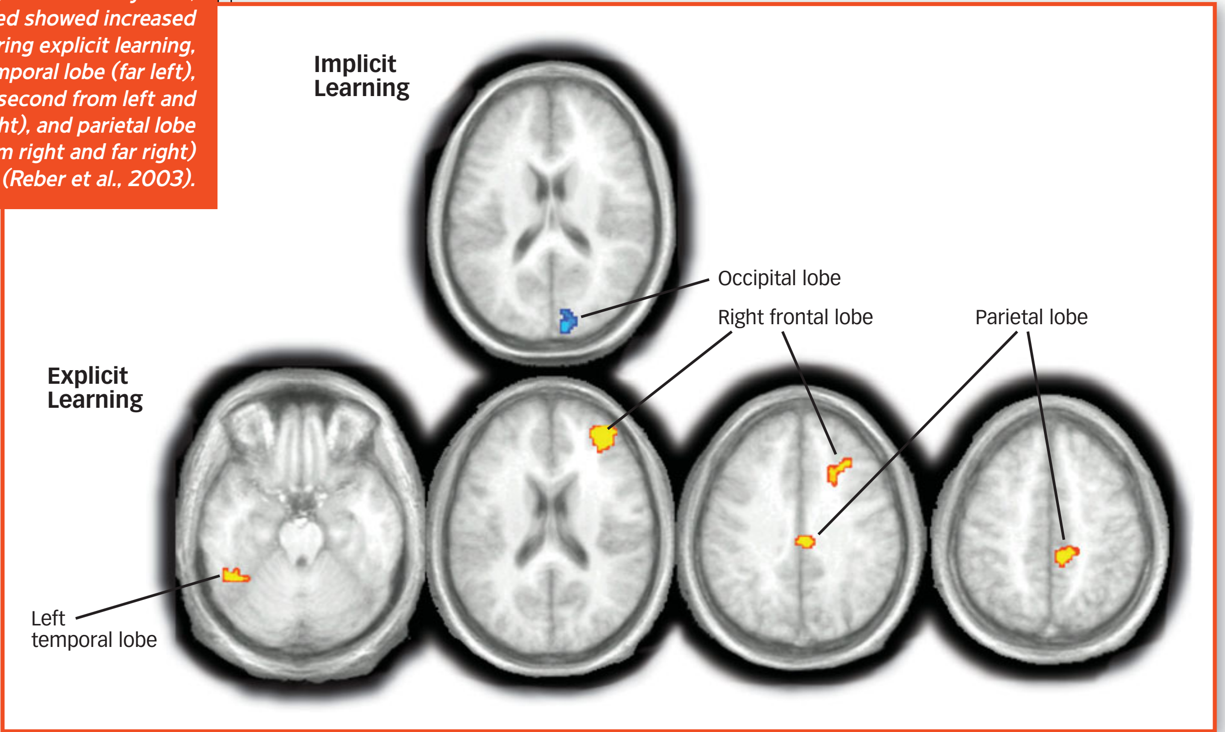

summary quiz [6.5]

15. Learning that takes place largely independent of awareness of both the process and the products of information acquisition is known as
 a. latent learning.
 b. implicit learning.
 c. unconscious learning.
 d. observational learning.

16. Which is true of implicit learning?
 a. People with high scores on intelligence tests are better implicit learners than those with low scores.
 b. Implicit learning decreases across the lifespan.
 c. People show large individual differences in implicit learning.
 d. Amnesic patients with *explicit* learning problems still show normal *implicit* learning.

17. Individuals who are given implicit instructions show *decreased* activity on which part of the brain?
 a. prefrontal cortex
 b. parietal cortex
 c. occipital region
 d. hippocampus

WhereDoYouStand?

Learning for Rewards or for Its Own Sake?

The principles of operant conditioning and the merits of reinforcement have more than found their way into mainstream culture. The least psychology-savvy parent intuitively understands that rewarding a child's good behavior should make that behavior more likely to occur in the future; the "law of effect" may mean nothing to this parent, but the principle and the outcome are readily appreciated nonetheless. If reward shapes good behavior, then more reward must be the pathway to exemplary behavior, often in the form of good grades, high test scores, and overall clean living. So, bring on the rewards!

