Chapter 7

LEARNING

CHAPTER PREVIEW QUESTIONS

- What are the three basic kinds of learning?
- How did Ivan Pavlov discover classical conditioning?
- Who made Little Albert afraid of the rat and how was it done?
- How are behaviors learned and unlearned in classical conditioning?
- What did Robert Rescorla, Leon Kamin, and John Garcia, add to our understanding of classical conditioning?
- Does classical conditioning explain all human learning?
- How are behaviors learned and unlearned in operant conditioning?
- How does the schedule of reinforcement affect operant conditioning?
- Can punishment be an effective method of eliminating undesirable behaviors?
- Are there problems or negative side effects associated with the use of punishment?
- Can we learn without undergoing classical or operant conditioning?
- Does watching violent shows make children more violent?

HOW DO PSYCHOLOGISTS DEFINE LEARNING?

The principle of learning can help explain much of our everyday behavior. We are born knowing how to perform some behaviors, such as breathing, sneezing, and coughing. Learning can modify even these biologically programmed responses, however. People can learn, for example, to cover their mouths when they sneeze or cough. One of the main things that makes us different from other animals is our ability to learn very complex behaviors. Humans can learn to play chess, dance a ballet, or write a novel. These skills are far beyond even the most intelligent of other species. It has been said that learning is the most hopeful of the fields of psychology. We hope that humans can learn to solve their personal problems, overcome their prejudices, and learn how to take care of the earth.

"Learning" as a technical term

learning
A relatively permanent change in behavior due to experience

Unlike some psychological terms, **learning** is a word that the average person uses all the time. Even small children will say, "Look, I 'learned' how to tie my shoes." As scientists, however, we need to have a more exact definition of the word "learning." Psychologists do have a formal definition of learning but the meaning is not very different from the way the word is used in everyday life. Here is the scientific definition of learning: *A relatively permanent change in behavior due to experience.* It seems obvious that a person's behavior must change if they learn something but we don't want to include *all* changes in behavior in our definition. First, the change must be "relatively permanent." When a person gets hungry, they may start to eat. When they are full, they will usually stop. These are changes in behavior, but we don't want to say that the person has "learned" to eat or stop eating several times each day. There are many temporary changes in behavior such as eating, sleeping, and getting angry, that don't qualify as learned behaviors. Since these are not "relatively permanent" changes, our definition excludes them.

Q: Why say that learned behaviors must be "due to experience?"

Can you think of some behaviors that are relatively permanent but wouldn't qualify as learning? During our early years, our ability to reach up a grab objects held high in the air increases steadily. This ability to reach to greater and greater heights is a change in behavior and it is relatively permanent. This change isn't learned, however. It is the result of maturation. By defining learned behaviors as resulting from experience, we exclude changes in behavior that are due to maturation, disease, drugs, or injury. These changes in behavior may be relatively permanent but we can't say that they are learned.

Two important terms—Stimulus and Response

stimulus
Information perceived through the senses

response
Any activity of the muscles or other identifiable behavior

Now that we have a scientific definition of learning, let's look at two important terms: *stimulus* and *response.* In order to understand the following sections on classical and operant conditioning, it is important that you become familiar and comfortable with these two scientific terms. A **stimulus** is anything that comes in through your senses. It could be something simple like a smell, a light, a bell, or a tone. There are also complex stimuli like the contents of a book or lecture. In psychology, we usually study simple stimuli. Responses can also be simple or complex. A **response** is anything that goes out through your muscles—anything you do. The responses we study in psychology also tend to be simple ones like blinking, salivating, or pressing a lever. In real life, we make complex responses when we drive to work, or read a book. Scientists often study limited situations in the laboratory. In the study of physics, for example, a scientist might study the behavior of a small ball rolling down a ramp. He or she hopes that the principles discovered in this simple experiment will apply to complex situations found in everyday life such as the movement of cars and airplanes. Psychologists hope that once we understand the relationship between simple stimuli and responses, we can begin to understand the more complex ones found in everyday human behavior outside the laboratory.

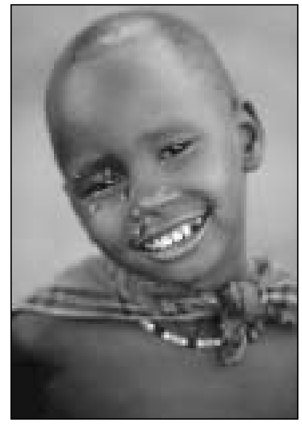

This African child has become habituated to the sensation of flies landing on his face and eyes

Habituation—A simple form of learning

In some parts of Africa, small flies gather around people's heads and often land on the surface of their eyes. Having a fly land on your eye will usually result in the blink reflex. The reflex is designed to protect your eyes from injury. In the parts of Africa where these flies are common, however, the flies often land and even walk around on the surface of people's eyes without triggering the blink reflex. As a scientist, you might wonder why. This ability to keep from blinking is due to a simple form of learning called **habituation.** It happens whenever a stimulus is presented over and over. As the stimulus continues to be presented, the response to it grows weaker. Notice that habituation fits our definition of learning very well. It is relatively permanent and it is due to experience. Habituation is a very simple form of learning. For habituation to occur, we need only to present a stimulus over and over. Our lives are full of repeated stimuli like ticking clocks, ventilating fans, the sound of our own breath, and our heartbeats. It is easy to see the survival value of being able to ignore these repeated stimuli and focus our attention on more important things. Although habituation is a form of learning, it is a very primitive kind. Most psychologists don't find it very interesting. Compared to the other kinds of learning, little research has been done on habituation. Most psychologists who study learning concentrate on the following three kinds: *classical conditioning, operant conditioning,* and *observational learning.* See Table 8.1 for a brief overview of these three kinds of learning.

