CHAPTER

7

Learning

Biologically and psychologically, you have a lot in common with your pet cat, a lowly sea slug, and a tap-dancing chicken. The biological resemblance is evident in the physiology you share with these creatures and the overlap in genetic patterns. But the psychological resemblance is also crucial, and includes an important set of similarities in the way all four of you detect relationships in the world and adjust your behavior accordingly. That is, there's a significant resemblance in the way that you and many other creatures learn.

In simple creatures and complex ones, learning comes in several forms. Both you and the cat look up when your roommate sneezes. But during allergy season, when she's sneezing all the time, you and the cat both learn that sneezes are part of the normal acoustic environment and cease responding to every "a-choo"—a form of learning called *habituation*.

Both you and the sea slug—and virtually every other animal—are also skilled at learning "what goes with what." If, just a couple of times, a light pressure on your skin is followed by a blast of cold air, soon you'll brace for the chill the moment you feel the pressure. Likewise, if the sea slug feels a mild poke and then, a moment later, a slight electric shock, the slug quickly learns to shift into a defensive posture as soon as it feels the poke. This form of learning is called *classical conditioning*.

And what about that dancing chicken? In nature, chickens scratch, peck, and waggle their heads, but they don't dance. Using *operant conditioning*, however, you can transform a chicken's natural antics into stomps, shuffles, and hops. First, identify a reward—such as corn—that your chicken likes. Then, through processes described in this chapter, you initially reward behaviors that vaguely resemble a tap-dance step—say, mere scratching—and then slowly shape these scratches into smooth moves.

You can use similar techniques to train one of your professors to lecture to only one side of the room—an exercise in behavior control cherished by generations of psychology students. In this case, the reward is not food, but the favor of your and your classmates' gaze. The procedure is simple: Conspire with your fellow students to look up with rapt attention whenever your professor addresses one side of the room, and to gaze downward and look bored whenever the prof turns in the other direction. After just a bit of this "training," your professor—like the chicken—will be producing the behavior you've selected.

These simple forms of learning—habituation, classical conditioning, and operant conditioning— are crucial for many aspects of our behavior and emotional responses, as they are for many other creatures on the planet. In this chapter, we'll look at how these types of learning proceed, and consider some of the biological mechanisms that—in you, the cat, the slug, or the chicken—make this learning possible. We'll then turn to some ways in which your learning differs from that of other creatures. For example, animals differ in how well they can learn just by watching their neighbors—and humans are especially skilled in this observational learning. At the same time, other creatures show feats of learning that humans can't match, such as easily learning to navigate in new environments. Why are some forms of learning shared by so many creatures, and why are other forms specific to just a few species? We'll tackle these questions and more in this chapter, as we consider what learning is and how it happens.

THE PERSPECTIVE OF LEARNING THEORY

What exactly is "learning"? In ordinary language, this term is applied to many different cases—the development of new skills, the acquisition of new knowledge, and more. According to some scholars, though, all learning involves the same basic processes.

As we discussed in Chapter 4, empiricist philosophers like John Locke (1632–1704) and George Berkeley (1685–1753) offered a simple account of how we come to understand our world. Perception, they argued, was massively influenced by learning; and (in their view) learning was just a matter of creating *associations* among ideas as a direct result of experience. Thus, for example, the sight of a stove might be followed immediately by the feeling of heat, creating an association between this sight and this feeling; and so you learned that stoves are hot. The sound of the word *flower* might be followed by the sight and smell of the flower, and so these ideas become associated.

But what about more complicated forms of learning? The answer, according to these philosophers, was easy: More complex learning simply involves more associations, built layer upon layer, so that complicated notions—and whole belief systems—are just the result of creating more and more links among individual ideas.

This is an appealing proposal, partly because similar conceptions have fared well in other sciences. Chemistry teaches us that complex molecules are built up by linking relatively simple atoms to each other, and then linking still other atoms to these, continuing in this way until huge combinations are created and the whole has properties that are often strikingly different from those of the individual components. Will a similar proposal work as our basic conception of learning?

A large number of researchers, called *learning theorists*, would answer this question with a firm yes, and in their research, they're guided by a striking implication of this view: If all learning depends on essentially the same mechanisms (i.e., mechanisms of association), then for research purposes it may not matter very much what forms of learning we choose to study, because the lessons we will draw from our research and the principles we will uncover should be the same whether we're scrutinizing simple cases of learning or far more complex ones.

Learning theorists also argue that it's sensible to focus experiments on simple organisms learning simple patterns; that way, the experiments will be easy to do, and the principles we're hoping to uncover should be immediately visible. Thus, rather than studying learning by asking how a college student masters calculus, we might choose to examine how less complicated organisms—rats, for example—form simple associations.

Does this strategy work? Are there uniform principles of learning that will emerge no matter what species and type of learning we examine? We know at the start that all animal species, as diverse as they are, have a lot in common biologically—in the structure of their nervous systems, for example, and in their evolutionary past. On this basis, perhaps all species do learn in essentially the same way; and, if so, we should be able to identify basic laws of learning that apply equally well to a dog learning to sit on command, a fish learning to navigate its way through a dense growth of algae, or a student learning to play a Mozart sonata.

As we'll see, this research strategy—an effort toward understanding *all* learning by studying *simple* learning—has led to many important discoveries. Indeed, in this chapter we'll discuss principles of learning that have amazing generality and apply to many species, situations, and types of behavior. These principles are also the basis for various useful techniques, including procedures used in treating phobias, techniques used to manage prison inmates' behavior, and more.

We'll also see that some forms of learning do not follow these general principles; so our overall discussion of learning will also need to take these distinctive forms of learning into account. Let's look at the data that underlie these important claims.

HABITUATION

habituation A decline in the response to a stimulus once the stimulus has become familiar.

dishabituation An increase in responsiveness when something novel is presented, following a series of presentations of something familiar.

Perhaps the simplest form of learning is **habituation**—the decline in an organism's response to a stimulus once the stimulus has become familiar. As a concrete case, imagine someone living in an apartment on a busy street. At first, he finds the traffic noises distracting and obnoxious. After he's lived in the apartment for a while, though, the noises bother him much less—and eventually he doesn't notice them at all. At this point, he's become habituated to the noises (Figure 7.1).

For the city dweller, habituation is obviously important. (Otherwise, with the traffic noise as it is, these people might never get any sleep!) But, more broadly, habituation produces a huge benefit: We want to pay attention to unfamiliar stimuli, because these may signal danger or indicate some unexpected opportunity. But, at the same time, we don't want to waste time scrutinizing every stimulus we run across. How, therefore, do we manage to be suitably selective? The answer is habituation; this simple form of learning essentially guarantees that we ignore inputs we're already familiar with and have found to be inconsequential, and focus instead on the novel ones.

Just as important as habituation is its opposite: **dishabituation**—an *increase* in responding, caused by a change in something familiar. Thus, the city dweller who's

7.1 Habituation Habituation has both benefits and costs. Thanks to habituation, the bison of Wyoming's Yellowstone Park have grown accustomed to the automobile traffic. This is helpful for the bison—but may be less so for the traffic. Similarly, many city residents have become habituated to the sight of homeless on the street—a sad fact that may undermine people's motivation to help the homeless.

seemingly oblivious to traffic noise will notice if the noise suddenly stops. Likewise, imagine an office worker who has finally gotten used to the humming of the building's light fixtures. Despite her apparent success in ignoring this humming, she's likely to notice immediately when, on one remarkable day, the humming ceases.

Dishabituation is obviously important because a change in stimulation often brings important news about the world. If the birds in the nearby trees suddenly stop chirping, is it because they've detected a predator? If so, the deer grazing in the meadow want to know about this. If the sound of the brook abruptly grows louder, has the water level suddenly increased? This, too, is certainly worth investigating. Thus, dishabituation serves the function of calling attention to newly arriving—and potentially useful—information, just as habituation serves the function of helping you ignore old news.

These simple forms of learning also provide powerful *research tools* for investigators. As one example, consider that adult speakers of Japanese have trouble hearing the distinction, obvious to English speakers, between the words *red* and *led*. This is because there's no equivalent distinction between these sounds in the Japanese language. But how did this perceptual difference between English and Japanese speakers come to be? In one experiment, 4-month-old Japanese infants heard the sound "la, la, la" repeated over and over. At first this sound stream caught their attention, but they soon habituated and stopped responding to the sounds. Then the researcher changed the sound to repetitions of "ra, ra, ra." Would the infants notice the change? In fact, they did: The infants showed immediate dishabituation and once again oriented to the sound. Apparently, Japanese infants heard the distinction between these sounds as readily as infants in English-speaking countries did. This seems, therefore, to be a case of "use it or lose it": All infants, no matter where they're born, can hear this acoustic difference. However, if the sound difference is not relevant to the language in the infant's surroundings (as, for example, in the case of Japanese), the infants stop paying attention to the distinction. By the time they're 12 months old, they've completely lost the ability to tell these sounds apart (for more on this finding, see Chapter 10). This is on its own an important finding, one that sheds light on how perception can change as a result of experience and helps us understand one aspect of language learning. But, for our purposes here, notice that this research relies on habituation and dishabituation as a means of finding out what sounds *different* to an infant and what sounds the *same*. Thus, these simple phenomena of learning provide a straightforward way to explore the perceptual capabilities of a very young (preverbal) child.

CLASSICAL CONDITIONING

Habituation is important, but it tells the organism about only a single stimulus—is the stimulus novel and so worth exploring, or familiar and therefore safe to ignore? Other forms of learning provide more information and, in particular, provide the organism with information about the *relationships* among events in the world.

Relationships come in many varieties. One event might *cause* another; an action might *prevent* some outcome; a certain circumstance might *magnify* an experience; and so on. Overall, though, relationships can usually be understood in terms of associations: Your dog learns to associate the sound of your footsteps with the possibility of a treat; you have learned to associate thunder with lightning; the farmer's cows learn to associate a certain time of day with milking. The importance of these associations was, as we've seen, highlighted by the empiricist philosophers, but the experimental study of these associations did not begin until the end of the late 1800s, when the work of the Russian physiologist Ivan Petrovich Pavlov (1849–1936) made a major contribution (Figure 7.2).

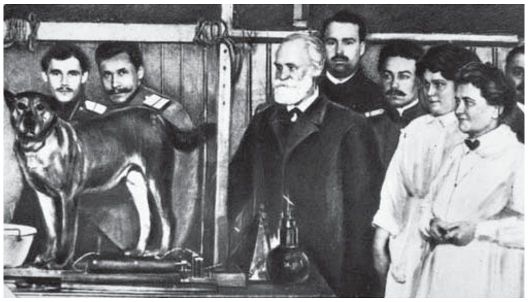

7.2 Ivan Petrovich Pavlov (1849–1936) Pavlov (center) in his laboratory, with some colleagues and his experimental subject.

unconditioned response (UR) A response elicited by an unconditioned stimulus without prior training.

unconditioned stimulus (US) A stimulus that reliably triggers a particular response without prior training.

Pavlov and the Conditioned Response

Pavlov's early work, for which he earned the Nobel Prize in medicine in 1904, was not in psychology. Instead, his research was concerned with digestive physiology, and many of his laboratory studies focused on the secretion of saliva in dogs. Pavlov knew from the start that salivation is triggered whenever food (especially dry food) is placed in the mouth. During his experiments, however, a new fact emerged: Salivation could be set off by a range of other stimuli as well, including stimuli that were at first totally neutral. Dogs that had been in the laboratory for a while would salivate in response to the mere sight of meat, or the sight of the dish that ordinarily held the meat, or even the sight of the person who usually brought the meat. Pavlov was intrigued by these effects because he realized that in these cases, the organism seemed to be developing new reflexes and changing its behavior in a fashion directly shaped by learning. He decided to refocus his research program to study this learning.

In his experiments, Pavlov created simple patterns for the animal to detect. For example, he would ring a bell and then give the animal food. Then, after a short wait, he would present another pair of stimuli: bell, then food. After another wait, he presented yet another pairing: bell, then food. After several such pairings, Pavlov observed what happened if the bell was sounded alone, without any food being given (Pavlov, 1927; Figure 7.3). The result was clear: The dog salivated in response to the bell.

To describe this pattern, Pavlov distinguished two types of responses: An **unconditioned response** (**UR**) was a biologically determined reflex, triggered by a certain stimulus independent of any learning. In Pavlov's terms, the trigger for an unconditioned response was an **unconditioned stimulus** (**US**). In the procedure described, the unconditioned stimulus (the US) is food in the animal's mouth; the unconditioned response (the UR) is salivation. The linkage that makes the US trigger a UR is something the animal brings into the situation,

7.3 Apparatus for salivary conditioning Here is an early version of Pavlov's apparatus for classical conditioning of the salivary response. The dog was held in a harness; sounds or lights acted as conditioned stimuli (CS), while meat powder in a dish was the unconditioned stimulus (US). The conditioned response (CR) was assessed with the aid of a tube connected to an opening in one of the animal's salivary glands.

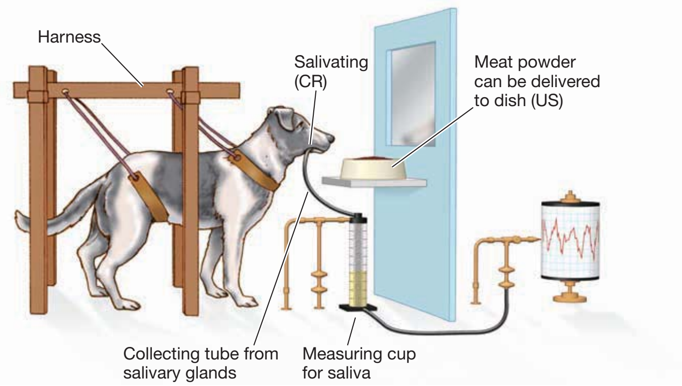

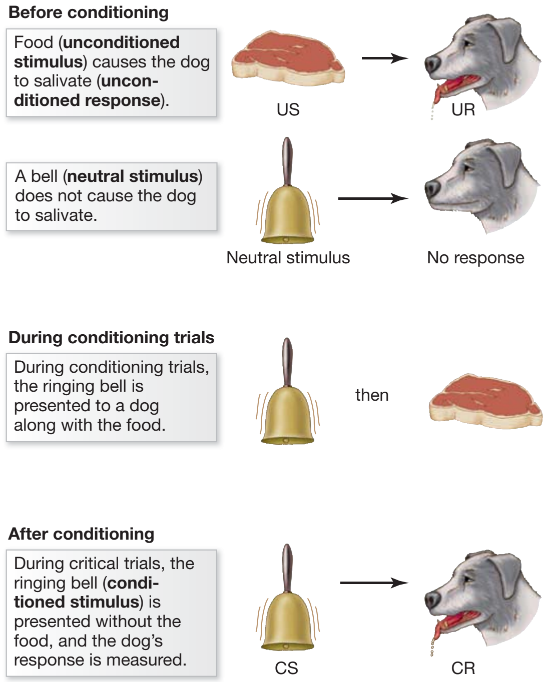

7.4 Relationships between CS, US, CR, and UR in classical conditioning

and so (in Pavlov's terms) is not a product of the learning process called "conditioning"; that's why the stimulus and response are said to be unconditioned.

The second type of response is a **conditioned response (CR)**, and it is a product of learning. Like the UR, the CR is triggered by a specific stimulus, but it's a stimulus that was neutral at the start of learning. In our example, this neutral stimulus is the bell, and it came to elicit the CR (salivation) only after several presentations in which this stimulus was followed by the US (food in the mouth). In Pavlov's terms, the bell is a **conditioned stimulus (CS)**—a stimulus that's initially neutral but becomes associated with the US during the experiment.

The relationships between US and UR, CS and CR, are summarized in Figure 7.4 and form the basis of the learning studied by Pavlov. In his honor, this type of learning is sometimes called *Pavlovian conditioning,* but it's more commonly known as **classical conditioning.**

Early research on classical conditioning focused on one conditioned response—salivation by dogs—and a narrow range of conditioned stimuli (the sound of bells—or in other experiments, the ticking of metronomes). Subsequent research, however, has made it plain that this form of learning occurs in a remarkable range of species and circumstances. Indeed, classical conditioning can be documented not just in humans but in species as diverse as ants and anteaters, cats and cockroaches, wolves and worms. By using the appropriate US, researchers have conditioned crabs to twitch their tail spines, fish to thrash about, and octopuses to change color. Responses conditioned in studies with humans include changes in heart rate or blood pressure (where the US is typically a loud noise or rap on the knee) and the reflexive eye blink (using a US of a puff of air on the open eye).

Outside of the laboratory, classical conditioning touches many aspects of our lives. We all tend to feel hungry at mealtime and less so in between; part of the reason is a conditioning process in which the CS is a particular time of day and the US is the presentation of food (which normally is paired with that time of day). Our emotional responses to certain songs, or certain smells, or even certain social situations can be understood in similar terms, and the response is likely to be the result of some previous pairing between these stimuli and some emotional experience. This type of learning is, for example, a plausible basis for some forms of anxiety as well as some phobias (Figure 7.5). Yet another example is sexual arousal, which can often be produced by an initially neutral word or gesture that has—through learning—acquired an erotic association. Clearly, then, classical conditioning is a process with wide application and great importance.

conditioned response (CR) A response elicited by an initially neutral stimulus—the conditioned stimulus (CS)—after it has been paired repeatedly with an unconditioned stimulus (US).

conditioned stimulus (CS) An initially neutral stimulus that comes to elicit a new response due to pairings with the unconditioned stimulus.

classical conditioning A form of learning in which one stimulus is paired with another so that the organism learns a relationship between the stimuli.

The Major Phenomena of Classical Conditioning

As we've seen, classical conditioning can be described in terms of a specific procedure: We need a US and UR that are biologically linked, and then we need a sequence of presentations in which the CS is paired with the US. But we can also describe this form of conditioning in terms of a set of properties that usually characterize this type of learning.