Maybe, maybe not. As you learned earlier in this chapter, the *overjustification effect* predicts that sometimes too much external reinforcement for performing an intrinsically rewarding task can undermine future performance. Rewarding a child for getting good grades or high test scores might backfire: The child may come to see the behavior as directed toward the attainment of rewards rather than for its own satisfying outcomes. In short, learning should be fun for its own sake, not because new toys, new clothes, or cash are riding on a set of straight A's.

Many parents seem to think differently. You probably have friends whose parents shower them with gifts whenever a report card shows improvement; in fact, the website www.rewardsforgrades.com lists organizations that will give students external reinforcements for good grades, high test scores, perfect school attendance, and other behaviors that students are usually expected to produce just because they're students. Krispy Kreme offers a free doughnut for each A, Blockbuster gives free kids' movie rentals, and Limited Too offers a $5 discount on merchandise if you present a report card "with passing grades" (which, in many school districts, might mean all D's).

Or if you happen to be enrolled at Wichita State University, you already might be familiar with the Cash for Grades initiative (www.cashforgrades.com). The proposal is that an 8%-per-credit-hour increase to student fees would be used to then reward good student performance: $624 to a student with a 3.5 GPA at the end of a semester, $804 for straight A's.

Where do you stand on this issue? Is this much ado about nothing or too much of a good thing? On the one hand, some proponents of rewarding good academic performance argue that it mirrors the real world that, presumably, academic performance is preparing students to enter. After all, in most jobs, better performance is reinforced with better salaries, so why not model that in the school system? On the other hand, shouldn't the search for knowledge be reward enough? Is the subtle shift away from wanting to learn for its own sake to wanting to learn for a doughnut harmful in the long run?

CHAPTER REVIEW

Summary

Defining Learning: Experience That Causes a Permanent Change

- Learning creates a permanent change in the learner.
- Habituation is a simple form of learning that isn't permanent.
- Learning is a behavioral, observable activity, as the behaviorists proposed, but it is also a mental activity.

Classical Conditioning: One Thing Leads to Another

- In classical conditioning, a neutral stimulus (the conditioned stimulus, CS) is paired with a meaningful event (the unconditioned stimulus, US) until it elicits a response (the conditioned response, CR).
- Classical conditioning was embraced by the behaviorists, since it explained how behavior could be learned without having to invoke higher-level functions, such as thinking or awareness.
- If the CS is no longer paired with the US, the CR may decline or extinguish, but the CR often shows spontaneous recovery if the CS is presented again after a delay. Learned responses may generalize to other, similar stimuli—or the organism may discriminate stimuli by giving different responses to each.
- Even complex responses, such as emotional reactions, can be conditioned by pairing a neutral CS with a US that evokes fear or anxiety.
- The amygdala is an important brain substrate of classical fear conditioning; species may have evolved to prepared to condition quickly to biologically relevant stimuli, such as those that pose a threat to the organism's health or survival.

Operant Conditioning: Reinforcements from the Environment

- Operant conditioning is a process in which behaviors are modified according to their consequences: behaviors that are reinforced tend to increase in frequency, and behaviors that are punished tend to decrease in frequency.
- Operant conditioning shares many features with classical conditioning, including extinction, discrimination, and generalization. Complex behaviors may be achieved through shaping, and superstitious behaviors increase in frequency if they are paired with reinforcers—even though they do not "cause" those reinforcers to appear.
- Schedules of reinforcement define the relationship between response and reinforcement. Under interval schedules, reinforcements are presented for the first response that occurs after a time interval has expired. Under ratio schedules, organisms must perform a given number of responses to obtain reinforcement.
- Reinforcement may be associated with activity in "pleasure centers" in the brain, such as the nucleus accumbens, medial forebrain bundle, and hypothalamus.
- Although the strict behaviorists tried to explain operant conditioning without considering cognitive or evolutionary mechanisms, organisms behave as though they have expectations about the outcomes of their actions, and animal species may differ in terms of what they find hard or easy to learn.

Observational Learning: Look at Me

- Observational learning is a process by which organisms acquire information simply by watching other organisms behave.
- Complex behaviors can be acquired more quickly by observation than through shaping by successive approximation, and the harmful consequences of trial-and-error learning can be avoided.
- Animal studies suggest that observational learning may involve biological preparedness and may depend on mirror neurons in the brain, which fire when an individual performs an action or watches someone else perform that action.