habituation
A decrease in response to a repeated stimulus

	KEY NAMES	TYPICAL BEHAVIOR LEARNED	HOW THE BEHAVIOR IS LEARNED	EXAMPLE
CLASSICAL CONDITIONING	Ivan Pavlov, J. B. Watson	Reflexes and Emotional Behaviors	Repeated pairing of two stimuli	Learning to fear the sound of the dentist's drill
OPERANT CONDITIONING	E. L. Thorndike, B. F. Skinner	Behaviors other than reflexes and emotional behavior	Following a response with reinforcement or punishment	Dog learning the sit in order to receive praise
OBSERVATIONAL LEARNING	Albert Bandura	Any observable behavior	Observing the behavior of others and its consequences	Becoming more violent by watching violent videos

TABLE 7.1

LEARNING: AN OVERVIEW

For each of the three kinds of learning, classical conditioning, operant conditioning, and observational learning, the behaviors learned and the conditioning procedure are different.

PAVLOV AND CLASSICAL CONDITIONING

Ivan Pavlov (center) and his staff demonstrating classical conditioning with a dog

The Russian physiologist, **Ivan Pavlov** devoted his entire life to science. Pavlov became famous for the scientific study of digestion in dogs. His experiments showed that digestion started in the mouth and that saliva was an important part of the digestive process. This led him to the discovery of the salivary gland and the salivary reflex. Pavlov found that putting food powder on a dog's tongue would trigger the salivary reflex (Pavlov, 1897/1902). He won a Nobel Prize in 1904 for this work. It was another discovery, though, that made him famous in the field of psychology. In his laboratory, Pavlov began to notice that some of his dogs were starting to salivate *before* he put the food powder on their tongues. Some would salivate at the sight of the food powder or even the sight of the spoon used to deliver the powder. Some of the dogs even began to salivate when they saw Pavlov's assistant bringing in the food (Vul'fson, 1898 as cited in Todes, 1997). When Pavlov investigated further, he found that the longer the dogs had been in the laboratory, the more likely they were to make these astonishing responses. New dogs would only respond to the food powder itself. Pavlov had thought of the salivary reflex as being like an electrical circuit. He thought that putting the food on the dog's tongue completed the circuit and caused the dog to salivate. Pavlov believed that the process was just like the way flipping a light switch completes a circuit and turns on the light. He was amazed, then, when the response occurred *before* the food powder arrived on the dog's tongue. Imagine how you would feel if you went to turn on the lights in your living room and you noticed the lights coming on before you even touched the switch. Pavlov was astounded by this new development. He concluded that the dogs were learning to respond to stimuli other than the food powder. Like many good scientists, he was so fascinated by an unexplained event in the laboratory that he changed his whole area of study. He stopped working on digestion and began studying what we now call **Classical Conditioning.** In Pavlov's honor, this kind of conditioning is sometimes called *Pavlovian conditioning.*

classical conditioning
Learning that results from the pairing of an unconditioned and a conditioned stimulus

Pavlov's experiments showed that his dogs could be conditioned to salivate to a number of other stimuli. He tried the sound of a metronome and a number of other stimuli including a small bell. All he had to do was present these other stimuli repeatedly along with the food powder. Repeatedly presenting food powder and ringing a bell at the same time, for example, will eventually result in a dog salivating to the sound of the bell. This is the basic form for all classical conditioning. One stimulus (in this case food) already produces the response (salivation). This stimulus is paired, or presented together with, a neutral stimulus (the bell). Before conditioning, the bell does not produce the response. Over time, as the two stimuli are presented together, the bell comes to produce the response (salivation). We could say that the dog has learned to salivate to the sound of the bell. Instead, however, we usually use a more scientific term. We say that the dog has been *conditioned* to salivate to the sound of the bell. Another scientific term we use in describing classical conditioning is the word "elicit." We say that before conditioning, the food *elicits* salivation. After conditioning, we say that the bell *elicits* salivation. When responses are elicited, they are automatic and involuntary.

Notice that there are four stimuli and responses here (two stimuli, two responses). The bell and the food powder are the two stimuli. Salivation to the food and salivation to the bell are the two responses (see Figure 7.1). We have scientific terms for these four elements of classical conditioning but the terms are sensible and easy to learn.

Because the food powder produces the response before we have done any conditioning, we could say that it works on an "unconditioned" animal. For

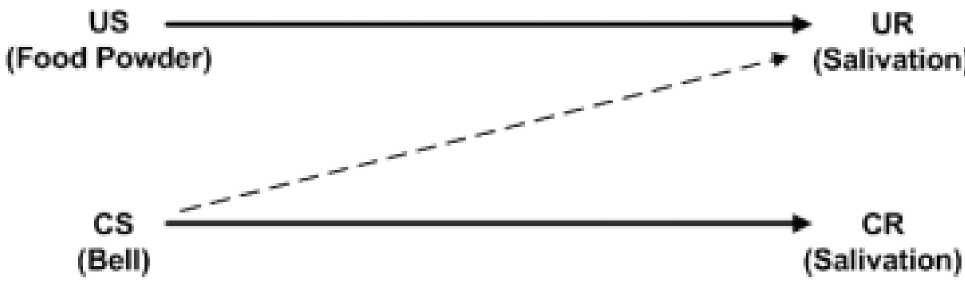

FIGURE 7.1

THE PROCESS OF CLASSICAL CONDITIONING

With repeated pairings of the conditioned stimulus and unconditioned stimulus, the conditioned stimulus comes to elicit the conditioned response

this reason, we call it the **unconditioned stimulus** (US). In every example of classical conditioning, one stimulus works before any conditioning has taken place. It is always called the unconditioned stimulus. The neutral stimulus (the bell in our example) won't elicit the response until after the conditioning process. It will only work on a conditioned animal. Naturally, we call it the **conditioned stimulus** (CS). The responses are named according to the stimulus that elicits them. Salivation to the unconditioned stimulus (the food powder) is called the **unconditioned response** (UR). Salivation to the conditioned stimulus (the bell) is called the **conditioned response** (CR).