7.5 SCIENTIFIC METHOD: Can classical conditioning establish an emotional response?

Method

1. "Little Albert"—a normal, 9-month-old infant—was presented with a white rat.

2. When Albert reached out to touch the creature, a researcher struck a steel bar with a hammer, producing a loud, startling noise.

3. This pairing (rat + loud noise) was repeated several times.

Results: Little Albert showed intense fear the moment the rat came into view. Albert also exhibited fear of other furry animals, like this rabbit.

CONCLUSION: Classical conditioning can establish strong fearful responses. This may provide insight into how some phobias develop.

SOURCE STUDY: Watson & Rayner, 1920

ACQUISITION OF CONDITIONED RESPONSES

At the start of a conditioning procedure, the conditioned stimulus (CS) does not elicit the conditioned response (CR). In fact, the CS may elicit no reaction at all beyond a general stop-and-look response that organisms produce whenever a new stimulus appears. But after several pairings with the unconditioned stimulus (US), things change, so that a previously neutral CS (say, the sound of a bell) now elicits a CR (salivation).

Let's emphasize, however, that learning doesn't take place all at once. Instead, the learning is gradual, and the strength of the CR slowly grows as the animal experiences more and more pairings of CS and US. This pattern is evident in the data shown in Figure 7.6. Once the CS-US relationship is solidly established, though, the CS can be used in other procedures to establish other conditioned stimuli. As one example, by using meat powder as the US, we can first condition a dog to salivate whenever it sees a light. Once this is done, we can sound a bell and follow that by the light, without ever introducing the food. After enough of these pairings, the bell itself will trigger salivation. In this setting, the bell has become a signal for the light, which we've already established as a signal for the appearance of food. This sequence is called **second-order conditioning**—a procedure in which a neutral stimulus (here, the bell) is paired with some already established CS (like the light), as shown in Figure 7.7.

Second-order conditioning considerably extends the power and importance of classical conditioning. For example, the sight of your dentist is often paired with the discomfort of feeling her drill; as a result, the sight of the dentist (the CS in this case) might become fearful. But other stimuli are in turn associated with the sight of the

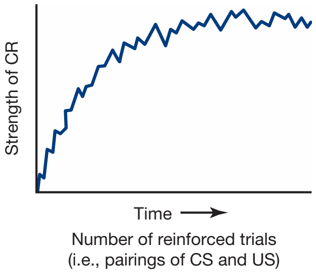

7.6 An idealized learning curve This graph plots CR strength against the number of reinforced (paired) trials. The CR gradually becomes stronger as the number of learning trials increases—but each trial adds less strength than the trial just before it.

second-order conditioning A form of learning in which a neutral stimulus is first made meaningful through classical conditioning. Then, that stimulus (the CS) is paired with a new, neutral stimulus until the new stimulus also elicits the conditioned response.

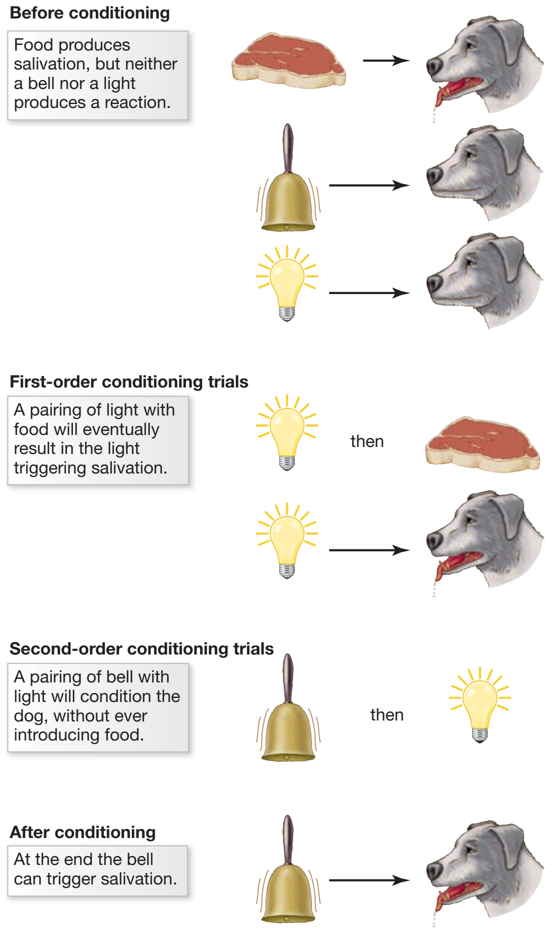

7.7 Second-order conditioning Before conditioning, neither the bell nor the light trigger salivation. During *first*-order conditioning, the light is paired with meat, and soon presentation of the light can trigger salivation. During *second*-order conditioning, the bell is paired with the light. As a result, presentation of the bell alone (which has never been paired with meat) will elicit salivation.

dentist—the sight of her office, the sound of her voice, the word *dentist,* and more. Through second-order conditioning, these stimuli, too, can become fearful—potentially leading to a fear of all things related to dentistry. In this way, second-order conditioning can produce widespread effects that, as in this example, can sometimes lead to the highly disruptive fears we call *phobias* (e.g., Gewirtz & Davis, 2000). More generally, though, mechanisms like higher-order conditioning allow the learning process to play a substantial role in shaping key aspects of our emotional lives.

EXTINCTION

Classical conditioning can have considerable adaptive value. Imagine a mouse that has several times seen the cat resting on a kitchen chair. It would serve the mouse well to learn about this association between the cat and a particular location; that way, the mouse will likely feel afraid whenever it nears the kitchen. This fear, in turn, will probably lead the mouse to avoid that room—a habit that could save the mouse's life!

At the same time, it would be unfortunate for the mouse if this association, once established, could never be undone. The cat might lose interest in that resting place or leave the household altogether. Either way, it would be useful for the mouse to lose its fearful response to the kitchen so it can return there to forage for food.

All is well for this mouse, though, because the effects of classical conditioning can be undone through a sequence of events similar to those that established the conditioning in the first place. Pavlov demonstrated that the CR will gradually disappear if the CS is presented several times by itself—that is, without the US. For example, repeated pairings of light plus a blast of cold air will create a conditioned response, so that the animal will shiver (the CR) whenever the light (the CS) is presented. But if the light is then presented several times on its own, the shivering response will be *extinguished.* **Extinction** is the undoing of a previously learned response so that the response is no longer produced (Figure 7.8).

Let's be clear, though, that extinction is not just the result of an animal forgetting what it learned earlier. Of course, animals (including humans) do eventually forget things they once learned, but that is not what's going on in extinction. This point is evident, for example, in the speed of extinction. As Figure 7.8 shows, a response can be extinguished in just a half-dozen trials over a period of only a few minutes. In contrast, forgetting is far slower: To demonstrate this, we can condition an animal, then leave it alone for several weeks, and then test it by presenting the CS. In this cicumstance, we have arranged for no extinction trials, but we have provided an opportunity for forgetting. The result of this procedure is clear: Even after a substantial delay, the animal is likely to exhibit a full-blown conditioned response (B. Schwartz, Wasserman, & Robbins, 2005). It seems, then, that classically conditioned responses are forgotten only very slowly.

The difference between extinction and forgetting is also clear in another procedure. First we condition an animal by repeated pairings of CS and US; then we extinguish the learning by presenting the CS on its own. In a third step, we *recondition* the same animal—by presenting some more learning trials, just like those in the first step of the procedure. What happens? The reconditioning usually takes much less time than the

extinction The weakening of a learned response that is produced if a conditioned stimulus is now repeatedly presented without the unconditioned stimulus.

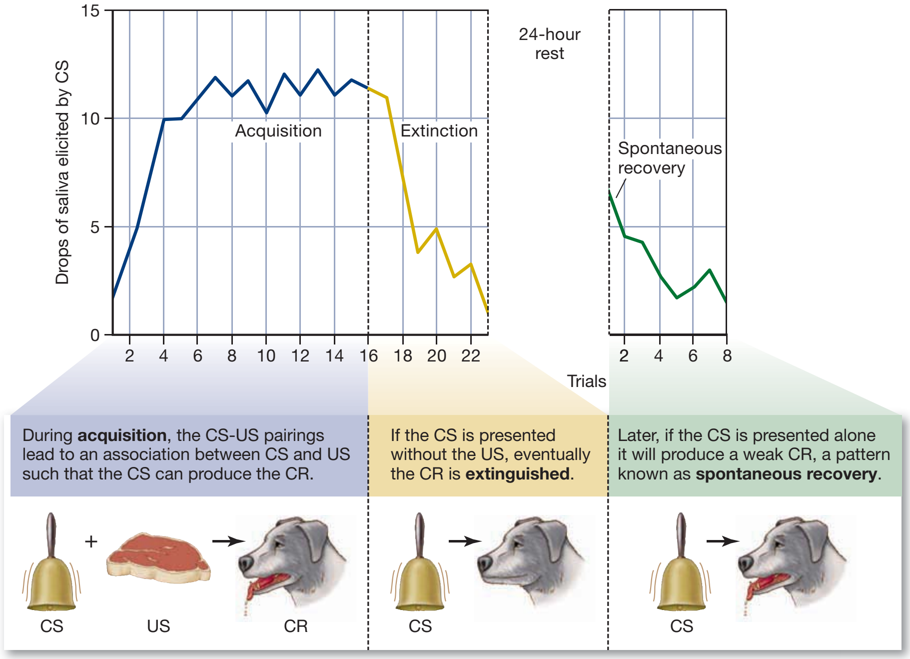

7.8 Extinction of a classically conditioned response The figure shows the decrease in the amount of saliva secreted (the CR) with increasing number of extinction trials—that is, trials on which the CS is presented without the US. However, if the animal then spends a little time away from the conditioning apparatus, the CR will reappear the next time the animal encounters the CS—a pattern known as spontaneous recovery.

initial conditioning did. The speed of relearning, in other words, is faster than the original rate of learning. Apparently, then, extinction doesn't "erase" the original learning and return the animal to its original naive state. Instead, the animal still has some memory of the learning, and this gives it a head start in the reconditioning trials.

We can draw similar conclusions about extinction from the phenomenon of **spontaneous recovery.** This phenomenon is observed in animals that have been through an extinction procedure and then left alone for a rest interval. After this rest period, the CS is again presented, and now the CS often elicits the CR—even though the CR was fully extinguished earlier (see Figure 7.8).

spontaneous recovery The reappearance of an extinguished response after a period in which no further conditioning trials have been presented.

According to one view of this effect, the extinction trials lead the animal to recognize that a once informative stimulus is no longer informative. The bell initially signaled that food would be coming soon; but now, the animal learns, the bell signals nothing. However, the animal still remembers that the bell was once informative; so when a new experimental session begins, the animal checks to see whether the bell will again be informative in this new setting. Thus, the animal resumes responding to the bell, producing the result we call spontaneous recovery (Robbins, 1990).

Like all aspects of conditioning, spontaneous recovery can easily be observed outside of the laboratory and in humans. For example, various anxiety disorders are often treated via *exposure therapy*—a process modeled after the extinction procedure. In this process, the person is repeatedly exposed to the specific stimulus or the particular situation that has, for that person, been a source of anxiety—heights, say, or enclosed spaces, or the sight of a snake. (For more on this procedure, see Chapter 17.) During these exposures, the person is kept safe and comfortable—and so there's no fearful US associated with the CS. As we'd expect, this sequence of events leads to extinction of the CR (the feelings of anxiety)—and with each exposure, the person feels less and less anxious.

When exposure therapy ends, however, people often relapse and again become anxious when exposed to the phobic stimulus. This relapse is not a sign that the therapy has failed. It's simply an example of spontaneous recovery of a CR—a sign that more treatment is needed to eliminate the anxiety.

stimulus generalization The tendency for stimuli similar to those used during learning to elicit a reaction similar to the learned response.

discrimination An aspect of learning in which the organism learns to respond differently to stimuli that have been associated with a US (or reinforcement), and stimuli that have not.

GENERALIZATION

In Pavlov's early experiments, animals were trained with a particular CS—the sound of a bell or metronome, for example—and then later tested with that same stimulus. But Pavlov understood that life outside the lab is more complicated. The master's voice may always signal food, but his tone of voice varies from one occasion to the next. The sight of an apple tree may well signal the availability of fruit, but apple trees differ in size and shape. Because of these variations, animals must be able to respond to stimuli that aren't identical to the original CS; otherwise, the animals may obtain no benefit from their earlier learning.

It's not surprising, therefore, that animals show a pattern called **stimulus generalization**—that is, they respond to a range of stimuli, provided that these stimuli are similar enough to the original CS. Here's an example: A dog might be conditioned to respond to a tone of a particular pitch. When tested later on, that dog will respond most strongly if the test tone is that same pitch. But the dog will also respond, although a bit less strongly, to a tone a few notes higher. The dog will also respond to an even higher tone, but the response will be weaker still. In general, the greater the difference between the new stimulus and the original CS, the weaker the CR will be. Figure 7.9 illustrates this pattern, called a *generalization gradient*. The peak of the gradient (the strongest response) is typically found when the test stimulus is identical to the conditioned stimulus used in training. As the stimuli become less like the original CS, the response gets weaker and weaker (so the curve gets lower and lower).

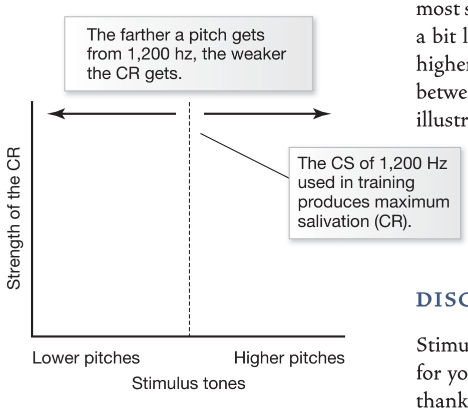

7.9 Generalization gradient of a classically conditioned response This figure shows the generalization of a conditioned blinking response in rabbits. The CS during the initial conditioning was a tone of 1,200 hertz, and the US was electric shock. After the conditioned response to the original CS was well established, generalization was measured by presenting various test stimuli and noting the percentage of trials in which the animals produced the CR.

DISCRIMINATION

Stimulus generalization is obviously beneficial, but it can be carried too far. It's sensible for you to feel fear when someone wearing an angry expression approaches you; and thanks to generalization, you'll feel this fear even if the person's face is a bit different from angry faces you've seen in the past. But if you generalize too much, you might feel fear in response to other facial expressions—and so end up being afraid in many social situations. What you need to do, therefore, is **discriminate**—respond in a way that's guided by the stimuli in your view.

The phenomenon of discrimination is easy to demonstrate. An experiment might, for example, use a loud boat horn as the US; this stimulus reliably produces a startle response when it's sounded. In the first part of the procedure, a red warning light is paired with the boat horn, and after a few pairings, this stimulus will reliably produce a conditioned response—you'll tense your muscles whenever the light appears. Once this CR is established, we proceed to the next step: Now we intersperse trials pairing the red light + horn with other trials in which a new stimulus—say, an orange light—is presented with no US (no boat horn). In this setup, you're likely to generalize at first by tensing up in response to both the red light (technically referred to as the CS$^+$) and the orange light (the CS$^-$). As the training continues, however—trials pairing the red light + horn, mixed with trials presenting the orange light alone—you'll learn to discriminate and will cringe only when you see the red light.

In a discrimination procedure you learn, of course, that the CS$^+$ signals the approach of the US. What about the CS$^-$? In our example, we might think you're learning nothing about the orange light because it's never followed by the horn, and for that matter never followed by any kind of US. Even so, the CS$^-$ (again, the orange light) does provide information. When it arrives, this signals a period in which the US is likely *not* to arrive. If the US is the horn, then the CS$^-$ indicates that the loud noise is not coming soon. If the US is food, then the CS$^-$ signals the start of a period in which no food will be available.

Essentially, then, CS$^-$ takes on a meaning opposite to that of the CS$^+$. It means "no noise," or "no food," or, in general, "no US." Correspondingly, the animal's response to the CS$^-$ tends to be the opposite of its response to the CS$^+$. If the US is a noise blast, then the CS$^+$ elicits fear and the CS$^-$ seems to inhibit fear—and so the animal is *calmer* in the presence of the CS$^-$ than it would be otherwise. If the US is food, then the CS$^+$ elicits salivation and the CS$^-$ causes the animal to salivate less than it ordinarily would. Overall, the CS$^-$ takes on the role of **inhibitor:** Whatever the response produced by the CS$^+$, the CS$^-$ makes that response less likely.

inhibitor A stimulus signaling that an event is not coming, which elicits a response opposite to the one that the event usually elicits.

THE CS AS A "SIGNAL"

In describing discrimination, it seems natural to speak about the CS$^+$ and CS$^-$ as though they were *signals* for the animal, providing information about things to come. And in fact this way of thinking about the CS helps us understand several aspects of conditioning, including why the rate at which conditioning develops depends on how the CS and US are related to each other in time.

Conditioning happens most efficiently when the CS precedes the US by some optimum interval—usually a half-second or so, or perhaps a few seconds at most. If the interval between the CS and US is increased beyond this optimum, the effectiveness of the pairing drops sharply. But we also don't want the interval between these stimuli to be *too short.* In fact, presenting the CS and US simultaneously is usually ineffective in establishing an association—and the backward procedure, presenting the US *before* the CS, is even worse (Rescorla, 1988; Figure 7.10).