Implicit Learning: Under the Wires

- Implicit learning is learning that occurs without explicit or conscious awareness on the part of the learner.
- Complex behaviors, such as grammar and social rules, can be learned through implicit learning. In the lab, implicit learning can be studied by asking participants to learn artificial grammars or motor tasks.
- Implicit and explicit learning differ from each other in a number of ways; for example, amnesic patients with explicit learning problems can exhibit intact implicit learning, and implicit and explicit learning recruit distinct brain structures.

Key Terms

learning (p. 162)
habituation (p. 163)
classical conditioning (p. 164)
unconditioned stimulus (US) (p. 164)
unconditioned response (UR) (p. 164)
conditioned stimulus (CS) (p. 164)
conditioned response (CR) (p. 164)
acquisition (p. 166)
extinction (p. 167)
spontaneous recovery (p. 167)
generalization (p. 167)
discrimination (p. 168)
biological preparedness (p. 172)
operant conditioning (p. 175)
law of effect (p. 175)
operant behavior (p. 175)
reinforcer (p. 175)
punisher (p. 175)
overjustification effect (p. 176)
shaping (p. 178)
fixed interval schedule (FI) (p. 180)
variable interval schedule (VI) (p. 180)
fixed ratio schedule (FR) (p. 180)
variable ratio schedule (VR) (p. 180)
intermittent reinforcement (p. 181)
intermittent reinforcement effect (p. 182)
latent learning (p. 183)
cognitive map (p. 184)
observational learning (p. 187)
implicit learning (p. 190)

Critical Thinking Questions

1. In habituation, repeated or prolonged exposure to a stimulus that initially evoked a response results in a gradual reduction of that response.

How might psychologists use the concept of habituation to explain the fact that today's action movies tend to show much more graphic violence than movies of the 1980s, which in turn tended to show more graphic violence than movies of the 1950s?

2. Little Albert was exposed to the sight of a rat paired with a distressing loud noise; with repeated pairings of the rat and the noise, he began to show a CR to the rat—crying and trembling.

Many people break into a cold sweat at the mere sound of a dentist's drill. How might this reaction be explained as a conditioned emotional response? (Hint: Assuming that human babies aren't born with a natural fear of drill sounds, then the cold sweat is a learned response [CR]. What are the CS and US?)

3. In operant conditioning, a reinforcer is a stimulus or event that increases the likelihood of the behavior that led to it, and a punisher is a stimulus or event that decreases the likelihood of the behavior that led to it.

Suppose you are the mayor of a suburban town, and you want to institute some new policies to decrease the number of drivers who speed on residential streets. How might you use punishment to decrease the behavior you desire (speeding)? How might you use reinforcement to increase the behavior you desire (safe driving)? Based on the principles of operant conditioning you read about in this section, which approach do you think might be most fruitful?

4. In fixed ratio (FR) schedules, reinforcement is delivered after a specific number of responses have been made. In variable ratio (VR) schedules, reinforcement is delivered after an average number of responses on average. Both FR and VR are examples of intermittent reinforcement schedules, because only some responses are followed by reinforcement, and they are both more resistant to extinction than continuous reinforcement schedules, in which a reinforcement is delivered after every response.

Imagine you own an insurance company, and you want to encourage your salespeople to sell as much merchandise as possible. You decide to give them bonuses, based on the number of items sold. How might you set up a system of bonuses using an FR schedule? Using a VR schedule? Which system do you think would encourage your salespeople to work harder, in terms of making more sales?

5. Observational learning takes place when one individual watches and learns from the actions of others. By contrast, in classical conditioning, learning takes place when an individual directly experiences the consequences (US) associated with a stimulus or event (CS).

Monkeys can be classically conditioned to fear objects such as snakes or flowers, if those objects are paired with an aversive US, such as electric shock. Monkeys can also learn to fear snakes through observational learning, if they see another monkey reacting with fear to the sight of a snake. But monkeys cannot be trained to fear flowers through observational learning—no matter how many times they watch another monkey who has been conditioned to fear the same flower. How does the principle of biological preparedness account for this finding?

Answers to Summary Quizzes

Summary Quiz 6-1
1. b ; 2. a ; 3. c

Summary Quiz 6-2
4. d; 5. c; 6. a; 7. b

Summary Quiz 6-3
8. d; 9. a; 10. b; 11. a

Summary Quiz 6-4
12. c; 13. b; 14. d

Summary Quiz 6-5
15. b; 16. d; 17. c