unconditioned stimulus (US)
A stimulus that elicits a response before any conditioning has occurred

conditioned stimulus (CS)
A stimulus that elicits a response after being paired with an unconditioned stimulus

unconditioned response (UR)
A response elicited by an unconditioned stimulus

conditioned response (CR)
A response elicited by a conditioned stimulus

Little Albert and the Rat

Let's look at a famous example of classical conditioning. In 1920, behaviorist **J. B. Watson** published a paper on "Conditioned Emotional Reactions." The paper was based on a series of scientific experiments Watson performed with his assistant, Rosalie Rayner. Watson did not believe that infants were naturally afraid of animals (Watson, 1919). He thought that their fears were learned through classical conditioning and set out to demonstrate that this was the case. Watson's experiments involved a small child he referred to as Albert B., but who is usually referred to as "Little Albert." Watson wanted to show how fear of specific objects or situations could be learned and Little Albert was a perfect subject. Before any conditioning took place, Little Albert was not at all afraid of rabbits, rats, or other furry animals. He was a very calm child and very seldom cried and he was especially fond of a white laboratory rat. Watson wanted to show how Little Albert could be conditioned to fear the rat.

Remember that classical conditioning always starts with a stimulus (the US) that produces the response before the start of the conditioning process. Since Watson wanted to produce the response of fear, he needed a stimulus that would scare Little Albert. Watson discovered that he could make Little Albert cry by standing behind him and hitting a steel bar with a hammer. As with all cases of classical conditioning, this example has two stimuli and two responses (see Figure 7.2). When Watson conditioned Little Albert, the stimuli were the loud noise (US) and the rat (CS). The responses were fear of the loud noise (unconditioned response) and fear of the rat (conditioned response).

Watson placed Little Albert on top of a table with the rat. When Little Albert looked at the rat, Watson would strike the steel bar with the hammer (pairing the unconditioned stimulus and the unconditioned response). After seven trials where the rat and the loud noise were presented together, Watson tried presenting the rat alone. Little Albert began crying and began crawling away from the rat so fast that Watson and his assistant had trouble catching him before he reached the edge of the table. We might say that Little Albert had "acquired" his fear of the rat. This part of classical conditioning is called

John Watson, Rosalie Rayner, and Little Albert

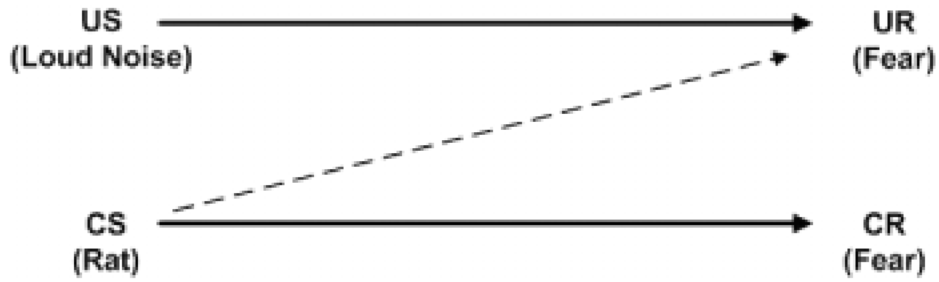

FIGURE 7.2

WATSON'S CONDITIONING EXPERIMENT

In Watson's experiment, the conditioned stimulus (the rat) came to elicit the conditioned response (fear)

acquisition because pairing the unconditioned stimulus and conditioned stimulus causes the acquisition of the response to the conditioned stimulus.

Shortly after Watson's experiment, Little Albert's mother removed him from the hospital where Watson was doing his experiments (Harris, 1979). Was he afraid of rats and other white furry animals for the rest of his life? We don't know. The experiment was traumatic for Little Albert and may have had a lasting effect. Because of this, it would be unethical to repeat Watson's experiment so we can only speculate on the long-term effects of this particular example of classical conditioning.

Watson was thrilled by the results of his experiment with Little Albert. Based on this, and other experiments he performed, he concluded that classical conditioning was the basis for all human learning. He believed that he could condition any behavior and that he could completely shape a person's life using only classical conditioning. We now know that Watson was too optimistic about the power of classical conditioning. Many human behaviors are learned in other ways. Still, classical conditioning does play an important role in conditioning reflexes and emotional behaviors. It gives a scientific explanation of why children become afraid of the dark and why many of us tense up at the sound of a dentist's drill. We'll see other examples as we discuss the details of classical conditioning.

Here is an example of classical conditioning that did not happen in a laboratory. A friend of mine, a psychologist, spent many years watching a famous TV news anchor on the evening news every night while eating dinner. To this day, when that particular news anchor appears on his television screen, he salivates. In this case, the news anchor is the conditioned stimulus. Salivation to the television image of the news anchor is the conditioned response (J. C. Megas, personal communication, May 3, 1998).