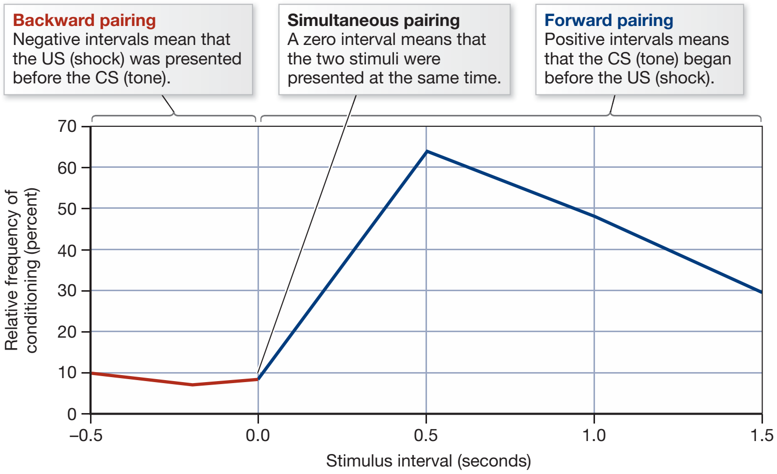

7.10 The CS-US interval in classical conditioning The figure shows the results of a study of the effectiveness of various CS-US intervals in humans. The CR was a finger withdrawal response, the CS a tone, and the US an electric shock. The time between CS and US is plotted on the horizontal axis; *negative* values indicates that the CS arrived after the US. The vertical axis indicates the strength of conditioning.

These facts make perfect sense if we think of the CS as a signal warning the organism that it should prepare itself for the upcoming US. To see why, imagine a mountain road that has a dangerous hairpin turn. How should drivers be warned about this turn? The best warning would be a "Caution" sign just before the turn (analogous to forward pairing with a short CS-US interval). This sign would be informative, and—crucially—would allow the driver enough time to prepare for the upcoming maneuver. But it's important not to place the sign too far ahead of the turn. Suppose a Caution sign is posted 100 miles before the turn (forward pairing with an extremely long CS-US interval). In that case the driver might not connect the sign with what it signifies—or just as bad, he might have forgotten about the sign by the time he reaches the curve. Things would be worse still, though, if the sign were prominently displayed right in the middle of the hairpin turn (simultaneous pairing), because now the warning comes too late to be of any use. Worst of all, the driver might suspect a degree of malevolence if he discovered the sign placed on the road just beyond the turn (backward pairing), although he'd probably be grateful that he didn't find the sign at the bottom of the ravine.

CONTINGENCY

The CS's role as a signal also has a crucial implication for what *produces* classical conditioning—that is, what the relationship between the CS and the US must be for learning to occur. To understand the issue, consider a dog in a conditioning experiment. Several times, it has heard a metronome and, a moment later, received some food powder. But many other stimuli were also present. At the same time it heard the metronome, the dog heard some doors slamming and some voices in the background. It saw the laboratory walls and the light fixtures hanging from the ceiling. At that moment, it could also feel various bodily sensations. What, therefore, should the dog learn? If it relies on mere contiguity, it will learn to associate the food powder with all of these stimuli—metronomes, light fixtures, and everything else on the scene—since they were all present when the US was introduced.

Notice, though, that many of these stimuli—even if contiguous with the US—give no information about the US. The light fixtures, for example, were on the scene just before the food powder arrived; but they were also on the scene during the many minutes when no food was on its way. So the sight of the light fixtures can't signal that food is coming soon, because the presence of the light fixtures has just as often conveyed the opposite message. Likewise for most of the sounds in the laboratory; they were present just before the food arrived, but they were also present during minutes without food. Therefore, none of these stimuli will help the animal predict when food is coming and when it's not.

To predict the US's arrival, the dog needs some event that reliably occurs when food is about to appear and doesn't occur otherwise. And, of course, the metronome beat in our example is the only stimulus that satisfies this requirement, since it never beats in the intervals between trials when food is not presented. Therefore, if the animal hears the metronome, it's a safe bet that food is on its way. If the animal cares about signaling, it should learn about the metronome and not about these other stimuli, even though they were all contiguous with the target event.

Are animals sensitive to these patterns? Said differently, what is it that leads to classical conditioning? Is it *contiguity*—the fact that the CS and US arrive close to each other in time? Or is it *contingency*—the fact that the CS provides information about the US's arrival? It turns out that contingency is the key, and in fact a CR is acquired only when the CS is informative about things to come.

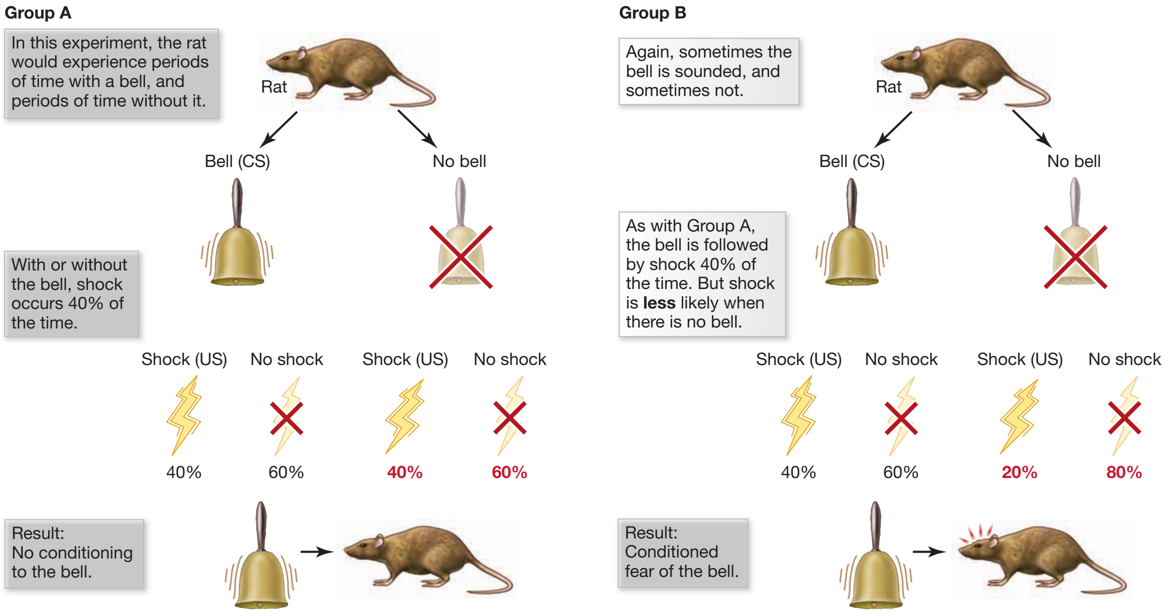

7.11 The effect of contingency on classical conditioning For both groups, there's only a 40% chance that bells will be followed by shock. However, for Group B, shock is *less likely* when no bell is sounded, and, for this group, the bell becomes a fearful stimulus.

In one experiment, rats were exposed to various combinations of a bell (CS) and a shock (US) (Figure 7.11). The bell was never a perfect predictor of the shock, but it did signal that shock was likely to arrive soon. Specifically, presentation of the bell signaled a 40% chance that a shock was about to arrive.

For some of the rats in this experiment (Group A in the figure), shocks also arrived 40% of the time without any warning. For these rats, therefore, the bell provided no information. The likelihood of a shock following the bell was exactly the same as the likelihood of shock in general. And in fact this situation led to no conditioning; instead, the rats simply learned to ignore the tone.

For another group of rats (Group B in the figure), the bell still signaled a 40% chance of shock, and shocks still arrived occasionally without warning. For these animals, though, the likelihood of a shock was only 20% when there was no bell. So in this setting, the bell was an imperfect predictor but it did provide some information, because shock was more likely after the bell than otherwise. And, in this situation, the rats did develop a conditioned response—they became fearful whenever the bell was sounded.

Let's be clear that in this experiment, the two groups of rats experienced the same number of bell-shock pairings, and so the degree of contiguity between bell and shock was the same for both groups. What differed between the groups, though, was whether the bell was informative or not—and it's this information value, not the contiguity, that matters for conditioning. Notice also that the bell was never a perfect predictor of shock: Bells were *not* followed by shock 60% of the time. Even so, conditioning was observed; apparently, an imperfect predictor is better than no predictor at all (Rescorla, 1967, 1988; Figure 7.12).

7.12 An imperfect predictor? The weather forecast is often wrong. Still, the chances of a hot day are greater if the forecast is for heat rather than snow. On that basis, we pay attention to the forecast, even if it's not always correct. Pavlovian conditioning seems to work the same way—and so a somewhat reliable signal seems to be better than no signal at all.

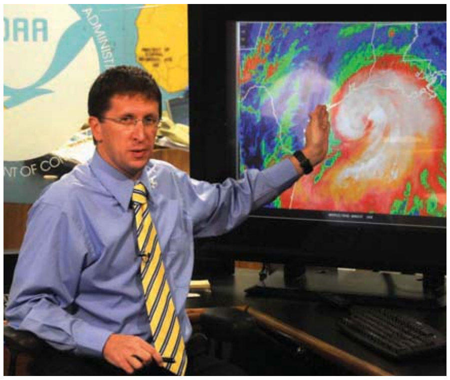

THE ABSENCE OF CONTINGENCY

The importance of contingency is also evident in another way—in an organism's reaction when there's no contingency at all. To explore this idea, let's imagine two different procedures. In both, the animal hears 40 presentations of a tone and receives 20 electric shocks. In the first procedure, these stimuli are presented randomly—and so there's no contingency, no relationship, between hearing the tone and receiving a shock. In the second procedure, the stimuli are arranged so that half of the tones are followed by shock and half are not, and shocks are never presented without a tone preceding them. In this setting, hearing the tone signals a 50% chance that a shock is about to arrive; in the absence of the tone, the chance of shock is zero.

Animals react very differently to these two procedures. In the first, the tone conveys no information—shock is just as likely with a tone as without. Not surprisingly, the tone in this situation does not become a fearful stimulus. Indeed, after just a few presentations of the tone—so that the novelty of this stimulus wears off—the animal's behavior doesn't change at all when the tone is sounded. In the second procedure, in contrast, the tone is informative, indicating that a shock quite likely is about to arrive. It's no wonder that in this case the animal soon shows a fear response whenever the tone is presented.

These two procedures also differ in another way. The first procedure—with its random arrangement of tones and shocks—is far more aversive to animals. This is clear, for example, if we monitor the animals' bodily state (e.g., their heart rates) during both procedures. The comparison tells us that the animals are more stressed in the noncontingent procedure. We can also, in effect, "ask" animals which of these procedures they find worse. We do this using a lab setup in which the animal can enter either of two chambers, one governed by the random procedure we've described and one governed by the contingent procedure. When given these options, animals reliably choose the second procedure.

What's going on here? In the second procedure, there is a clear "danger signal" for the animal (the tone), and this signal reliably produces fear. But this setup also provides a clear "safety signal"—namely, the *absence* of the tone. When the tone isn't sounded, the animal knows no shock is coming and it can relax. In the first (noncontingent) procedure, in contrast, there's never a danger signal, nothing to indicate when a shock is coming, and therefore no specific trigger for fear. But, in this setting, there's also no indication when the animal is safe. As a result, in this procedure the animal is constantly afraid and constantly on guard.

Results like these highlight the importance of contingency—the relations among stimuli that allow us to anticipate upcoming events. When there is some contingency (i.e., one event allows predictions about another event), animals learn this; and it seems to be crucial for classical conditioning. But when there's no contingency, animals learn this, too—and learn that their environment is unpredictable. These points are crucial for, say, dogs in a conditioning experiment, but they're no less important for humans. For example, think about why terrorist activities are so frightening: Terrorists hope to convey the message that they can strike you anytime, anywhere, so that you're never safe. It's this absence of contingency that makes terrorist threats so scary. Similarly—but on a more personal level—think about the distinction between fear and anxiety. According to some theorists, fear is a state triggered by a specific situation or object; anxiety, on the other hand, is chronic, has no object, and occurs in many situations. Some authors suggest that such unfocused anxiety is partly caused by unpredictability—that is, by an absence of safety signals (B. Schwartz et al., 2005; Seligman, 1975).

THE ROLE OF SURPRISE

In the experiments we've been describing, animals seem keenly sensitive to *comparisons among probabilities*. Will a dog learn to salivate in response to a beeper if the sound is followed by food only 30% of the time? The answer is yes, *if* food arrives less than 30% of the time when no beeper is sounded. Will a mouse shiver in response to a tone that, in the past, has been followed only half the time by a blast of cold air? Again, the answer is yes—provided that cold blasts arrive less than half the time without the warning tone.

This sensitivity to probabilities is, in fact, a widespread feature of classical conditioning and can be observed in creatures as diverse as humans and rabbits, pigeons and rats. In all cases, conditioning depends on whether the probability of the US *after* the CS is different from the probability of the US *without* the CS. But how is this possible? The dog and the mouse are obviously not standing by with calculators, computing these probabilities. For that matter, when we study classical conditioning in humans, our participants don't appear to be tallying up the various types of trials and computing the relevant ratios. How, then, are the test subjects influenced by these probabilities?

One proposal starts with the idea that—for any organism, human or otherwise—associations can provide a basis for *expectations,* and learning can then take the form of an *adjustment in expectations whenever a surprise occurs.* To see how this plays out, imagine an animal in a learning experiment. At the very start, there's no association between (say) the sound of the metronome and the delivery of food, and so the animal has no expectation that this CS will be followed by this US. When the food does arrive, therefore, it's a surprise; and this causes the animal to adjust its expectations. Specifically, the surprise leads the animal to a tentative expectation that, in the future, other metronome sounds might also be followed by food. This expectation will be weak at first, and so the next time this CS is followed by the US, the animal will still be a little bit surprised. This will cause another adjustment in the animal's expectation—and so an increase in the strength of association between CS and US.

This process will then be repeated over and over as the learning experiment continues. In each trial, the CS is presented; and this will trigger certain expectations in the animal for what's going to happen next. If these expectations are correct, then the arrival (or nonarrival) of the US will be no surprise, and so there's no reason to adjust the expectations (Figure 7.13). But if the expectations are wrong, then some adjustment is called for. If no US is expected but one arrives anyhow, this surprise will lead to a strengthening of the CS-US association and thus a stronger expectation the next time around. Conversely, if the US is expected but does not arrive, then the CS-US association will be weakened and so there will be a weaker expectation on the next trial.

If things continue in this way, trial after trial, the expectations will be adjusted each time they're out of step with actual events and left alone when they're correct. Put differently, the expectations will be tuned and retuned until they're quite accurate—fully in line with the circumstances. Thus, if the animal is in a situation in which a CS is followed by the US 90% of the time, the animal will end up with strong expectations for what's going to happen next, whenever it

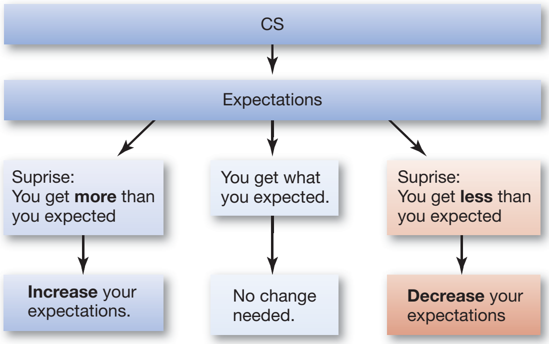

7.13 The role of expectations and surprise This figure shows the (automatic, unconscious) process through which expectations can be adjusted, trial by trial, in a classical conditioning experiment. The one complication not shown here is that bigger surprises (greater departures from expectations) will trigger larger adjustments; smaller surprises will lead to smaller adjustments.

experiences the CS. If the animal is in an environment in which a CS is followed by the US only 80% of the time, or 60%, the animal's expectations will be accordingly weaker.

Notice, then, that we can predict the animal's behavior—in particular, whether the animal will learn or not—by keeping track of the various probabilities in the animal's situation. But the animal has no direct knowledge about these probabilities. Instead, the animal has a set of expectations, based on its experience. The key, though, is that these expectations have been shaped, trial by trial, by an adjustment process that brings the expectations into line with reality, as shown in Figure 7.13. This is the mechanism through which the animal's behavior ends up fully in accord with the probabilities in its environment (Kamin, 1968; Rescorla & Wagner, 1972; but see R. Miller, Barnet, & Grahame, 1995; Pearce & Bouton, 2001).

Is this proposal correct? Is this how classical conditioning proceeds, with a trial-by-trial calibration of expectations? One way to find out is to scrutinize the role of surprise, which—in the account just sketched—plays a crucial role in learning. And in fact the importance of surprise is easily demonstrated. In studies of the **blocking effect**, animals are exposed to a three-part procedure (Figure 7.14). In stage 1, the animals—rats, for example—hear a hissing noise that's followed by a US (let's say, the sight of a sexually receptive mate). As one might expect, this noise becomes a CS for sexual arousal. In stage 2, the hissing sound is still followed by this same US; but now the hissing sound is reliably accompanied by another stimulus—a bright light. The sequence in this stage, therefore, is the "package" of hiss plus light, followed by the sight of the potential mate. Then, in stage 3, it's time for the crucial test: Now we present the light by itself and observe the rat's response.

In stage 2 of this procedure, the light was reliably followed by the US, and we might expect this to produce learning—so that the light will now trigger the CR. But that's not what happens. Instead, the rat doesn't respond to the light at all. This is because in stage 2 the light provided only redundant information. The rat already knew from the hissing noise that the US was about to be presented, and so it wasn't at all surprised

blocking effect A result showing that an animal learns nothing about a stimulus if the stimulus provides no new information.