Let's look at one more example of the acquisition phase of classical conditioning. Suppose that you are tired of just reading about classical conditioning and, as a scientist, you decide that you actually want to try it. First, you need a subject. You don't have a dog but you do have a little sister and you decide that she will be your subject. Now you need a response to condition. You don't want to scare your little sister and you think that salivation is too messy. You know that classical conditioning works well on reflexes and try to think of a reflex that you can condition. You decide on the blink reflex. You blow a puff of air into your little sister's eye and every time you do so, she blinks. You have your unconditioned stimulus and unconditioned response. Now you need a neutral stimulus to serve as the conditioned stimulus. You remember that Pavlov used a bell and search your house for a bell. You can't find one but you finally think of your doorbell. You drag your little sister into the doorway and make sure that she doesn't already blink to the doorbell. Now you present the doorbell and the puff of air together for a number of training trials. Soon, your little sister is blinking every time someone rings the doorbell. You have performed classical conditioning and your scientific experiment is a success.

Locking It In

1. In classical conditioning, the CS will elicit the _____.
2. ______________ started out studying digestion in dogs.
3. ______________ performed conditioning experiments on 'Little Albert.'
4. In classical conditioning the unconditioned response is elicited by the _____.
5. Critical Thinking: Would Watson's results have been different if he tried to teach 'Little Albert' to be afraid of something other than a rat? Why?

Answers: 1) conditioned response, 2) Ivan Pavlov, 3) J. B. Watson, 4) unconditioned stimulus

Classical Extinction

Q: Can classically conditioned responses be "unlearned?"

If Little Albert had not been removed from the hospital, Watson might have wanted to help him lose his fear of the white rat. You might begin to feel sorry for your little sister who now blinks whenever the doorbell rings. Is there a scientific method that will undo classical conditioning? In classical conditioning, the response will get weaker and weaker if we present the CS over and over without the US. This process is called **extinction** (pronounced "ex-STINK-shun"). If Watson had repeatedly put Little Albert together with the white rat without making the loud noise, it is likely that Little Albert would have gradually lost his fear of the rat. If you ring the doorbell over and over without the puff of air, your little sister's blink response will grow weaker and weaker. Let's return for a moment to Pavlov's experiment. During acquisition, the US (food) was paired with the CS (bell). What would extinction look like in this example? In extinction, we present the CS alone. As we ring the bell (the CS) over and over, the response gets weaker and weaker and eventually disappears (see Figure 7.3). We now have a scientific definition of Classical extinction: *The presentation of the conditioned stimulus over and over without the unconditioned stimulus.*

extinction
Eliminating a previously conditioned response

Spontaneous Recovery

Q: Once the response fails to follow the CS, is it gone for good?

As you can see from Figure 7.3, a response will often reappear after a period of rest. This is called spontaneous recovery. It is normal in classical conditioning (and in operant conditioning as well). It is wise to remember that **spontaneous recovery** is a normal part of the extinction process. Another important thing to note is that during the extinction process, the response will be occurring quite often. If you are trying to eliminate a habitual behavior in a child, a pet, or even yourself, it is normal for the response to occur many times and to reappear occasionally in the future. This happens in both classical and operant conditioning. It doesn't mean that the subject is trying to be difficult or uncooperative. It just means that they are following the normal scientific process of extinction. It may help you to be patient in such situations if you remember how the process works.

spontaneous recovery
The reappearance of a conditioned response following extinction

Generalization

As a scientist, you might wonder what would have happened if Pavlov's bell had broken and he couldn't find another one exactly like it? Would the dogs salivate to another bell that was very similar to the bell that was used during acquisition?

FIGURE 7.3

ACQUISITION, EXTINCTION, AND SPONTANEOUS RECOVERY

This graph shows a response being acquired, going through extinction, and bouncing back during spontaneous recovery after a rest period

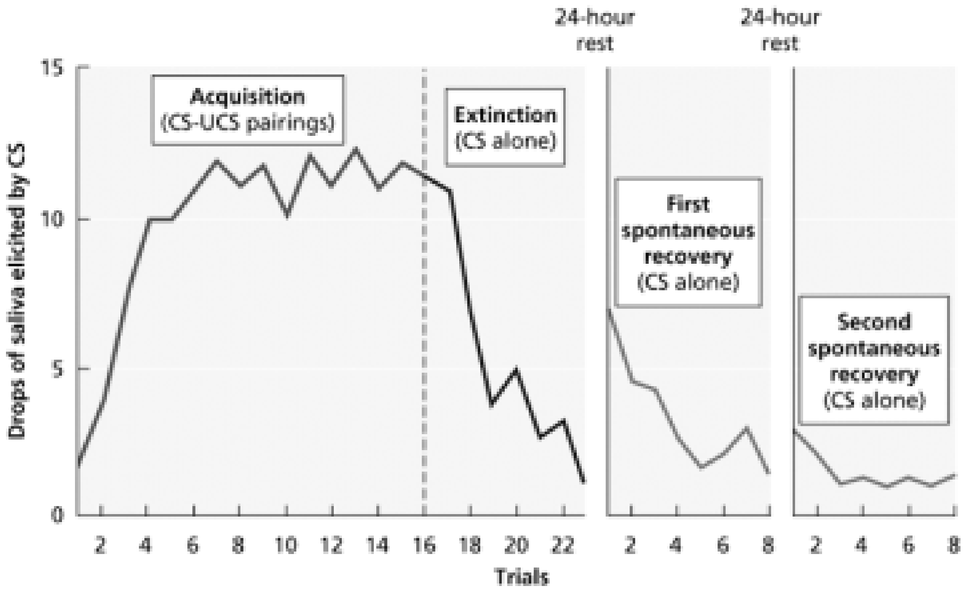

generalization
The tendency to respond to another stimulus that is similar to the training stimulus

As you can probably guess, they would. This is called **generalization** and it applies to any classically conditioned response. Animals trained to respond to a particular stimulus will respond to any other stimulus that is similar to it. The more similar the stimulus, the stronger the response will be. A bee stung my niece on the sole of her foot. The bee was the CS and her fear of bees is the conditioned response. She's not just afraid of that particular bee, though. She is afraid of all bees. The more similar they are to the bee that stung her, the more afraid she is.

Imagine that we conditioned a dog to salivate to the sound of middle C on the piano. Through the process of generalization, the dog would also salivate to the nearby keys on the keyboard. As we play keys further and further away from middle C, we would expect the response to grow weaker and weaker. Little Albert also showed generalization after learning to be afraid of the white rat. Five days after his training, Little Albert reacted with fear to other furry stimuli including a white rabbit, a dog, and even a white Santa Claus mask (Harris, 1979).

Discrimination

discrimination
The tendency to respond differently to two or more stimuli

discrimination training
Training an animal to respond differently to two different stimuli

In generalization, an animal makes the same response to two similar stimuli. If the animal responds to one stimulus but not to another, we say that the animal is discriminating between the two stimuli. This is called **discrimination.** Discrimination is the opposite of generalization. Suppose that our dog has been trained to salivate to the sound of middle C on the piano. Imagine that the dog salivates equally to C and to C#, the key next to middle C. As a scientist, you might wonder whether this is because of generalization, or because the dog can't hear the difference between these sounds? How could we tell? One way to find the answer to our question is to use a scientific technique called **discrimination training.** In discrimination training, one stimulus is presented with the US and the other is presented without it. In our example, we might continue to present Middle C with the food powder some of the time. Other times, we would present C# by itself. If the dog can tell the two stimuli apart, its response to middle C should grow stronger. At the same time, its response to C# should go through extinction since we are presenting the CS without the US. Dogs *can* distinguish between the tones made by different piano keys. Discrimination training allows scientists to test the discrimination abilities of animals and babies who can't use language to tell us whether they can tell the difference between similar stimuli.

Locking It In

1. Eliminating a classical response by presenting the CS without the US is called __________________.
2. When an animal responds to a new stimulus similar to the training stimulus, it's called ___________________.
3. If an animal responds differently to two different stimuli, it's called ________________.
4. When a response comes back after extinction, it's called ___________ ___________.
5. Critical Thinking: Why does spontaneous recovery occur?

Answers: 1) extinction, 2) generalization, 3) discrimination, 4) spontaneous recovery

Conditioned Responses in Everyday Life

Does this picture make you nervous?

We said earlier that classical conditioning can result in reflexes and emotional responses in our daily lives. Let's look at a few examples. We know that very young children are often unafraid of snakes, heights, spiders, rats, dentist's drills, and a number of other stimuli that can frighten us as adults. When I list these common fears in class and ask my students to raise their hands if they are afraid of at least one of these, usually every hand in the class goes up. How would a scientist explain these fear responses? Based on what we know of classical conditioning, it seems likely that these stimuli were paired with unconditioned stimuli that caused fear earlier in the person's life. It's easy to imagine, for example, a small child held in a parent's arms when a snake, rat, or spider appears unexpectedly. The parent might yell or let the child fall an inch or two. These unconditioned stimuli can elicit fear in a small child and would explain the child's fear of the conditioned stimulus. These conditioned emotions are called **conditioned emotional responses** and they can sometimes be quite strong. Some people are so conditioned that they begin to perspire and shake with fear when they see a photograph or drawing of a snake, rat, or spider.

conditioned emotional response
An emotional response learned through classical conditioning

Often, conditioned emotional responses can be learned in a single trial and can last for a long time. In 1994, a man with an iron bar attacked Olympic Skater Nancy Kerrigan. It happened during skating practice. The man came from behind and struck her right knee with the bar. In an interview in 1998, the skater reported that for years afterward, "If someone suddenly popped into her peripheral vision, in that certain spot slightly behind her right shoulder, then the awful memory would come flooding back and Nancy Kerrigan's defense mechanisms would be charged, racing." (Dupont, 1998)

Ice skater Nancy Kerrigan was the victim of a physical assault at the Olympic Games in 1994. As a result of the attack, she developed a conditioned emotional response that is triggered whenever she experiences stimuli similar to those present during her attack

You may have had a frightening car accident at some time in your life. It is likely that you were afraid to drive for some time afterward. Because most of us continue to drive, a response like this usually goes through extinction over time. Whether the response is extinguished depends on how often the frightening stimulus occurs in our lives. The more often we encounter the conditioned stimulus, (CS) without the fear-producing unconditioned stimulus (US), the more likely the response will grow weak and finally disappear.

Q: What about stimuli that you can avoid. Do these go through extinction?

Often they do not. People with a fear of snakes, for example, will usually go out of their way to avoid snakes. If their local zoo has a snake exhibit, they won't look at it. In addition, when they do encounter a snake, they usually run away from it. This reduces their fear and makes them feel better. Feeling better rewards them for being afraid and for running away. This is why phobias, irrational fears of objects or situations, can be so persistent.