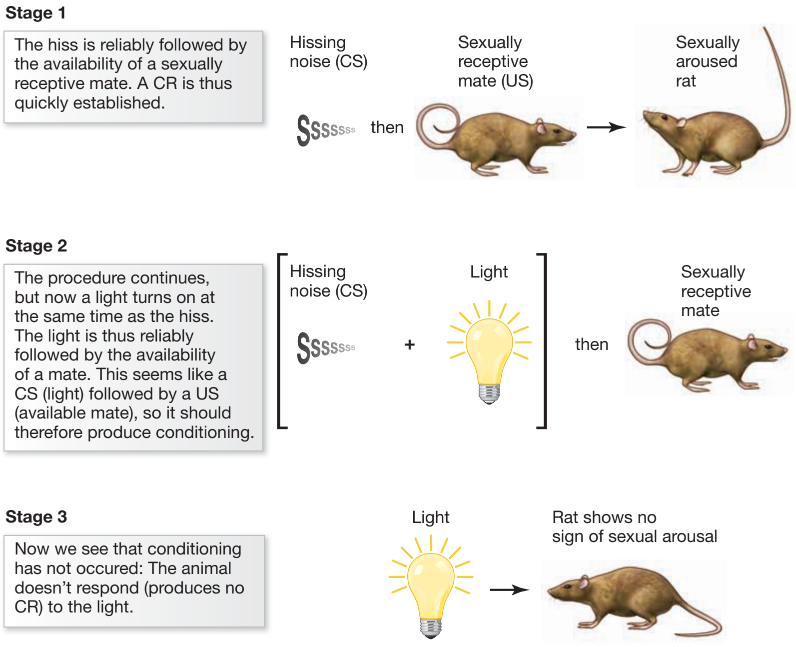
Stage 1

The hiss is reliably followed by the availability of a sexually receptive mate. A CR is thus quickly established.

Stage 2

The procedure continues, but now a light turns on at the same time as the hiss. The light is thus reliably followed by the availability of a mate. This seems like a CS (light) followed by a US (available mate), so it should therefore produce conditioning.

Stage 3

Now we see that conditioning has not occured: The animal doesn't respond (produces no CR) to the light.

7.14 The blocking effect procedure What prevented the conditioning in this sequence? In Stage 2, the light provided no new information, because the hiss told the animal that the US would soon arrive. Conditioning does not occur with uninformative stimuli!

when the US did arrive. As a result, the rat learned nothing about the light—clear confirmation of the proposal that learning does in fact depend on surprise and, more broadly, on the information value of the stimuli.

The blocking effect can easily be demonstrated in experiments with humans (Beckers, Miller, De Houwer, & Urushihara, 2006; Kruschke & Blair, 2000), and analogous phenomena can be documented outside of the laboratory. In the United States, for example, temperatures are routinely reported in degrees Fahrenheit—unlike the rest of the world, which uses the metric centigrade scale. For some years, weather forecasters tried to teach Americans the alternative scale by routinely describing the temperature with *both* scales: "The high tomorrow will be 75 degrees, or 24 centigrade." "Watch out—it's going to be cold tonight—in the teens, or –10 degrees centigrade." This effort failed—and Americans learned nothing about the metric scale. Why? Because the centigrade number was redundant with the Fahrenheit temperature, and so the centigrade number provided no information, no surprise. As a result, Americans ignored the number—just as a rat in the laboratory ignores the redundant (and therefore uninformative) light.

The Relationship between the CR and the UR

In Pavlov's original studies, dogs salivated both when they heard the CS and when they experienced the US. We might conclude from this that the conditioned response (CR) and unconditioned response (UR) are essentially the same behavior—and simply triggered by different inputs. In truth, though, the CR and UR are rarely identical and are sometimes quite different from each other.

Consider, for example, a rat that has been conditioned to associate a flashing light with electric shock. When the shock is actually presented, the animal jumps and squeals, and its heart beats faster; this is the UR, triggered by the shock itself (the US). When the animal sees the flashing light, though, its response (the CR) is different. The animal freezes and tenses its muscles, and its heartbeat slows. This is not an escape-from-shock reaction; instead, it's a display of fearful anticipation.

From a biological perspective, this result makes sense. A flashing light causes no physical damage, and so there's no need to spend energy escaping from the light. Likewise, an animal can't eat a tone that has been paired with food, so there's no point in opening its mouth and chewing in response to the tone. In these and many other cases, it's sensible for the animal to react somewhat differently to the CS than it does to the US.

So, what is the relationship between the conditioned response (CR) and the unconditioned response (UR)? We've already seen that animals interpret the CS as a signal indicating that the US is about to arrive. Plausibly, then, we can think of the CS as telling the animal to "Get ready!"—and, from this perspective, the CR is just the set of adjustments the animal makes in *preparation* for the US. If the sound of a bell has been reliably followed by food, then the sound now signals that the animal should moisten its mouth so it will be ready to eat when the food arrives. If the sight of a light has been followed by a shock, then the light is a signal that the animal should stop moving around and stay at "full alert" so it will be ready to jump as soon as the shock begins (Domjan, 2005; P. Holland, 1984; Hollis, 1984).

Does this preparation help the animal? It surely does. Evidence suggests, for example, that sexual activity is more likely to lead to offspring if a CS announced the imminent arrival of the US (the sight of a sexually receptive partner). Presumably, this is because the CS allowed the animal to prepare itself for mating (Domjan, 2005). Likewise, digestion is more efficient if a CS announced the imminent arrival of the US (food in the mouth); again, the CS allowed the animal to prepare to ingest and digest the food (Domjan, 2005; Woods & Ramsay, 2000).

CONDITIONING AND COMPENSATORY RESPONSES

Preparation for a US can take many forms. As a remarkable illustration, consider the situation of a heroin addict. The heroin itself is a potent stimulus with many biological effects: It decreases pain sensitivity and lifts the user's mood. It also causes other changes throughout the body, including drying out the mouth and various mucous membranes.

The first few times someone uses heroin, all of these effects are strong. If drug use continues, though, these effects diminish. This phenomenon, which we described in Chapter 6, is called drug tolerance—a decrease in the response to a drug, usually resulting from continued use. If the person wants to keep getting the same impact from the drug, they'll need a larger and larger dose to offset the effect of the tolerance.

Continued use of heroin also leads to *drug dependence* and drug *cravings*—an inability to function without the drug and an overwhelming desire for yet another dose, yet another injection. The cravings are accompanied by their own set of effects—including an *increased* sensitivity to pain, a depression in the person's mood, and an overproduction of fluid in the person's mouth and mucous membranes.

What's going on here? What produces drug tolerance and drug craving? Why does the craving take the form that it does? Answers to these questions turn out to include a key role for classical conditioning. To understand how this works, however, we need to begin with a point we'll discuss more fully in Chapter 12 where we consider the importance of *homeostasis.* This term refers to the remarkably stable environment that exists inside of our bodies—a body temperature that's kept at an almost constant level, a nearly constant pH in the bloodstream, a consistent level of glucose and oxygen in the blood, and so on.

A drug such as heroin changes the body's status in many ways. In other words, heroin disrupts homeostasis; and so, when this drug is present in the body, a range of mechanisms come into play, all seeking to restore homeostatic stability. These mechanisms involve many cellular and biochemical changes throughout the body—all designed to repair the "disruption" caused by the drug.

But, of course, *preventing* a problem is always preferable to solving the problem after it arises; *avoiding* disruption is better than repairing the disruption once it's in place. Rather than waiting until heroin disrupts the body's state and then responding to this disruption, it would be better if the body had a way of dealing with the heroin disruption *as it happens,* so that homeostasis is never lost in the first place. This is where classical conditioning enters the scene: Thanks to conditioning, the body begins a series of adjustments to offset the heroin's effects even before the drug arrives.

Let's put this in concrete terms. For a heroin user, the US is the drug itself; the UR is the body's natural response to the drug (Figure 7.15). The CS is complex and includes all the stimuli that signal the drug is about to arrive—the sight of the needle, thoughts about the drug, possibly the sights and smells of the physical environment in which the drug is injected, and so on. But what is the CR? Here as always, the CR will be a response that prepares the organism for the US, and so we need to ask: What is the right preparation for heroin? If the goal is homeostasis, then the CR should include a depression of mood in order to cancel out heroin's positive effects on mood; it should also include an increase in pain sensitivity to cancel out heroin's analgesic effects; it should include an increase in moisture in the mucous membranes to offset heroin's tendency to dry out these membranes. In point after point, the CR should simply be the opposite of the UR, so that the two will cancel each other out, leaving no overall effect—and thus preserving homeostasis. A CR like this is referred to as a **compensatory response**—one that "compensates" for the effects of the upcoming US.

Of course, when someone is exposed to heroin for the first time, there hasn't yet been an opportunity for learning. Therefore, the US (the heroin) will produce the UR

compensatory response A response that offsets the effects of the upcoming unconditioned stimulus.

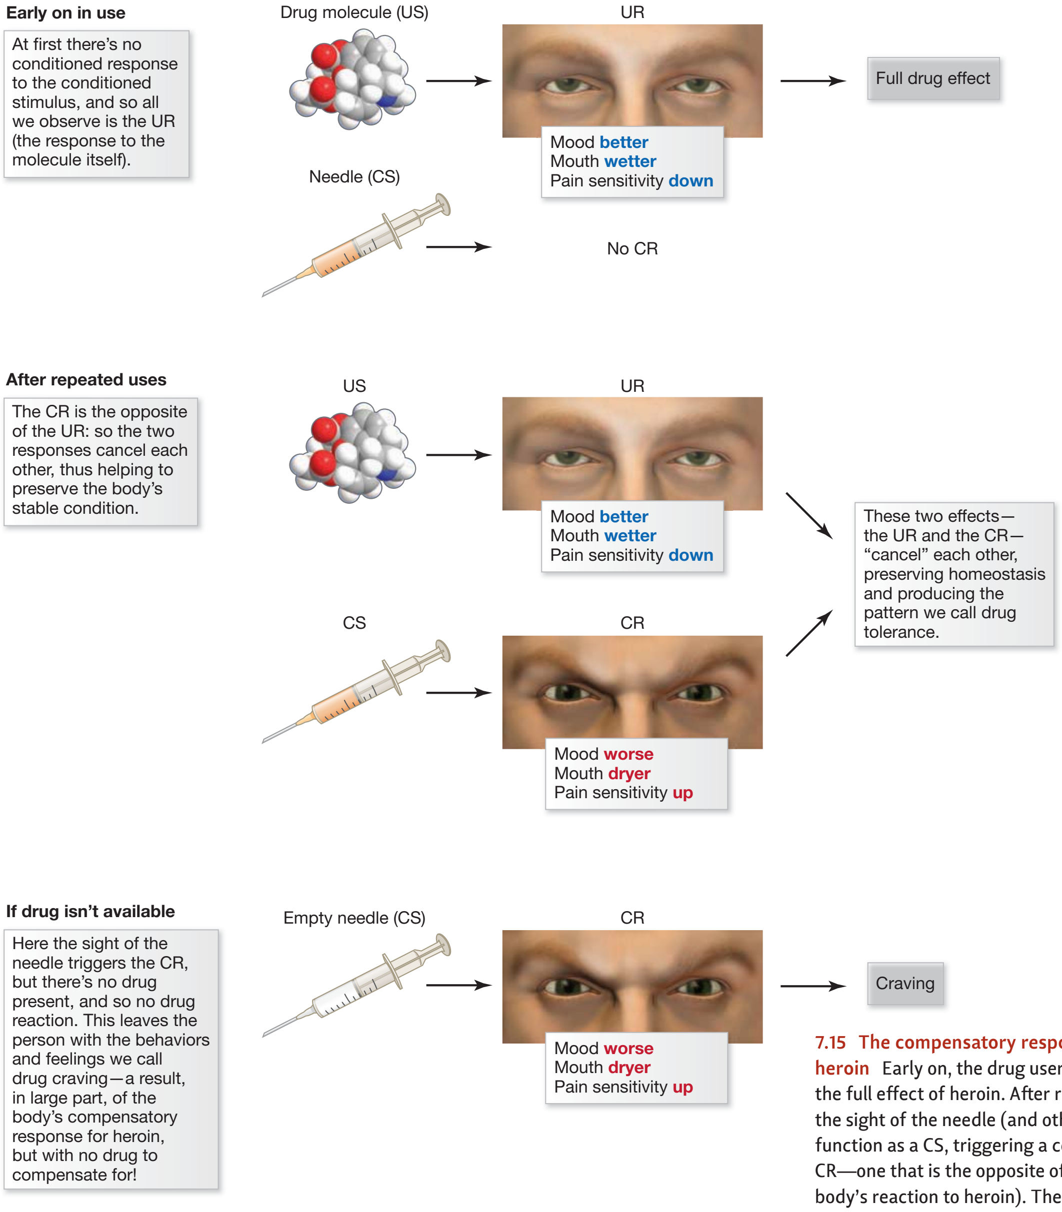

7.15 The compensatory response to heroin Early on, the drug user experiences the full effect of heroin. After repeated uses, the sight of the needle (and other stimuli) function as a CS, triggering a compensatory CR—one that is the opposite of the UR (the body's reaction to heroin). The compensatory response "cancels out" the UR, and so helps maintain homeostasis.

(the biological reaction to the drug); but there will be no CR (because no learning has taken place yet). As a result, the person will experience the drug's full effects, and no compensatory response will be in place (Figure 7.15, top panel). After repeated exposures to the drug, though, learning will have taken place. At this point the US still produces the UR; but now, in addition, the CS elicits the compensatory CR. This learned response, we've proposed, functions to "cancel out" the heroin effects and preserve homeostasis. As a result, the UR will be less evident—resulting in the diminished drug response that we call drug tolerance (Figure 7.15, middle panel).

What happens if an addict sees an empty hypodermic needle, or visits the place where he ordinarily buys or injects the drug, but no heroin is available (Figure 7.15, bottom panel)? In these cases the drug-associated signals all indicate that the drug will arrive soon, but then it never arrives—so there's a CS (the signals), but no US (no heroin). With the CS present, the CR will be produced; but with no US, there will be no UR. As a result, the person will experience the CR on its own—and the CR, we've proposed, is the opposite of the UR. Hence the person experiences the depression, the pain sensitivity, and so on—exactly the pattern referred to as drug craving.

Many experiments have confirmed these claims about heroin tolerance and craving. Among other points, the data indicate that tolerance shows not only the pattern of generalization and discrimination that we would expect with classical conditioning but also the familiar patterns of extinction and spontaneous recovery. In other words, tolerance shows the standard profile of classical conditioning, and this is strong support for the account we've just sketched. Moreover, related studies show a similar role for conditioning in tolerance observed with other drugs, including insulin, nicotine, caffeine, and amphetamines (Domjan, 2005; S. Siegel, 1977, 1983; S. Siegel & Allan, 1998; S. Siegel, Kim, & Sokolowska, 2003; Sokolowska, Siegel, & Kim, 2002). Overall, these points provide a powerful argument that the CR is indeed best understood as a *preparation for* the US; but it's important to note that this notion of preparation must be understood broadly.

INSTRUMENTAL CONDITIONING

instrumental conditioning A form of learning in which the participant receives a reinforcer only after performing the desired response, and thereby learns a relationship between the response and the reinforcer.

Habituation and classical conditioning are both general forms of learning that are relevant to many species (including, of course, humans) and many different responses (including a variety of overt behaviors, a range of subjective feelings, and a broad set of bodily responses). An equally important type of learning is **instrumental conditioning** (also called *operant conditioning*). This form of learning involves behaviors that aren't triggered automatically by some stimulus. Instead, instrumental conditioning is concerned with behaviors initiated by—and presumably under the control of—the organism itself. In other words, while classical conditioning essentially involves the creation of new reflexes, instrumental conditioning involves the learning of new voluntary behaviors.

Thorndike and the Law of Effect

The experimental study of instrumental conditioning began a century ago and was sparked by the debate over Darwin's theory of evolution by natural selection. Supporters of Darwin's theory emphasized the continuity among species, both living and extinct: Despite their apparent differences, a bird's wing, a whale's fin, and a human arm, for example, all have the same basic bone structure; this similarity makes it plausible that these diverse organisms all descended, by a series of incremental steps, from common ancestors. But opponents of Darwin's theory pointed to something they perceived as the crucial discontinuity among species: the human ability to think and reason—an ability they claimed animals did not share. Didn't this ability, unique to our species, require an altogether different (non-Darwinian) type of explanation?

In response, Darwin and his colleagues argued that there is, in fact, considerable continuity of mental prowess across the animal kingdom. Yes, humans are smarter in some ways than other species; but the differences might be smaller than they initially seem. In support of this idea, Darwinian naturalists collected stories about the intellectual achievements of various animals (Darwin, 1871). These stories painted a flattering picture, as in the reports of cunning cats that scattered breadcrumbs on the lawn to lure birds into their reach (Romanes, 1882). In many cases, however, it was hard to tell whether these reports were genuine or just bits of folklore. Even if they were genuine, it was unclear whether the reports had been polished by the loving touch of a proud pet owner. What was needed, therefore, was more objective and better documented research—research that was made possible by a method described in 1898 by Edward L. Thorndike (1874–1949; Figure 7.16).

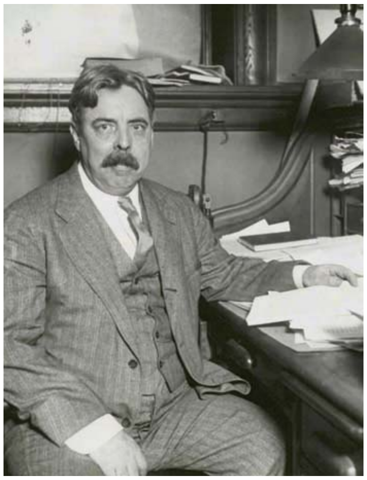

7.16 Edward L. Thorndike (1874–1949) An important early behaviorist, Thorndike was the first to formulate the law of effect.

CATS IN A PUZZLE BOX

Thorndike's method was to set up a problem for an animal to solve. In his classic experiments, he placed a hungry cat inside a box with a latched door. The cat could open the door—and escape from the box—only by performing some simple action such as pulling a loop of wire or pressing a lever (Figure 7.17); and once outside the box, the cat was rewarded with a small portion of food. Then the cat was placed back into the box for another trial so that the procedure could be repeated over and over until the task of escaping the box was mastered.

On the first trial, the cats had no notion of how to escape—and so they meowed loudly and clawed and bit at their surroundings. This continued for several minutes until finally, purely by accident, the animal hit upon the correct response. Subsequent trials brought gradual improvement, and the animal took less and less time to produce the response that unlocked the door. By the time the training sessions were completed, the cats' behavior was almost unrecognizable from what it had been at the start. When placed in the box, they immediately approached the wire loop or the lever, yanked it or pressed it with businesslike dispatch, and hurried through the open door to enjoy the well-deserved reward.

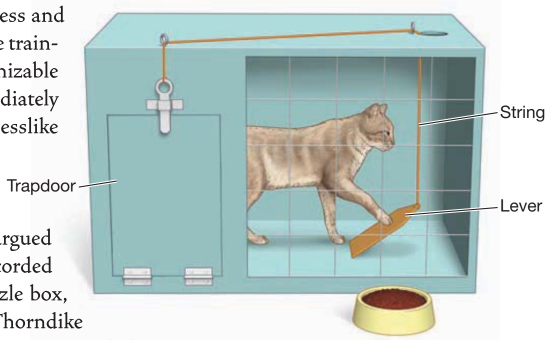

7.17 Puzzle box This box is much like those used by Edward Thorndike. By stepping on a lever attached to a rope, the animal releases the latch and so unlocks the door.

If you observed only the final performance of these cats, you might well credit the animals with reason or understanding. But Thorndike argued that the cats solved the problem in a very different way. As proof, he recorded how much time the cats required on each trial to escape from the puzzle box, and he charted how these times changed over the course of learning. Thorndike found that the resulting curves declined quite gradually as the learning proceeded (Figure 7.18). This isn't the pattern we would expect if the cats had achieved some understanding of how to solve the problem. If they had, their curves would show a sudden drop at some point in the training, when they finally got the point. ("Aha!" muttered the insightful cat, "it's the lever that lets me out," and henceforth howled and bit

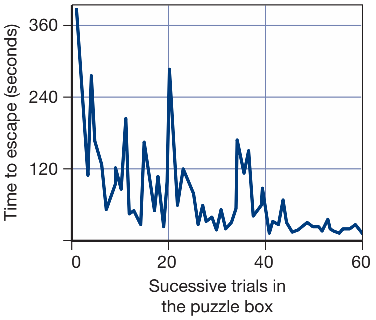

7.18 Learning curve of one of Thorndike's cats To get out of the box, the cat had to move a wooden handle from a vertical to a horizontal position. The figure shows the gradual decline in the animal's response latency (the time it takes to get out of the box). Notice that the learning curve is not very smooth; this feature is common in the learning curves of individual subjects. Smooth learning curves are generally produced by averaging the results of many individual subjects.

no more.) Instead, these learning curves suggest that the cats learned to escape in small increments; they displayed no evidence at all of understanding and certainly no evidence of any sudden insight into the problem's solution.

THE LAW OF EFFECT

In Thorndike's procedure, the cats' initial responses in the puzzle box—biting at the latch, clawing at the walls—all led to failure. As the trials proceeded, though, the cats' tendency to produce these responses gradually weakened. At the same time, the animals' tendency to produce the correct response was weak at first; but, over the trials, this response gradually grew stronger. In Thorndike's terms, the correct response was gradually "stamped in," while futile ones were "stamped out."

law of effect Thorndike's theory that a response followed by a reward will be strengthened, whereas a response followed by no reward (or by punishment) will be weakened.

But what causes this stamping in or stamping out? Thorndike's answer was the **law of effect.** Its key proposition is that if a response is followed by a reward, that response will be strengthened. If a response is followed by no reward (or, worse yet, by punishment), it will be weakened. In general, the strength of a response is adjusted according to the response's consequences (Figure 7.19). In this view, we do not need to suppose that the cat's performance required any sophisticated intellectual processes. We likewise do not need to assume that the animal noticed a connection between its acts and the consequences of those acts. All we need to assert is that, if the animal made a response and a reward followed soon after, that response was more likely to be performed later.

Notice that Thorndike's proposal suggests a clear parallel between how an organism learns during its lifetime and how species evolve, thanks to the forces of natural selection. In both cases, variations that "work"—behaviors that lead to successful outcomes, or individuals with successful adaptations—are kept on. In both cases, variations that are less successful are weakened or dropped. And, crucially, in both cases

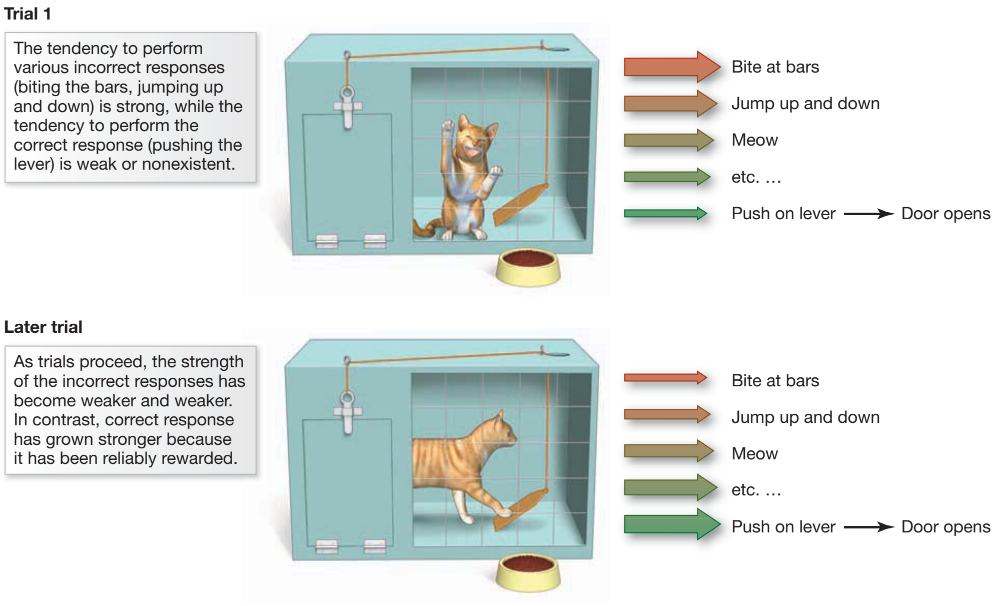

The tendency to perform various incorrect responses (biting the bars, jumping up and down) is strong, while the tendency to perform the correct response (pushing the lever) is weak or nonexistent.

As trials proceed, the strength of the incorrect responses has become weaker and weaker. In contrast, correct response has grown stronger because it has been reliably rewarded.

7.19 The law of effect A schematic presentation of Thorndike's theory of instrumental learning.

the selection involves no guide or supervisor to steer the process forward. Instead, selection depends only on the consequences of actions or adaptations and on whether these serve the organism's biological needs or not.

Skinner and Operant Behavior

Thorndike initiated the experimental study of instrumental behavior; but, unquestionably, the psychologist who shaped the way most modern learning theorists think about the subject was B. F. Skinner (1904–1990; Figure 7.20). Skinner was one of the first theorists to insist on a sharp distinction between classical and instrumental conditioning. He noted that in classical conditioning, the animal's behavior is elicited by the US. Salivation, for example, is set off by an event outside the organism. But in instrumental conditioning, Skinner argued, the organism is much less at the mercy of external factors. Its reactions are emitted from within, as if they were what we ordinarily call "voluntary." Skinner called these instrumental responses **operants:** They operate on the environment to bring about some change that leads to some consequence. And, in Skinner's view, these consequences are crucial. Like Thorndike, Skinner argued that an operant followed by a positive consequence was more likely to be emitted in the future, while an operant followed by a negative consequence was less likely to be emitted again (Skinner, 1938).

Skinner believed, however, that Thorndike's procedure for studying learning was inefficient. Rather than placing animals in a puzzle box (which required many minutes for each learning trial), Skinner sought a procedure in which the instrumental response could be performed repeatedly and rapidly, so that data could be gathered more easily. Many of his studies therefore employed an experimental chamber (popularly called the Skinner box) in which a rat presses a lever or a pigeon pecks at a lighted key in order to gain a reward (Figure 7.21). In these situations, the animal stays in the chamber for a set interval—perhaps an hour at a time—and during that interval, we track the animal's behavior by recording its *response rate*—the number of lever presses or key pecks per unit of time.

operant In Skinner's system, an instrumental response that is defined by its effect (the way it operates) on the environment.

reinforcer A stimulus delivered after a response that makes the response more likely in the future.

7.20 B. F. Skinner (1904–1990) Unmistakably the most influential of the learning theorists, Skinner made a sharp distinction between classical and operant conditioning.

The Major Phenomena of Instrumental Conditioning

As Skinner noted, classical and instrumental conditioning are different in important ways: Classical conditioning builds on a response (UR) that's automatically triggered by a stimulus (US); instrumental conditioning involves behaviors that appear to be voluntary. Classical conditioning involves learning about the relation between two stimuli (US and CS); instrumental conditioning involves learning about the relation between a response and a stimulus (the operant and a reward). Even with these differences, modern theorists have argued that the two forms of conditioning have a lot in common. This makes sense because both involve learning about *relationships* among simple events (stimuli or responses).

It's perhaps inevitable, then, that many of the central phenomena of instrumental learning parallel those of classical conditioning. For example, in classical conditioning, learning trials typically involve the presentation of a CS followed by a US. In instrumental conditioning, learning trials typically involve a response by the organism followed by a reward or **reinforcer.** The reinforcement often involves the presentation of something good, such as grain to a hungry pigeon. Alternatively, reinforcement may involve the termination or prevention of something bad, such as the cessation of a loud noise.

7.21 Animals in operant chambers This rat is trained to press a lever for food reinforcement. Reinforcement consists of a few seconds' access to a feeder located near the response lever.

In both forms of conditioning, the more such pairings there are, the stronger the learning. And if we discontinue these pairings so that the CS is no longer followed by the US or the response by a reinforcer, the result is extinction.

GENERALIZATION AND DISCRIMINATION

An instrumental response is not directly triggered by an external stimulus, the way a CR or UR is. But that doesn't mean external stimuli have no role here. In instrumental conditioning, external events serve as *discriminative stimuli*, signaling for an animal what sorts of behaviors will be rewarded in a given situation. For example, suppose a pigeon is trained to hop onto a platform to get some grain. When a green light is on, hopping on the platform pays off. But when a red light is on, hopping gains no reward. Under these circumstances, the green light becomes a positive discriminative stimulus and the red light a negative one (usually labeled S$^+$ and S$^-$, respectively). The pigeon swiftly learns this pattern and so will hop in the presence of the first and not in the presence of the second.

Other examples are easy to find. A child learns that pinching her sister leads to punishment when her parents are on the scene but may have no consequences otherwise. In this situation, the child may learn to behave well in the presence of the S$^+$ (i.e., when her parents are there) but not in other circumstances. A hypochondriac may learn that loud groans will garner sympathy and support from others but may bring no benefits when others are not around. As a result, he may learn to groan in social settings but not when alone.

Let's be clear, though, about the comparison between these stimuli and the stimuli central to classical conditioning. A CS$^+$ tells the animal about events in the world: "No matter what you do, the US is coming." The S$^+$, on the other hand, tells the animal about the impact of its own behavior: "If you respond now, you'll get rewarded." The CS$^-$ indicates that no matter what the animal does, no US is coming. The S$^-$, in contrast, tells the animal something about its behavior—namely, that there's no point in responding right now.

Despite these differences, generalization in instrumental conditioning functions much the way it does in classical conditioning, and likewise for discrimination. One illustration of these parallels lies in the generalization gradient. We saw earlier that if an organism is trained with one CS (perhaps a high tone) but then tested with a different one (a low tone), the CR will be diminished. The greater the change in the CS, the greater the drop in the CR's strength. The same pattern emerges in instrumental conditioning. In one experiment, pigeons were trained to peck at a key illuminated with yellow light. Later, they were tested with lights of varying wavelengths, and the results showed an orderly generalization gradient (Figure 7.22). As the test light became less similar to the original S$^+$, the pigeons were less inclined to peck at it (Guttmann & Kalish, 1956).

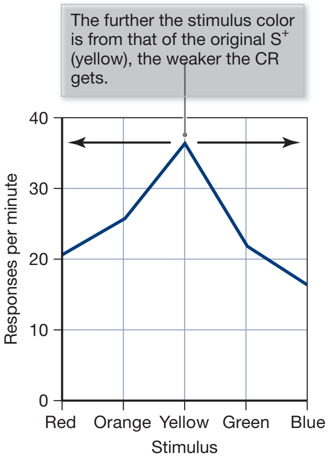

7.22 Stimulus generalization of an instrumental response Pigeons were originally reinforced to peck at a yellow light. When later tested with lights of various colors, the birds showed a standard generalization gradient, pecking more vigorously at colors more similar to yellow (such as green and orange) than at colors farther removed (such as red and blue).

The ability to distinguish an S$^-$ from an S$^+$ obviously allows an organism to tune its behavior to its circumstances. Thus, the dolphins at the aquarium leap out of the water to get a treat when their feeders are around; they don't leap up in the presence of other people. Your pet dog sits and begs when it sees you eating, in hopes that you'll share the snack; but the dog doesn't beg when it sees you drinking. In these and other cases, behaviors are emitted only when the available stimuli indicate that the behavior will now be rewarded.

In fact, animals are quite skilled at making discriminations and can use impressively complex stimuli as a basis for controlling their behavior. In one study, pigeons were trained to peck a key whenever a picture of water was in view, but not to peck otherwise.

Some of the water pictures showed flowing streams; some showed calm lakes. Some pictures showed large bodies of water photographed from far away; some showed small puddles photographed close up. Despite these variations, pigeons mastered this discrimination task and were even able to respond appropriately to new pictures that hadn't been included in the training trials (Herrnstein, Loveland, & Cable, 1976). Apparently, pigeons are capable of discriminating relatively abstract categories—categories not defined in terms of a few simple perceptual features. Similar procedures have shown that pigeons can discriminate between pictures showing trees and pictures not showing trees; thus, for example, they'll learn to peck in response to a picture of a leaf-covered tree or a tree bare of leaves, but not to peck in response to a picture of a telephone pole or a picture of a celery stalk. Likewise, pigeons can learn to peck whenever they're shown a picture of a particular human—whether she's photographed from one angle and close up or from a very different angle, far away, and wearing different clothes (Herrnstein, 1979; Lea & Ryan, 1990; for other examples of complex discriminations, see Cook, K. Cavoto, & B. Cavoto, 1995; Giurfa, Zhang, Jenett, Menzel, & Srinivasan, 2001; Lazareva, Freiburger, & Wasserman, 2004; D. Premack, 1976, 1978; A. Premack & D. Premack, 1983; Reiss & Marino, 2001; Wasserman, Hugart, & Kirkpatrick-Steger, 1995; Zentall, 2000).

SHAPING

Once a response has been made, reinforcement will strengthen it. Once the dolphin has leapt out of the water, the trainers can reward it, encouraging further leaps. Once the pigeon pecks the key, food can be delivered, making the next peck more likely. But what causes the animal to perform the desired response in the first place? What leads to the first leap or the first peck? This is no problem for many responses. Dolphins occasionally leap with no encouragement from a trainer, and pecking is something pigeons do all the time. If the trainer is patient, therefore, an opportunity to reinforce (and thus encourage) these responses will eventually arrive.

But what about less obvious responses? For example, rats quite commonly manipulate objects in their environment, and so they're likely to press on a lever if we put one within reach. But what if we place a lever so high that the rat has to stretch up on its hind legs to reach it? Now the rat might never press the lever on its own. Still, it can learn this response if its behavior is suitably shaped. This **shaping** is accomplished by a little "coaching," using the method of *successive approximations*.

shaping The process of eliciting a desired response by rewarding behaviors that are increasingly similar to that response.

How could we train a rat to press the elevated lever? At first, we reinforce the animal merely for walking into the general area where the lever is located. As soon as the rat is there, we deliver food. After a few such trials, the rat will have learned to remain in this vicinity most of the time, so we can now increase our demand. When the rat is in this neighborhood, sometimes it's facing one way, sometimes another; but from this point on, we reinforce the rat only if it's in the area and facing the lever. The rat soon masters this response too; now it's facing in the right direction most of the time. Again, therefore, we increase our demand: Sometimes the rat is facing the lever with its nose to the ground; sometimes it's facing the lever with its head elevated. We now reinforce the animal only when its head is elevated—and soon this, too, is a well-established response. We continue in this way, counting on the fact that at each step, the rat naturally varies its behavior somewhat, allowing us to reinforce just those variations we prefer. Thus, we can gradually move toward reinforcing the rat only when it stretches up to the lever, then when it actually touches the lever, and so on. Step by step, we guide the rat toward the desired response.

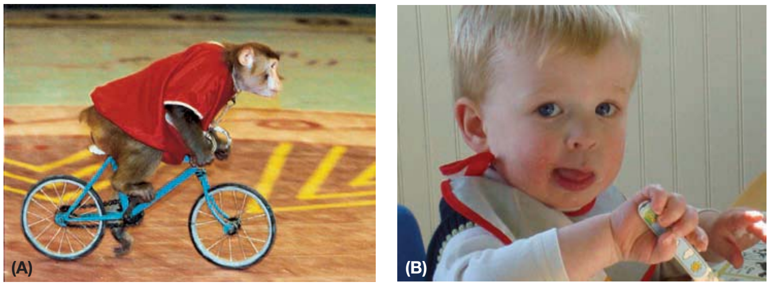

7.23 Shaping (A) Nhat, a four-year-old monkey with the Hanoi Circus, was trained using techniques essentially identical to those described in the text—a process of gradual shaping that leads to the desired response. (B) Parents use shaping to teach their children to eat with utensils. As the parents' demands gradually increase, the child's behavior comes closer and closer to the "standard" pattern of using the spoon and fork.