Q: Could we use extinction to get rid of these fears?

systematic desensitization
A gradual process of extinguishing a fear response through classical extinction

Yes we could, in fact there is a common scientific technique called **systematic desensitization** based on the principle of classical extinction. It was developed by Mary Cover Jones in 1924 to help a boy named Peter who was terrified of rabbits. To help Peter get over his fear of rabbits, Jones planned to have a rabbit present while Peter ate his meals. She hoped that Peter's fear of the rabbit would go through extinction and that he would come to have the same positive feelings for the rabbit that he had toward food. Fearing that the rabbit's presence might frighten Peter enough to prevent him from eating, Jones came up with the idea of keeping the rabbit at a distance at first, then bringing it closer and closer. Not only did Peter lose his fear of rabbits, he developed positive feelings toward the rabbit and they became close friends.

Mary Cover Jones (1896–1987)

We will discuss systematic desensitization in more detail in Chapter 16 when we look at the various ways of treating phobias. Before we leave the topic, however, let's look at one more example of how this technique can be used to eliminate another common childhood fear. Many children are afraid of the dark. Sometimes their fear will keep them awake at night and may keep them from entering dimly lit rooms in their own homes. One way of helping children overcome their fear of the dark is to play a game called "treasure hunt." Prizes are hidden all over the house with the best prizes in the darkest rooms. The only rule is that the child can't turn on any lights. For children, the game is exciting and pleasant and they rapidly begin to associate the pleasant emotion of finding a prize with the stimulus of darkness. Notice that the child is in complete control of the frightening stimulus throughout the desensitization process. They are never forced to go into a dark room. During the game, being in the dark is completely up to them. This is a common feature of successful plans to help children with their fears. Any attempt to force a frightening stimulus on a child usually results in increasing rather than decreasing the fear. The same is true of adults.

Q: Do advertising agencies try to use classical conditioning to sell products?

In Chapter one, we said that J. B. Watson lost his job as a psychologist when he had an affair with his assistant, Rosalie Rayner. He worked at a number of jobs. For a while, he did market research by going from door to door asking people what kind of rubber boots they wore. Eventually, though, he got a job with an advertising agency. Over time, he became a successful advertising executive. Watson used the scientific principles of classical conditioning to sell products (Buckley, 1982, 1989). If fact, he is probably the father of modern advertising since these techniques are used constantly today, especially on television. Many television advertisements are designed to make us feel good. During the advertisement, we are shown the product or its logo and the advertisers hope that, if we see that advertisement often enough, we will come to associate that positive emotional state with the product itself. Seeing the product in the real world then makes us feel good. As a result, we want the product near us and often end up buying it.

The conditioned fear of body odor and sweat stains helps encourage people to buy antiperspirant products

Another approach used in advertising is to make us associate negative emotions with some condition. The advertisements then convince us that that the product will protect us from that condition. As a result of advertising, people are afraid of appearing in public with dandruff, unruly hair, wrinkled skin, the normal smell of our bodies, and stained teeth, among other things. We spend millions of dollars a year on products that do nothing more than protect us from these imaginary problems. Dandruff, for example, is a result of the normal shedding of skin on our heads. It is completely harmless and natural. Successful advertisements, however, have made most of us quite nervous at the thought of finding dandruff on our shoulders.

A New Look at Classical Conditioning—Rescorla, Kamin and Garcia

Q: Is classical conditioning simply a matter of putting two stimuli together?

Psychologists used to think so. For a long time, psychologists thought that classical conditioning was a mechanical process that occurred whenever we paired an unconditioned stimulus with a neutral stimulus. Pavlov (1927/1960) claimed that, "any natural phenomenon chosen at will may be converted into a conditioned stimulus." Learning researcher Gregory Kimble made a similar claim in 1956. Kimble wrote that "just about any activity of which the organism is capable can be conditioned and … these responses can be conditioned to any stimulus that the organism can perceive" (p. 195). Now, however, we know that classical conditioning is actually a much more complicated process. Kimble confessed in 1981 that scientific reports had proven him wrong. Sometimes, stimuli presented together don't result in classical conditioning. Sometimes, the stimuli don't need to be presented together for successful classical conditioning. Researchers have also found that some stimulus-response pairings work better than others do.

It is particularly easy to condition a human to fear a rat. Psychologist Martin Seligman (1971) noticed that irrational fears of rats, spiders and heights are fairly common. He also noted that very few people have irrational fears of stairs, electric outlets, or knives even though these things are much more likely to have been associated with a painful experience. You may remember the story from Chapter 4 of my younger brother, who, at the age of four, tried to plug a television antenna into an electric outlet and received a painful shock. He was quite frightened at the time but the experience does not seem to have had any lasting effect. As an adult, he is afraid of deep water but he is not at all afraid of electric outlets. He actually installed some of the outlets in his house. Seligman proposes that humans are biologically prepared to fear certain objects and situations that threatened our distant ancestors. This **biological preparedness** doesn't apply to relatively recent inventions like electric outlets, knives, and cars. These new stimuli are much more dangerous to us than rats or snakes but we are not biologically prepared to fear them. This is because fear of them did not help our ancestors survive.

biological preparedness
The biological readiness to learn some responses more easily than others

Q: How could a fear of rats be that important for survival?

In the fourteenth century, the black plague killed millions of people in Europe and Asia. In some areas, over half of the local population died from the plague (Bailey, 1996). As you probably know, the disease was transmitted to humans by fleas that were carried by rats (Platt, 1996). During the plague years, many children who had no fear of rats and did not try to avoid them probably did not live long enough to reproduce. Many of us are descended from people who lived in areas affected by the plague. It is likely that we have a number of ancestors whose fear of rats helped them avoid illness. This is a possible scientific explanation of why we are biologically prepared to fear rats.

We have made the case that biological preparedness exists in animals because it has survival value. This is a theory and although it seems sensible, as scientists, we need some way of testing it. In Chapter two, we argued that the best way to test a theory is to see if the theory's predictions are accurate. Can we make a scientific prediction that would let us test the theory that the classical conditioning process is custom designed to protect us from specific dangers in our environment? If biological preparedness has survival value, we could predict that different animals might be conditioned to different stimuli than others. We know for example, that birds are highly visual and depend on their vision to get food. Rats, on the other hand, have very weak eyes and depend more on taste and smell. As a result, we would predict that birds could be more easily conditioned to visual cues and rats to smell and taste cues. In support of this prediction, researchers compared classically conditioned responses in various species of birds and rats (Staddon & Ettinger, 1989; Wilcoxon et al., 1971). As predicted, they found that birds are much more easily conditioned using visual cues while rats are more easily conditioned to taste and smell stimuli.

Recent research has shown us that biological preparedness is not the only exception to the traditional view of classical conditioning. Important research by Robert Rescorla, Leon Kamin, and John Garcia has also changed our view of classical conditioning. In any science, new information is constantly being gathered. Traditional theories are often modified or rejected because of the new information.

Rescorla's Work Imagine that your telephone gives you a small shock when the weather is very dry. Imagine also that just before the shock occurs, you hear the sound of static on the phone. It wouldn't be long before you would flinch whenever you heard the static. What if the shock sometimes came without the static? Even if the static was always followed by the shock, you might not learn to flinch to the sound because the static wouldn't be a reliable indicator of the shock. In fact, you might not even notice the connection as long as the shock occurred often without the static.

In a series of important scientific experiments, **Robert Rescorla** and his associates discovered that presenting two stimuli together doesn't always lead to classical conditioning. If we shock a rat every time a tone sounds, the rat will quickly learn to fear the tone. What if we take a second rat and give it the same experience with one important difference? Suppose that for the second rat, we also present the shock some of the time without the tone? The second rat will receive the shock and the tone together the same number of times as the first rat. Will the second rat learn to fear the tone? No it will not (Rescorla, 1968, 1988).

Rescorla's work suggests that the rat will only learn to fear the tone if the tone is reliably connected with the shock. Since, for the second rat, the tone is often presented without the shock; it is not reliable as a predictor of the shock. It is tempting to say that because the shock occurs without the tone some of the time, the second rat "knows" that the two can't be related and doesn't learn to fear the tone. As scientists, though, we can't know what the rat is thinking so this explanation is unscientific.

Kamin's Experiments—Blocking In a related series of experiments, **Leon Kamin** (1968) discovered another exception to the traditional form of classical conditioning. First, Kamin conditioned rats to fear a tone by presenting it with a painful shock. Once the rats had learned to fear the tone, Kamin added a light that was presented along with the tone. Normally, presenting a light together with a shock would condition the rats to fear the light. In this case, though, the rats had already learned that the shock and the tone were related. They showed no conditioning to the new stimulus (the light). This effect is called **blocking.** The conditioned response to the tone *blocked* the effect of the light. It is as if the rats' brains had already established a warning signal for the shock and refused to respond to any new stimuli even though they were paired with the shock.

blocking
The ability of an existing conditioned response to block the learning of a new conditioned response

Garcia's Research on Taste Aversion Many years ago, I ate dinner at an Asian restaurant with some friends. I ordered Kung Pao Chicken, one of my favorite dishes. About eight hours after eating, I came down with the flu. I had chills, nausea, and a high fever. I was sick for several days. I know for sure that the chicken had nothing to do with my illness but, to this day, even the sight of Kung Pao Chicken will make me feel a little ill. Perhaps you have had a similar experience. In my case, a neutral stimulus (Kung Pao Chicken) was paired with an unconditioned stimulus (stomach flu)—the result: classical conditioning. Notice, however, that the two stimuli did not occur together in time. The illness came a number of hours after the meal.

In a series of classic experiments, **John Garcia** and his colleague, Robert Koelling (1966) demonstrated a similar effect in rats. They fed rats artificially sweetened water. Several hours later, they injected the rats with a drug that made them very sick. After the rats had recovered from their illness, Garcia and Koelling offered them the sweetened water. The rats were thirsty and readily drank unsweetened water but they refused to drink the flavored water. The rats had developed a **taste aversion** to the sweetened water as a result of classical conditioning. Two things about this experiment surprised people at the time. First, the conditioning occurred after a single pairing of the US and CS. More startling, the US was presented not seconds or minutes but *hours* after the CS. Psychologists were very skeptical of Garcia's results. They seemed to violate the principles of classical conditioning. Several leading journals in psychology refused to publish Garcia's research. They insisted that the results were impossible (Garcia et al., 1974). Garcia's findings have been replicated many times (once by me with Kung Pao Chicken). They are now widely accepted by psychologists. One experiment even showed that the conditioning could occur even when the drug causing the illness was presented 24 hours after the rats drank the flavored water (Etscorn & Stephens, 1973).

taste aversion
A learned tendency to avoid certain foods after eating them is followed by illness

Garcia conducted another famous experiment that supports the concept of biological preparedness (Garcia & Koelling, 1966) Garcia presented rats with three stimuli together—a clicking noise, a bright light, and flavored water. They divided the rats into two groups. One group received an electrical shock to the feet. The other group received lithium chloride or X-rays—stimuli that that made them sick. If any stimulus can be connected to any response as Pavlov suggested, both groups should have learned to fear and avoid all three stimuli. Instead, the researchers saw a very different result. The rats that became sick avoided the flavored water but showed no reaction to the light or the clicking noise. The rats that were shocked showed a reaction to both the clicking noise and the light. Clearly, classical conditioning is not just a mechanical process depending on the pairing of stimuli. Avoiding the taste of foods that made you sick has survival value. It also makes sense that a light or a clicking noise might be associated with a painful stimulus.