Using this technique, people have trained animals to perform all kinds of complex behavior; many human behaviors probably come about in the same way (Figure 7.23). For example, how do parents in Western countries teach their children to eat with a spoon and fork? At first, they reward the child (probably with smiles and praise) just for holding the spoon. This step soon establishes a grasp-the-spoon operant; and, at that point, the parents can require a bit more. Now they praise the child just when she touches food with the spoon—and, thanks to this reinforcement, the new operant is quickly established. If the parents continue in this way, gradually increasing their expectations, the child will soon be eating in the "proper" way.

Similar techniques are used in therapeutic settings to shape the behavior of the hospitalized mentally ill. Initially, the hospitalized patients might be rewarded just for getting out of bed. Once that behavior is established, the requirement is increased so that, perhaps, the patients have to move around a bit in their room. Then, later, the patients are rewarded for leaving the room and going to breakfast or getting their medicine. In this way, the behavior therapist can gradually lead the patients into a more acceptable level of functioning (see Chapter 17).

WHAT IS A REINFORCER?

We've now said a great deal about what reinforcement does; it encourages some responses, discourages others, and even—through the process of shaping—creates entirely new responses. But what is it that makes a stimulus serve as a reinforcer?

Some stimuli serve as reinforcers because of their biological significance. These *primary reinforcers* include food, water, escape from the scent of a predator, and so on—all stimuli with obvious importance for survival. Other reinforcers are social—think of the smiles and praise from parents that we mentioned in our example of teaching a child to use a spoon and fork.

Other stimuli are initially neutral in their value but come to act as reinforcers because, in the animal's experience, they've been repeatedly paired with some other, already established reinforcer. This kind of stimulus is called a *conditioned reinforcer*, and it works just like any other reinforcer. A plausible example is *money*—a reward that takes its value from its association with other more basic reinforcers.

Other reinforcers, however, fall into none of these categories, so we have to broaden our notion of what a reinforcer is. Pigeons, for example, will peck in order to gain *information* about the availability of food (e.g., G. Bower, McLean, & Meachem, 1966; Hendry, 1969). Monkeys will work merely to open a small window through which they can see a moving toy train (Butler, 1954). Rats will press a lever to gain access to an exercise wheel (D. Premack, 1965; but also Timberlake & Allison, 1974; Timberlake, 1995). And these are just a few examples of reinforcers.

But examples like these make it difficult to say just what a "reinforcement" is, and in practice the stimuli we call *reinforcements* are generally identified only after the fact. Is a glimpse of a toy train reinforcing? We can find out only by seeing whether an animal will work to obtain this glimpse. Remarkably, no other, more informative definition of a reinforcer is available.

BEHAVIORAL CONTRAST AND INTRINSIC MOTIVATION

Once we've identified a stimulus as a reinforcer, what determines how effective the reinforcer will be? We know that some reinforcers are more powerful than others—and so an animal will respond more strongly for a large reward than for a small one. However, what counts as large or small depends on the context. If a rat is used to getting 60 food pellets for a response, then 16 pellets will seem measly and the animal will respond only weakly for this puny reward. But if a rat is used to getting only 4 pellets for a response, then 16 pellets will seem like a feast and the rat's response will be fast and strong (for the classic demonstration of this point, see Crespi, 1942). Thus, the effectiveness of a reinforcer depends largely on what other rewards are available (or have recently been available); this effect is known as **behavioral contrast.**

behavioral contrast A response pattern in which an organism evaluates a reward relative to other available rewards or those that have been available recently.

Contrast effects are important for their own sake, but they may also help explain another (somewhat controversial) group of findings. In one study, for example, nursery-school children were given an opportunity to draw pictures. The children seemed to enjoy this activity and produced a steady stream of drawings. The experimenters then changed the situation: They introduced an additional reward so that the children now earned an attractive "Good Player" certificate for producing their pictures. Then, later on, the children were again given the opportunity to draw pictures—but this time with no provision for "Good Player" rewards. Remarkably, these children showed considerably less interest in drawing than they had at the start and chose instead to spend their time on other activities (see, for example, Lepper, Greene, & Nisbett, 1973; also Kohn, 1993).

Some theorists say these data illustrate the power of behavioral contrast. At the start of the study, the activity of drawing was presumably maintained by certain reinforcements in the situation—perhaps encouragement from the teachers or comments by other children. Whatever the reinforcements were, they were strong enough to maintain the behavior; we know this because the children were producing drawings at a steady pace. Later on, though, an additional reinforcement (the "Good Player" certificate) was added and then removed. At that point the children were back to the same rewards they'd been getting at the start, but now these rewards seemed puny in comparison to the greater prize they'd been earning during the time when the "Good Player" award was available. As a consequence, the initial set of rewards was no longer enough to motivate continued drawing.

Other theorists interpret these findings differently. In their view, results like this one suggest that there are actually two different types of reward. One type is merely tacked onto a behavior and is under the experimenter's control; it's the sort of reward that's in play when we give a pigeon a bit of food for pecking a key, or hand a factory worker a paycheck for completing a day's work. The other type of reward is intrinsic to the behavior and independent of the experimenter's intentions; these rewards are in play when someone is engaging in an activity just for the pleasure of the activity itself.

In addition, these two forms of reward can interfere with each other. Thus, in the study with the "Good Player" certificates, the children were initially drawing pictures for an intrinsic reward. Drawing, in other words, was a form of *play* engaged in for its own sake. However, once the external rewards (the certificates) entered the situation, the same activity became a form of *work*—something you do for a payoff. And once the activity was redefined in this way, then the absence of a payoff meant there was no longer any point in drawing. (For more on this topic, see Chapter 12.)

Debate continues about which of these interpretations is preferable—the one based on behavioral contrast or the one based on intrinsic motivation. (It also seems plausible that *both* interpretations may capture aspects of what's going on here.) Clearly, there's more to be learned about reinforcement and the nature of motivation. (For further exploration, see Bowles, 2008; Deci, Koestner, & Ryan, 1999a, 1999b; Eisenberger, Pierce, & Cameron, 1999; Henderlong & Lepper, 2002.)

SCHEDULES OF REINFORCEMENT

Let's now return to the issue of how extrinsic reinforcements work, since—by anyone's account—these reinforcements play a huge role in governing human (and other species') behavior. We do, after all, work for money, buy lottery tickets in hopes of winning, and act in a way that we believe will bring us praise. But notice that in all of these examples, the reinforcement comes only occasionally: We aren't paid after every task we do at work; we almost never win the lottery; and we don't always get the praise we seek. Yet we show a surprising resistance to extinction of those behaviors. About some things, we have learned that if you don't succeed, it pays to try again—an important strategy for achieving much of what we earn in life.

This pattern, in which we're reinforced for only some of our behaviors, is known as **partial reinforcement.** In the laboratory, partial reinforcement can be provided according to different **schedules of reinforcement**—rules about how often and under what conditions a response will be reinforced. Some behaviors are reinforced via a **ratio schedule,** in which you're rewarded for producing a certain number of responses (Figure 7.24). The ratio can be "fixed" or "variable." In a "fixed-ratio 2" (FR 2) schedule, for example, two responses are required for each reinforcement; for an "FR 5" schedule, five responses are required. In a variable-ratio schedule, the number of responses required changes from trial to trial. Thus, in a "VR 10" schedule, 10 responses are required *on average* to get a reward—so it might be that the first 5 responses are enough to earn one reward, but 15 more are needed to earn the next.

Other behaviors are reinforced on an **interval schedule,** in which you're rewarded for producing a response after a certain amount of time has passed (Figure 7.25). Thus, on an "FI 3-minute" schedule, responses made during the 3-minute interval aren't reinforced; but the first response after the 3 minutes have passed will earn a reward. Interval schedules can also be variable: For a "VI 8-minute" schedule, reinforcement is available on average after 8 minutes; but the exact interval required varies from trial to trial.

partial reinforcement A learning condition in which only some of the organism's responses are reinforced.

schedule of reinforcement The rules about how often and under what conditions a response will be reinforced.

ratio schedule A pattern of delivering reinforcements only after a certain number of responses.

interval schedule A pattern of delivering reinforcements only after a certain amount of time has passed.

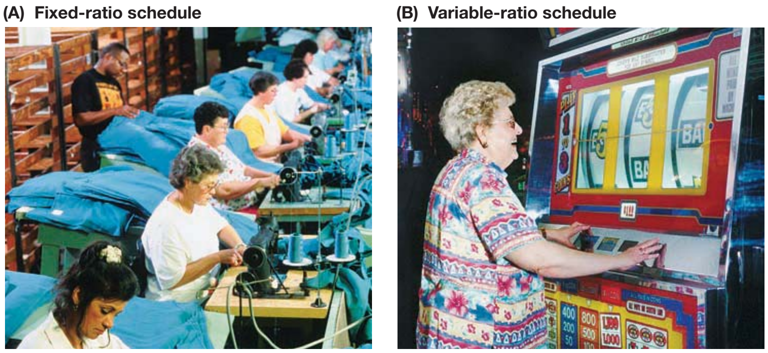
(A) Fixed-ratio schedule (B) Variable-ratio schedule

7.24 Ratio reinforcement (A) Workers in a garment factory are usually paid a certain amount for each item of clothing completed, so they're rewarded on a fixed-ratio schedule of reinforcement. (B) By law, slot machines pay out on a certain percentage of tries. But these machines pay out at random, so that people feeding coins to a machine are rewarded on a variable-ratio reinforcement schedule.

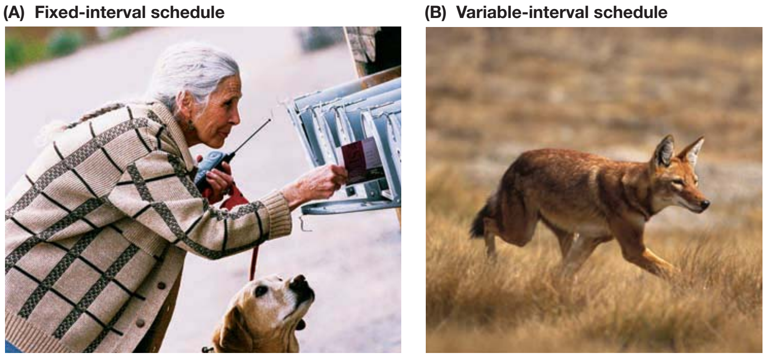
(A) Fixed-interval schedule **(B) Variable-interval schedule**

7.25 Interval reinforcement (A) The behavior of checking the mailbox isn't rewarded until the mail has actually been delivered. Then, the first check of the mailbox after this delivery will be rewarded (by the person actually getting some mail!). This is a fixed-interval schedule of reinforcement. (B) When a wolf prowls through a meadow, the rodents it seeks will all retreat into hiding; so, if the wolf returns soon, he'll find no prey. Eventually, though, the rodents will come back, and then another hunting trip by the wolf will pay off—he'll find his dinner. Therefore, a return visit by the wolf will be rewarded only after some time has passed—and hence this is an interval schedule. But the rodents may sometimes return a little sooner, and sometimes a little later—and so this is a variable interval schedule.

Changing Behaviors or Acquiring Knowledge?

We've almost finished our discussion of instrumental conditioning, except for one crucial question: What is it exactly that animals learn in an instrumental conditioning procedure? The law of effect implies that the learning is best understood as a change in behavior, in which responses are either being strengthened or weakened by the mechanical effects of reinforcement. From the earliest days of learning theory, however, there was an alternative view of conditioning—one asserting that behavior change isn't the key; what matters instead is the acquisition of new knowledge.

One of the most prominent proponents of this alternative view was Edward C. Tolman (1886–1959; Figure 7.26), and many forms of evidence support his position. For example, consider cases of **latent learning**—learning that takes place without any corresponding change in behavior. In one experiment, rats were allowed to explore a maze, without any reward, for 10 days. During these days, there was no detectable change in the rats' behavior; and so, if we define learning in terms of behavior change, there was no learning. But in truth the rats *were* learning—and in particular, they were gaining knowledge about how to navigate the maze's corridors. This became obvious on the 11th day, when food was placed in the maze's goal box for the first time. The rats learned to run to this goal box, virtually without error, almost immediately. The knowledge they had acquired earlier now took on motivational significance, so the animals swiftly displayed what they had learned (Tolman & Honzik, 1930; also H. Gleitman, 1963; Tolman, 1948).

latent learning Learning that occurs without a corresponding change in behavior.

In this case, the knowledge the rats had gained can be understood as a *mental map* of the maze—an internal representation of spatial layout that indicates what is where and what leads to what. Other evidence suggests that many species rely on such maps—to guide their foraging for food, their navigation to places of safety, and their choice of a path to the watering hole. These maps can be relatively complex and are typically quite accurate (Gallistel, 1994; J. Gould, 1990).


7.26 Edward C. Tolman (1886–1959) Tolman was an early advocate for the idea that learning involves a change in knowledge rather than a change in overt behavior.

CONTINGENCY IN INSTRUMENTAL CONDITIONING

To understand latent learning or cognitive maps, we need to emphasize what an organism *knows* more than what an organism *does*. We also need to consider an organism's cognition for another reason: Recall that, in our discussion of classical conditioning, we saw that learning doesn't depend only on the CS being paired with the US; instead, the CS needs to *predict* the US, telling the animal when the US is more likely and when it's less likely. Similarly, instrumental conditioning doesn't depend only on responses being paired with rewards. Instead, the response needs to predict the reward, so that (for example) the probability of getting a pellet after a lever press has to be greater than the probability of getting it without the press.

What matters for instrumental conditioning, therefore, is not merely the fact that a reward arrives after the response is made. Instead, what matters is the *relationship* between responding and getting the reward, and this relationship actually gives the animal some control over the reward: By choosing when (or whether) to respond, the animal itself can determine when the reward is delivered. And it turns out that this control is important, because animals can tell when they're in control and when they're not—and they clearly prefer being in control.

One line of evidence comes from a study in which infants were placed in cribs that had colorful mobiles hanging above them. Whenever the infants moved their heads, they closed a switch in their pillows; this activated the overhead mobile, which spun merrily for a second or so. The infants soon learned to shake their heads about, making their mobiles turn. They evidently enjoyed this, smiling and cooing at their mobiles, clearly delighted to see the mobiles move.

A second group of infants was exposed to a similar situation, but with one important difference: Their mobile turned just as often as the mobile for the first group; but it was moved for them, not by them. This difference turned out to be crucial. After a few days, these infants no longer smiled and cooed at the mobile, nor did they seem particularly interested when it turned. This suggests that what the first group of infants liked about the mobile was not that it moved, but that they made it move. Even a 2-month-old infant wants to be the master of his own fate (J. S. Watson, 1967; Figure 7.27).

This study with infants illustrates the joy of mastery. Another series of studies demonstrates the despair of no mastery at all. These studies focus on **learned helplessness**—an acquired sense that one has lost control over one's environment, with the sad consequence that one gives up trying (Seligman, 1975).

learned helplessness A condition of passivity apparently created by exposure to inescapable aversive events. This condition inhibits or prevents learning in later situations in which escape or avoidance is possible.

The classic experiment on learned helplessness used two groups of dogs, A and B, which received strong electric shocks while strapped in a hammock. The dogs in group A were able to exert some control over their situation: They could turn the shock off whenever it began simply by pushing a panel that was placed close to their noses. The dogs in group B had no such power. For them, the shocks were inescapable. But the number and duration of the shocks were the same as for the first group. This was guaranteed by the fact that, for each dog in group A, there was a corresponding animal in group B whose fate was yoked to that of the first dog. Whenever the group A dog was shocked, so was the group B dog. Whenever the group A dog turned off the shock, the shock was turned off for the group B dog. Thus, both groups experienced exactly the

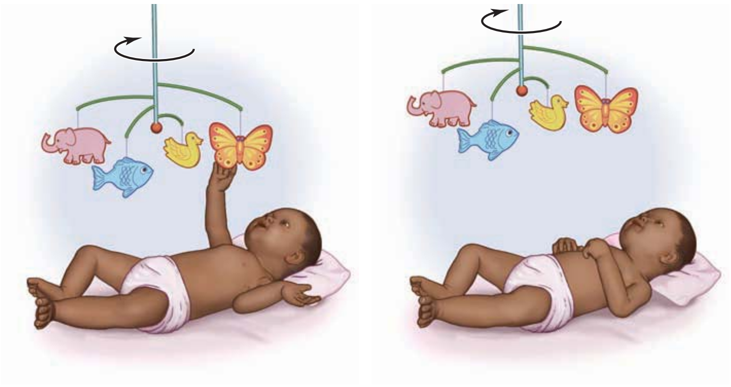

7.27 **Response control** Infants who could make a mobile move smiled and cooed at it; those who had no control over its motion quickly lost interest in the mobile.

same level of physical suffering; the only difference was what the animals were able to do about it. The dogs in group A had some control; those in group B could only endure.

What did the group B dogs learn in this situation? To find out, both groups of dogs were next presented with a task in which they had to learn to jump from one compartment to another to avoid a shock. The dogs in group A learned easily. During the first few trials, they ran about frantically when the shock began but eventually scrambled over the hurdle into the other compartment, where there was no shock. Based on this experience, they soon learned to leap over the hurdle the moment the shock began, easily escaping the aversive experience. Then, with just a few more trials, these dogs learned something even better: They jumped over the hurdle as soon as they heard the tone signaling that shock was about to begin; as a result, they avoided the shock entirely.

Things were different for the dogs in group B, those that had previously experienced the inescapable shock. Initially, these dogs responded to the electric shock just like the group A dogs did—running about, whimpering, and so on. But they soon became much more passive. They lay down, whined, and simply took whatever shocks were delivered. They neither avoided nor escaped; they just gave up. In the earlier phase of the experiment, they really had been objectively helpless; there truly was nothing they could do. In the shuttle box, however, their helplessness was only subjective because now they did have a way to escape the shocks. But they never discovered it, because they had learned to be helpless (Seligman & Maier, 1967).

Martin Seligman, one of the discoverers of the learned helplessness effect, asserts that depression in humans can develop in a similar way. Like the dog that has learned to be helpless, Seligman argues, the depressed patient has come to believe that nothing she does will improve her circumstances. And Seligman maintains that, like the dog, the depressed patient has reached this morbid state by experiencing a situation in which she really was helpless. While the dog received inescapable shocks in its hammock, the patient found herself powerless in the face of bereavement, some career failure, or serious illness (Seligman, Klein, & Miller, 1976). In both cases, the outcome is the same—a belief that there's no contingency between acts and outcomes, and so there's no point in trying. (For more on this theory of depression, see Chapter 16.)*

OBSERVATIONAL LEARNING

We've now considered three types of learning—habituation, classical conditioning, and instrumental conditioning, each of which can be observed in many different species. Indeed, across this chapter we've mixed together examples of how these learning principles apply to humans and how they apply to a range of other organisms. In contrast, a fourth type of learning was for many years regarded as exclusively human; but more recent evidence indicates that this learning, too, can be observed in many

---

*Before moving on, we should mention that studies of learned helplessness in dogs—like many studies mentioned in this chapter—raise ethical questions. Is it ethically acceptable to deliver electric shock to animals, knowing that the shock produces pain and may produce a depression-like state? These are difficult questions, but the link between these studies and human depression may well help us to understand depression and may lead to more effective forms of treatment for, or even prevention of, this illness. These considerations persuade many investigators that the animal work, as troubling as it is, is necessary to further the goal of diminishing human suffering. In many countries, an Animal Care and Use Committee at each research center weighs these issues before approving any animal experiment, ensuring that every experiment is scrutinized and evaluated with close attention to the protection of animal welfare.

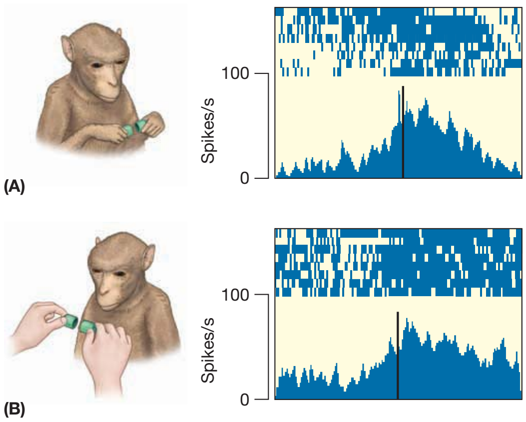

7.28 Mirror neurons Panel A shows the respones of a neuron in a monkey's motor cortex when the animal breaks a peanut. Panel B shows the remarkably similar pattern of activity when the monkey watches someone else open a peanut.

observational learning The process of watching how others behave and learning from their example.

vicarious conditioning A form of learning in which the learner acquires a conditioned response merely by observing another participant being conditioned.

mirror neurons Neurons that fire whenever an animal performs an action, such as stretching out its arm or reaching toward a target, and also whenever the animal watches another performing the same action.

species. At issue here is **observational learning**—a process through which we watch how others behave and learn from their example.

Other species are certainly capable of observational learning (e.g., Bugnyar & Kotrschal, 2002; White & Galef, 1998). For example, monkeys are capable of **vicarious conditioning**—acquiring a fear response, triggered by a specific stimulus, merely by watching another monkey show fear in response to that stimulus (Mineka & Ben Hamida, 1998; Öhman & Mineka, 2001). This learning is probably supported by a specific type of neuron found in the brains of many species: These **mirror neurons** are located in the frontal lobe, near the motor cortex, and fire whenever an animal performs an action such as stretching out its arm or reaching toward a target (Figure 7.28). Remarkably, the same neurons fire when the animal observes someone else performing the same action (Rizzolatti & Craighero, 2004); and several theorists have suggested that these neurons play an essential role both in understanding others' behavior and in imitating that behavior.

Observational learning can also be documented in non-primate species. In one experiment, pigeons were allowed to observe other pigeons getting rewarded either for pecking at a disk or for stepping on a lever. When the "watchers" were then placed in the same experimental chamber, they tended to make the same response—pecking or stepping—they had observed earlier (Zentall, Sutton, & Sherburne, 1996).

Observational learning—and imitation in particular—also plays a central role for humans, even when we're very young (e.g., Bandura, 1977, 1986). Indeed, infants less than a month old imitate the facial expressions of people in their environment (e.g., Meltzoff & Moore, 1977); not much later, they start imitating a range of other behaviors. Sometimes human imitation takes the form of outright mimicry—duplicating, as best we can, the exact behaviors we have observed in others. At other times, the imitation is more sophisticated: We observe others and draw general conclusions about what sorts of behaviors are permissible in that situation.

We can see both forms of imitation at work in a classic study by Bandura (1969, 1977; Figure 7.29). In that study, children watched while an adult punched and kicked a large, inflated doll. Later on, the children were allowed to play in the room with various toys; and those who had observed the adults' aggression tended to mimic it—kicking and punching the doll exactly as the adult had. But these children also showed a broader form of imitation: They became aggressive toward other toys, as if they'd figured out by observation that aggression was permitted and maybe even appropriate in this environment.

As a further illustration of observational learning, consider the impact of media violence—seen on TV or in video games—on children. Research has made it clear that media violence—whether it's watching cartoon characters punch each other or guiding videogame characters as they rip each other apart—does encourage violence in child viewers (e.g., C. Anderson & Bushman, 2002; Bushman & Anderson, 2009; Carnagey & Anderson, 2005; Feshbach & Tangney, 2008), Indeed, the evidence is compelling enough that six major professional societies (including the American Psychological Association and the American Medical Association) issued a joint statement noting that studies "point overwhelmingly to a causal connection between media violence and aggressive behavior in some children" (Joint Statement, 2000, p. 1). This effect probably involves several different mechanisms; still, it's a compelling example of just how important observational learning can be—and it reminds us that this learning can have both good effects (the acquisition of new skills) and bad.

7.29 SCIENTIFIC METHOD: Can children learn to be aggressive through observational learning?

Method

1. Two groups of children were each shown a different short film of an adult playing with a large inflatable “Bobo” doll.

2. In one video the adult played quietly with the doll. In the other the adult hit and kicked the doll.

Results

When children were allowed to play with the doll later, those who had seen the aggressive play were much more likely to hit and kick the doll, as they had seen the adult do.

CONCLUSION: Observational learning can have a powerful effect on aggression.

SOURCE STUDIES: Bandura, 1969, 1977

VARIETIES OF LEARNING

Overall, the attempt to find general principles of learning—principles that apply to virtually all species—has been rather successful. This is why we can gain insights into human depression by studying helplessness in dogs, and it’s how we’ve increased our understanding of human drug addiction—thanks to research on classical conditioning in rats. Other examples of learning phenomena shared across species are easy to find.

We need to acknowledge, though, that there are also important differences from one species to the next in how learning proceeds. These differences are often best understood by taking a biological perspective on learning—a perspective that highlights the actual *function* of learning in each species’ natural environment (see, for example, Bolles & Beecher, 1988; Domjan, 2005; Rozin & Schull, 1988).

Biological Influences on Learning: Belongingness

In the early days of learning theory, investigators widely believed that animals (both humans and others) are capable of connecting any CS to any US in classical conditioning, and of associating virtually any response with any reinforcer in instrumental conditioning. A child could be taught that a tone signaled the approach of dinner or that a flashing light or a particular word did. Likewise, a rat could be trained to press a lever to get food, water, or access to a sexually receptive mate.

But a great deal of evidence speaks against this idea; instead, each species seems predisposed to form some associations and not others. The predispositions put *biological constraints* on that species’ learning, governing what the species can learn easily and what it can learn only with difficulty. These associative predispositions are probably hardwired and likely to be a direct product of our evolutionary past (Rozin & Kalat, 1971, 1972; Seligman & Hager, 1972).

TASTE AVERSION LEARNING

taste aversion learning A form of learning in which an organism learns to avoid a taste after just one pairing of that taste with illness.

A central example of the biological constraints on learning comes from studies of **taste aversion learning.** These studies make it clear that, from the organism's viewpoint, some stimuli belong together and some do not (Domjan, 1983, 2005; Garcia & Koelling, 1966).

To understand this phenomenon, we need to begin with the fact that when a wild rat encounters a novel food, it generally takes only a small bite at first. If the rat suffers no ill effects from this first taste, it will return (perhaps a day or two later) for a second helping and will gradually make the food a part of its regular diet. But what if this novel food is harmful, either because of some natural toxin or an exterminator's poison? In that case, the initial taste will make the rat sick; but because it ate only a little of the food, the rat will probably recover. Based on this experience, though, the rat is likely to develop a strong aversion to that particular flavor, so it never returns for a second dose of the poison.

This sort of learning is easily documented in the laboratory. The subjects, usually rats, are presented with a food or drink that has a novel flavor—perhaps water with some vanilla added. After drinking this flavored water, the rats are exposed to X-ray radiation—not enough to injure them, but enough to make them ill. After they recover, the rats show a strong aversion to the taste of vanilla and refuse to drink water flavored in this way (Figure 7.30).

7.30 **Taste aversion** The text describes taste-aversion learning in the lab, but the same pattern is easily observed in other settings. Young adults often develop a taste aversion to particular forms of alcohol. Usually this is the consequence of a single episode of overindulgence: After drinking far too much vodka, for example, the person ends up quite sick, and that experience can lead to a lifelong distaste for vodka.

This learned taste aversion is actually based on classical conditioning. The flavor (here, vanilla) serves as the CS, and the sensation of being sick serves as the US. This is, however, a specialized type of classical conditioning that is distinct from other forms in the sheer speed of learning: One pairing of a taste + illness is all it takes to establish the connection between them. This *one-trial learning* is obviously much faster than the speed of ordinary classical conditioning. What's more, this form of conditioning is distinctive in its timing requirements. In most classical conditioning, the CS must be soon followed by the US; if too much time passes between these two stimuli, the likelihood of conditioning is much reduced (see Figure 7.10). In taste aversion learning, in contrast, conditioning can be observed even if several hours elapse between the CS and the US.

Learned taste aversions are also remarkable for their specificity. In one early study, thirsty rats were allowed to drink sweetened water through a tube. Whenever the rats licked the nozzle of this tube, a bright light flashed and a loud clicking noise sounded. Thus, the sweetness, bright light, and loud noise were always grouped together; if one was presented, all were presented. One group of these rats then received an electric shock to the feet. A second group was exposed to a dose of X-rays strong enough to cause illness.

Notice, then, that we have two different USs—illness for one group and foot shock for the other. Both groups also have received a three-part CS: sweet + bright + noisy. The question is: How will the animals put these pieces together—what will get associated with what?

To find out, the experimenters tested the rats in a new situation. They gave some of the rats sweetened water, unaccompanied by either light or noise. Rats that had received foot shock showed no inclination to avoid this water; apparently, they didn't associate foot shock with the sweet flavor. However, rats that had been made ill with X-rays refused to drink this sweetened water; they associated their illness with the taste (Figure 7.31).

Another group of rats were tested with unflavored water accompanied by the light and sound cues that were present during training. Now the pattern was reversed. Rats that had become ill showed no objection to this water. For them, the objectionable (sweet) taste was absent from this test stimulus, and they didn't associate their illness with the sights and sounds that were present during the test. However, rats that had been shocked earlier refused to drink this water; in their minds, pain was associated with bright lights and loud clicks (Garcia & Koelling, 1966).

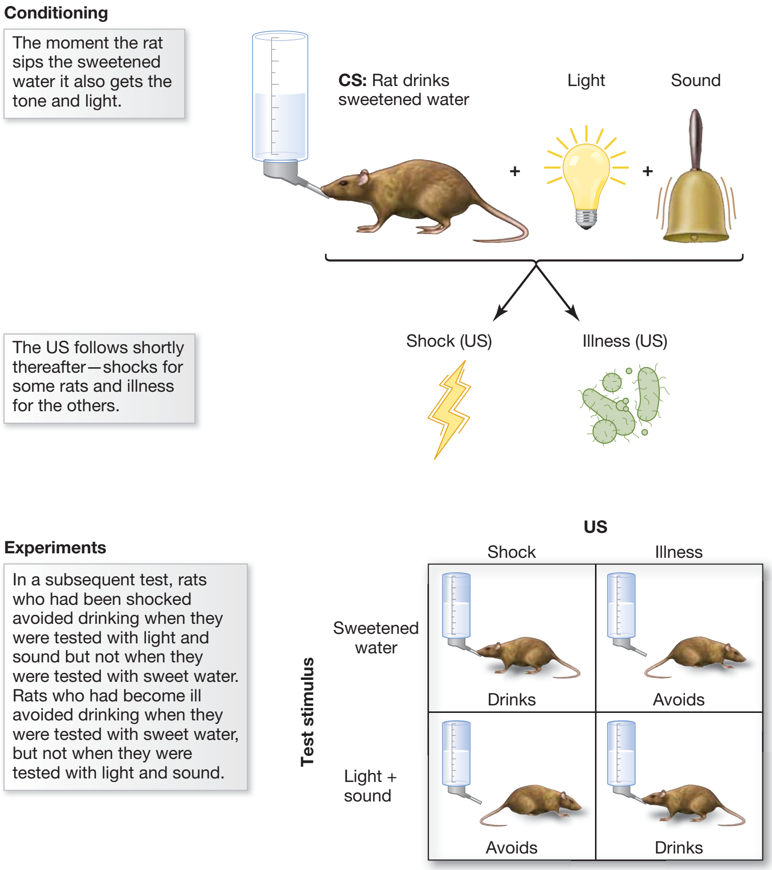

7.31 The specificity of taste aversion Rats that had been shocked associated the shock with the lights and sounds (but not the taste) that had accompanied the painful experience. Rats that had become ill associated the illness with a taste (but not with the lights and sounds).

For the rat, therefore, taste goes with illness, and sights and sounds go with externally induced pain. And for this species, this pattern makes biological sense. Illness in wild rats is likely to have been caused by harmful or tainted food, and rats generally select their food largely on the basis of flavor. So there's survival value in the rats being able to learn quickly about the connection between a particular flavor and illness; this will provide useful information for them as they select their next meal, ensuring that they don't resample the harmful berries or poisoned meat.

Using this logic, one might expect species that choose foods on the basis of other attributes to make different associations. For example, many birds make their food choices from a distance, relying on the food's visual appearance. How will this behavior affect the data? In one study, quail were given blue, sour water to drink and were then given a low dose of poison—enough to make them ill, but not enough to harm them. Some of the birds were later tested with blue, unflavored water; others were tested with water that was sour but colorless. The results showed that the quail had developed a strong aversion to blue water but no aversion to the sour water. They learned which water was safe based on its color rather than its taste (Wilcoxin, Dragoin, & Kral, 1971).

Clearly, what belongs with what depends upon the species. Birds are predisposed to associate illness with visual cues. Rats (and many other mammals) associate illness with taste. In each case, the bias makes the animal more prepared to form certain associations and far less prepared to form others (Seligman, 1970).

prepared learning Learning that occurs without extensive training because of an evolved predisposition to the behavior.

We should also mention that taste aversion learning, as important as it is, is just one example of **prepared learning** (Figure 7.32). We mentioned a different example in the Prologue: Humans in one experiment were shown specific pictures as the CS and received electric shocks as the US. When the pictures showed flowers or mushrooms, learning was relatively slow. When the pictures showed snakes, learning was much quicker. The implication is that humans (and many other primates) are innately prepared to associate the sight of a snake with unpleasant or even painful experiences (Öhman & Mineka, 2003; Öhman & Soares, 1993; also Domjan, Cusato, & Krause, 2004).

These results may help us understand why so many people are afraid of snakes and why strong phobias for snakes are relatively common. Perhaps it's not surprising that many cultures regard snakes as the embodiments of evil. All these facts may simply be the result of prepared learning in our species—our innate tendencies toward making certain associations but not others.

7.32 Prepared learning Humans seem prepared to associate aversive outcomes with the sight of snakes. This may be one reason many find the story of the Garden of Eden so compelling.

BIOLOGICAL CONSTRAINTS ON INSTRUMENTAL CONDITIONING

Prepared learning can also be demonstrated in instrumental conditioning because, from an animal's viewpoint, certain responses belong with some rewards and not others (Shettleworth, 1972). For example, pigeons can easily be taught to peck a lit key to obtain food or water, but it's extremely difficult to train a pigeon to peck in order to escape electric shock (Hineline & Rachlin, 1969). In contrast, pigeons can easily be taught to hop or flap their wings to get away from shock, but it's difficult to train the pigeon to produce these same responses in order to gain food or water.

Once again, this pattern makes good biological sense. The pigeon's normal reaction to danger is to hop away or break into flight, so the pigeon is biologically prepared to associate these responses with aversive stimuli such as electrical shock. Pecking, in contrast, is not part of the pigeon's innate defense pattern, so it's difficult for the pigeon to learn pecking as an escape response (Bolles, 1970). Conversely, since pecking is what pigeons do naturally when they eat, the pigeon is biologically prepared to associate this response with food or drink; it's no wonder, then, that pigeons easily learn to make this association in the psychologist's laboratory.

Different Types of Learning

It seems therefore that we need to "tune" the laws of learning on a case-by-case basis to accommodate the fact that a given species learns some relationships easily, others only with difficulty, and still others not at all. This tuning builds some flexibility into the laws of learning; but it allows us to retain the idea that there *are* general laws, applicable (with the appropriate tuning) to all species and to all situations. Other evidence suggests, though, that we must go further than this, because some types of learning follow their own specialized rules and depend on specialized capacities found in that species and few others. On this basis, we need to do more than adjust the laws of learning. We may also need some entirely new laws—laws that are specific to the species that does the learning and to what's being learned (Gallistel, 1990; Roper, 1983).

7.33 The Clark's nutcracker

As one example, consider the Clark's nutcracker, a bird that makes its home in the American Southwest (Figure 7.33). In the summer, this bird buries thousands of pine nuts in various hiding places over an area of several square miles. Then, throughout the winter and early spring, the nutcracker flies back again and again to dig up its thousands of caches. The bird doesn't mark its cache sites in any special way. Instead, it relies on memory to find its stash—a remarkable feat that few of us could duplicate.

The Clark's nutcracker has various anatomical features that support its food-hoarding activities—for example, there's a special pouch under its tongue that it fills with pine nuts when flying to find a hiding place. The bird's extraordinary ability to learn a huge number of geographical locations, and then to remember these locations for the next few months, is probably a similar evolutionary adaptation. Like the tongue pouch, this learning ability is a specialty of this species: Related birds like jays and crows don't store food in this way; and, when tested, they have a correspondingly poorer spatial memory (D. Olson, 1991; Shettleworth, 1983, 1984, 1990).

Many phenomena of animal learning—in birds, fish, and mammals—reveal similar specializations. In each case, the organism has some extraordinary ability not shared even by closely related species. In each case, the ability has obvious survival value and seems quite narrow. The Clark's nutcracker, for example, has no special skill in remembering pictures or shapes; instead, its remarkable memory comes into play only in the appropriate setting—when hiding and then relocating buried pine nuts. Similarly, many birds show remarkable talent in learning the particular songs used by their species. This skill, however, can be used for no other purpose: A zebra finch easily masters the notes of the zebra finch's song but is utterly inept at learning any other (nonmusical) sequence of similar length and complexity. Truly, then, these are specialized learning abilities—only one or a few species have them, and they apply only to a particular task crucial for their members' survival (Gallistel, 1990; Marler, 1970).

But what about humans? Throughout this chapter, we've emphasized that a great deal of human behavior—just like the behavior of every animal species—is governed by principles of habituation as well as classical and operant conditioning. But, even so, some of our behavior is the product of distinctly human forms of learning. One example involves the processes through which humans learn language. These processes seem controlled by innate mechanisms that guide the learning and make it possible for us to achieve remarkable linguistic competence by the time we're 3 years old. (We'll say much more about language learning in Chapter 10.) Humans also have remarkable inferential abilities that allow us to gain broad sets of new beliefs, based on events we've observed or information we've received from others; and these new beliefs can profoundly affect our behavior. In these ways, humans are capable of distinctive forms of learning at the same time that they're powerfully shaped by the more general principles discussed in this chapter. (We'll have much more to say about human inferential capacities, and how we're influenced by others, in Chapters 8 and 9.)

Similarities in How Different Species Learn

In short, there are certainly differences—as well as crucial similarities—in how species learn, and, as we've noted, the differences make good biological sense. After all, each species lives in its own distinctive environment and needs its own set of skills, and so it may need to learn in its own ways. But what about the similarities? After all, the rats and pigeons we study in the laboratory don't gather food the way a human does. They don't communicate with their fellows the way a human does. Their nervous systems are much simpler than ours. It wouldn't be surprising, therefore, if they learned in different ways than we do. Yet, as we've repeatedly noted, the major phenomena of both classical and instrumental conditioning apply across species—whether we're considering humans, rats, pigeons, cats, dogs, fish, or even some types of snails (Couvillon & Bitterman, 1980; E. Kandel, 2007).

How should we think about this point? Why do such diverse creatures share certain types of learning? The answer lies in the fact that all of these creatures, no matter what their evolutionary history or ecological niche, share certain needs. For example, virtually all creatures are better off if they can prepare themselves for upcoming events, and to do this they need some way of anticipating what will happen to them in the near future. It's no wonder, then, that many species have developed nervous systems that support classical conditioning.

Similarly, in the world we all inhabit, important outcomes are often influenced by one's behavior, so it pays for all species to repeat actions that have worked well in the past and to abandon actions that haven't succeeded. Hence we might expect natural selection to have favored organisms capable of learning about the consequences of their actions and able to adjust their future behavior accordingly. In other words, we'd expect natural selection to have favored organisms capable of instrumental conditioning.

Of course, people are different from pigeons, and pigeons from sea slugs; no one is questioning these points. Even so, it seems that there are some types of learning that all species need to do. And this is why, in this chapter, we've easily identified principles of learning that apply to an extraordinary range of species and settings.

THE NEURAL BASIS FOR LEARNING

We just suggested that different species all need to learn the same kinds of lessons—including how events in the world are related to each other, and what the consequences are of the organism's own actions. But do all organisms accomplish this learning through the same neural mechanisms? If we zoom in for a closer look, will we find the biology of learning is the same from one species to the next? As is often the case, there are important commonalities between species; but there are also significant differences: The biological mechanisms that allow learning in mammals are somewhat different from the mechanisms crucial for reptiles, amphibians, or invertebrates (Macphail, 1996; Woolf, 1998). Indeed, even within a single species, the biological mechanisms needed for learning can vary and seem to depend on the CS, the US, and the procedure (R. Clark, Manns, & Squire, 2003; Fanselow & Poulos, 2005). Thus, the brain circuits underlying fear conditioning (with electric shock as the US) are centered in the amygdala; the brain circuits underlying eyeblink conditioning (with a puff of air to the eye as the US) are centered in the cerebellum. Conditioning with a long delay between the CS and US typically involves the hippocampus, while conditioning with a shorter delay may not (Berman & Dudai, 2001; Lattal, Honarvar, & Abel, 2004).

Even with these variations, some biological principles do apply to all cases of learning. In all cases, learning depends on *neural plasticity*—the capacity for neurons to change the way they function as a result of experience. In all cases, the plasticity involves changes at the synapse—that is, changes in the way neurons communicate with each other. These changes, in turn, can involve any of three adjustments: Some neurons, after learning, end up sending a stronger signal than they did before. Other neurons become more sensitive to the signals they've been receiving all along. And, third, learning can lead to the creation of entirely new connections among neurons—new synapses—allowing for new lines of communication within the nervous system.

Evidence for these points comes from many sources, including studies of the marine mollusk *Aplysia*. Because the nervous systems of these creatures contain a mere 20,000 neurons, they're good candidates for detailed analysis. Researchers have been able to document that, after conditioning, the *Aplysia*'s sensory neurons—the neurons that receive the CS—literally release more neurotransmitter into the synapse than they did before the conditioning trials. This is a crucial part of why, at the end of learning, these neurons are able to trigger a response—the CR—that they couldn't trigger at the start. This process, which can be documented in many species in addition to *Aplysia*, produces an increase in the neural signal being sent, and is called **presynaptic facilitation** (e.g., Lisman, 2003; Pittenger & Kandel, 2003).

presynaptic facilitation A process, documented in studies of *Aplysia*, that underlies many kinds of learning. It occurs when learning results in an increased release of neurotransmitter into the synapse.

Other forms of neural plasticity, in other organisms, involve postsynaptic changes—that is, they influence the receiving side of the synapse. A particularly important mechanism in this category is **long-term potentiation** (**LTP**; T. Bliss & Lomo, 1973; T. Bliss, Collingridge, & Laroche, 2006; Martinez & Derrick, 1996)—*potentiation* because the mechanism involves an increase in the responsiveness of a neuron (an increase in the neuron's potential for firing) and *long term* because this potentiation lasts for days, perhaps even weeks.

long-term potentiation (LTP) A long-lasting increase in a neuron's response to specific inputs, caused by repeated stimulation.

LTP is produced when one neuron activates another neuron over and over. The repeated stimulation causes the postsynaptic neuron to become more sensitive to this input than it was before, so it's more likely to respond to this input in the future. In addition, the increased responsiveness can spread to other nearby neurons. Here's an example: Let's suppose that within some brief period of time neuron A in Figure 7.34 repeatedly causes neuron C to fire. This will cause neuron C to become more responsive to A than it was initially. But, of course, C also has other synapses that receive input from other neurons, such as neuron B in the figure. These receptors too will become more sensitive as a result of neuron A's repeated activity, provided that these other neurons fire at the same time as neuron A. In other words, the spread of potentiation is *activity dependent* and so will spread to neuron B only if B was active at the same time as the neuron that caused the potentiation in the first place—in this case, neuron A (W. Levy & Steward, 1979; McNaughton, Douglas, & Goddard, 1978). In this way, LTP provides a cellular mechanism through which associations—in our example, the association between A's and B's activity—can be detected and recorded in the brain (Fanselow & Poulos, 2005; E. Kandel & Hawkins, 1992; Koekoek et al., 2003; Martinez & Derrick, 1996).

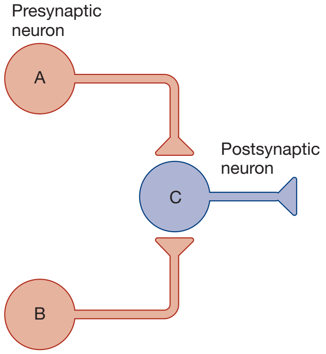

7.34 **Long-term potentiation** If neuron A repeatedly causes neuron C to fire, the strength of their connection (synapse) will increase (become potentiated).

Presynaptic facilitation and LTP both involve changes in how efficiently a synapse functions—the first involves changes in how neurons send signals, and the second involves changes in how neurons receive signals. A third form of neural plasticity involves the creation of entirely new synapses. These changes seem to take place largely on the dendrites of the postsynaptic neurons; as we mentioned in Chapter 3, the dendrites can grow new *dendritic spines* (Moser, 1999; Woolf, 1998). These spines are the "receiving stations" for most synapses; so growing more spines means that, as learning proceeds, the neuron is gaining new lines of communication with its cellular neighbors.

SOME FINAL THOUGHTS: LEARNING THEORY AND BEYOND

The investigations begun by Pavlov and Thorndike more than 100 years ago have led to many important discoveries, and it's clear that some of the principles they uncovered apply to every species we have studied. We can demonstrate these principles in the laboratory, and we can employ them outside the lab. For example, behavior modification techniques, building on the lessons we've learned in this chapter, are widely used in hospitals for the mentally ill and in one-on-one therapy; many prisons use instrumental conditioning to shape prisoners' behavior; parents employ similar techniques to teach their children good manners, household chores, and interpersonal skills. In short, these principles of learning are extensively used and enormously useful, and they're an effective means of altering or maintaining specific patterns of behavior.

At the same time, we need to place these principles of learning in a broader context. On the one side, we need to connect our understanding of learning to the relevant biology. This allows us to explore the mechanisms that make learning possible—and, with that, to see the *diversity* of mechanisms (even within a single species) that underlie learning. A link to biology also calls our attention to the evolutionary roots of learning, and this helps us comprehend why rats are biologically prepared to associate illness with tastes—and presumably, why humans are biologically prepared to associate snakes with aversive outcomes.

On the other side, we also need to connect our conception of learning to the relevant cognition. As we've seen, learning involves more than a change in behavior; it also involves the acquisition of new knowledge. In classical conditioning, this knowledge involves the relationships among events in the world; that's why the CS serves as a *signal* for things to come. In instrumental conditioning, the knowledge concerns the relationships between one's own actions and events in the world, but it can also include other, broader topics: knowledge about complicated spatial layouts, knowledge about whether an organism has control over its environment or not. But what is this knowledge? How is it stored in memory? How is the knowledge *used* as a basis for solving problems or making decisions? We plainly are not going to understand learning unless we tackle these issues, and this sets part of the agenda for us in the next two chapters.

SUMMARY CHAPTER 7

THE PERSPECTIVE OF LEARNING THEORY

- The empiricist philosophers argued that learning involves the forming of simple associations. More complex learning just involves a great many associations, each layered upon the others. From this perspective, all learning depends on the same mechanisms, and so all learning should be governed by the same principles.

HABITUATION

- The simplest form of learning is *habituation,* a decline in the response to stimuli that have become familiar through repeated exposure. In *dishabituation,* the organism learns that a previously predictable stimulus has now changed, causing the organism to renew its attention to the stimulus.

CLASSICAL CONDITIONING

- In *classical conditioning,* animals learn about the association between one stimulus and another. Before conditioning, an *unconditioned stimulus* (US, such as food) elicits an *unconditioned response* (UR, such as salivation). After repeated occasions on which the US follows a *conditioned stimulus* (CS, such as a buzzer), this CS alone will begin to evoke the *conditioned response* (CR; here again, salivation).
- When a CS-US relationship is well established, the CS can be preceded by a second, neutral stimulus to produce *second-order conditioning.*
- Trials in which the CS is presented without the US lead to *extinction.* However, the phenomenon of *spontaneous recovery* shows that the CR is masked, not abolished, by extinction.
- Because of *stimulus generalization,* the CR can also be elicited by stimuli that are similar to the CS. To train the animal to *discriminate* among stimuli, one stimulus (CS^+) is presented with the US, while another (CS^-) is presented without the US.
- Several lines of evidence suggest that the CS serves as a *signal* for upcoming events. This fits with the fact that learning is less likely if the CS is simultaneous with the US, or (worse) follows it. The signal value of the CS is also evident because learning occurs only if there is some contingency between CS and US; mere contiguity between these stimuli isn't enough. Animals can also learn about the *absence* of contingency, and animals clearly prefer environments in which stimuli are predictable.
- Animals seem sensitive to relationships among probabilities, but this is not because they're tracking the probabilities directly. Instead, animals develop expectations about upcoming events and adjust their expectations whenever events surprise them. In this way, the expectations are gradually adjusted until they are accurately in tune with probabilities in the environment. The role of surprise is directly evident in the *blocking effect.*
- The CR is not identical to the UR. Instead, the CR seems to be a means of *preparing* for the US. Sometimes the preparation takes the form of a compensatory response, and this point may be crucial in understanding drug addiction and drug tolerance.

INSTRUMENTAL CONDITIONING

- When training an animal using *instrumental* (or *operant*) *conditioning,* the trainer delivers a reward or *reinforcement* only after the animal gives the appropriate response. According to Thorndike, learning in this situation is governed by the *law of effect,* which states that the tendency to perform a response is strengthened if it's followed by a reward and weakened if it's not.
- *Operants* are voluntary responses, strengthened by reinforcement; but acquiring them may call for some initial *shaping,* through a method of *successive approximations.*
- Some reinforcers are stimuli whose reinforcing power is unlearned. Other conditioned reinforcers acquire their power from prior presentations with stimuli already having that capacity. The magnitude of a reinforcer depends on several factors, including the magnitude of other reinforcers that might be available. This effect, which is reflected in the phenomenon of *behavioral contrast,* may be one source of findings sometimes attributed to *intrinsic motivation.* Many theorists, however, believe that intrinsic motivation involves a separate set of principles, different from those that govern operant conditioning.
- During *partial reinforcement,* the response is reinforced only some of the time. The rule that determines when a reinforcer is given is called a *schedule of reinforcement.* In *ratio schedules,* reinforcement is delivered after a number of responses; the ratio used may be fixed or variable. In *interval schedules,* reinforcers are delivered for the first response made after a given interval since the last reinforcement; this interval, too, can be fixed or variable.
- Learning involves more than a change in behavior; it also involves the acquisition of new knowledge. This principle is evident in many settings, including *latent learning.*
- Operant conditioning results when reinforcement is contingent on a response, not just when reinforcement happens to be contiguous with responding. Organisms' sensitivity to contingency can be demonstrated in the phenomenon of *learned helplessness,* where animals seem to learn they have no control over the events they're experiencing.

OBSERVATIONAL LEARNING

- Many animals can learn simply by watching other individuals and learning from their example. This is evident in *vicarious conditioning,* and it's also clear in learning through *imitation.* However, the impact of media violence reminds us that imitation can be a source of undesired behaviors as well as a source of new skills.

VARIETIES OF LEARNING

- According to the early learning theorists, just about any CS can become associated with any US, and just about any response can be strengthened by any reinforcer. This assertion is challenged by the fact that certain conditioned stimuli are more readily associated with some unconditioned stimuli than with others, as shown by studies of *taste aversion learning.* These studies suggest that animals are biologically prepared to learn certain relations more readily than others. Similar effects occur in instrumental conditioning, where some responses are more readily strengthened by some reinforcers than by others.
- Certain forms of learning are species specific. In humans, specialized forms of learning may include our capacity for learning language as well as our remarkable ability to learn by observing others.
- Animals vary in how (and what) they learn; but in some aspects of learning, diverse species also exhibit striking similarities. These similarities probably arise because all organisms live in the same world, and the nature of the world creates a need in many species for the forms of learning we call classical and operant conditioning.

THE NEURAL BASIS FOR LEARNING

- In recent years, investigators have made considerable progress in understanding the neural bases for learning. These bases involve diverse mechanisms, such as *presynaptic facilitation,* and *postsynaptic changes,* such as *long-term potentiation* (*LTP*). Still another mechanism involves the creation of new synapses—made possible by the growth of new dendritic spines, which act as "receiving stations" for synapses and open new lines of communication between cellular neighbors.

ONLINE STUDY TOOLS

Go to StudySpace, **wwnorton.com/studyspace**, to access additional review and enrichment materials, including the following resources for each chapter:

Organize

- Study Plan
- Chapter Outline
- Quiz+ Assessment

Learn

- Ebook
- Chapter Review
- Vocabulary Flashcards
- Drag-and-Drop Labeling Exercises
- Audio Podcast Chapter Overview

Connect

- Critical Thinking Activity
- Studying the Mind Video Podcasts
- Video Exercises
- Animations
- **ZAPS** Psychology Labs