Does This Mean That the Animal Is Thinking? How can we explain the findings of Rescorla, Kamin, and Garcia? The process of classical conditioning exists in animals for a reason. If it didn't help our ancestors survive long enough to reproduce, it wouldn't exist. It comes as no surprise to most people that we are more easily conditioned to fear rats, snakes, and spiders, than pillows, spoons, and beach balls. It is easy to see how fear of painful or dangerous stimuli or the desire to avoid food that causes illness would help an animal survive. Think how much more useful it would be, however, if the process were tuned to help us avoid the true cause of the pain or illness rather than whatever stimuli happened to be present. The research of Rescorla, Kamin, and Garcia shows that nature has provided us with a sophisticated system for learning which stimuli we should respond to. Conditioning works best when the stimulus and the response are related in some meaningful way. Simply presenting them together is not always enough. Robert Rescorla (1988) points out a number of problems with the simple traditional view of classical conditioning. The stimuli need not be presented together and presenting stimuli together doesn't always result in conditioning. The response to the conditioned stimulus is seldom the same as the response to the unconditioned stimulus. Clearly, the process is much more complicated than Pavlov and Watson thought. As scientists, though, we should be careful about assuming that the animal is "thinking" about the conditioning situation and "deciding" to become conditioned. In order to draw such a conclusion, we would need some clear empirical evidence of this "thinking." The fact that the process is complicated doesn't necessarily mean that we have to give up the idea that classical conditioning is a mechanical process that is built in by nature.

Locking It In

1. When an animal gets sick after eating a certain food and avoids that food in the future, it's called ________ ________.
2. Leon Kamin is known for his research on ________.
3. Martin Seligman believes that fears of heights, rats, and spiders are common because of ________ ________.
4. Critical Thinking: How did conditioned taste aversions help our ancestors survive?

Answers: 1) taste aversion, 2) blocking, 3) biological preparedness

Applications PROBLEM SOLVING WITH CLASSICAL CONDITIONING

Deterring the Coyotes

In the 1970's, sheep ranchers in the western United States had a serious problem with coyotes killing their sheep. They tried to solve their problem by killing the coyotes. Gustavson and other researchers (1974) found a scientific way to save both the sheep and the coyotes. They found a solution that used taste aversion to make the coyotes lose interest in eating the sheep. They set out lamb carcasses that had been injected with chemicals that made the coyotes very ill but did not kill them. After eating the tainted lamb once or twice, the coyotes completely lost their taste for lamb and left the flocks alone. The same scientific technique was later used by park rangers to keep coyotes from scavenging at campsites in a national park (Cornell & Cornley, 1979). The rangers put out a variety of camping foods full of lithium chloride, the same chemical that had been used on the lamb carcasses. Once the coyotes grew sick a few times from eating the camp food, they stopped visiting the campground. Even two years later, they were still staying away.

Reducing Nausea From Chemotherapy

One of the situations that can lead to taste aversion in humans is chemotherapy. Because chemotherapy makes people feel ill, cancer patients who receive chemotherapy can develop aversions to foods they have eaten in the several hours before their treatment (Bovbjerg et al, 1992). Some cancer patients find that they have come to dislike some of their favorite foods if they eat them before their chemotherapy. Ilene Bernstein (1988) found a scientific way to help fight this problem. She found that if the patients ate some unusual or novel food before their chemotherapy, the taste aversion would focus on that particular food. When this was done, the patients tended not to develop an aversion to other, more common foods they had eaten. As a result of this research, patients are advised not to eat preferred or nutritious foods before chemotherapy. Instead, they should eat some unusual food that is not a regular part of their diet. Notice that this is another demonstration of the biological significance of classical conditioning. When we experience an upset stomach, avoiding an unusual or new food we have eaten recently is a better survival strategy than giving up something that we eat regularly.

Operant Conditioning

Q: Can classical conditioning be used on any behavior?

Suppose that you want to use classical conditioning to teach your dog to sit. First, you need an unconditioned stimulus that will make the dog sit. You might try pushing down on the back end of the dog. The trouble with this is that the dog isn't really sitting. Many dogs will respond to this stimulus by pushing back, which is the opposite of sitting. There are many other stimuli you might try such as backing the dog into a corner, holding food over the dog's head, or pulling back on the dog's lead. Remember though that an unconditioned stimulus should make the response occur every time without fail. None of these stimuli will really do the job. The reason that it is so hard to come up with an unconditioned stimulus for sitting is that there isn't one. Classical conditioning works well with involuntary behavior like reflexes and emotional behaviors. Sitting isn't a reflex, though, and it's certainly not an emotional behavior. This means that we can forget using classical conditioning to teach a dog to sit. We must resort to another kind of conditioning called *operant conditioning.* Operant behaviors "operate" on the environment. Unlike classically conditioned behaviors, they are not elicited by specific stimuli.

A friend of mine was visiting some neighbors who were completely deaf. During the visit, their five-year-old (who could hear fine) was crawling around under a table. The child forgot that he was under a table and stood up. My friend heard the loud thump of the child's head hitting the table and waited for the sound of crying. To his surprise, the child didn't make a sound. He crawled out from under the table and moved to where his parents could see him. Once his parents were looking at him he screwed up his face and began holding his head in pain. He looked for all the world like a child sobbing hysterically but he didn't make a sound. Apparently, he had learned that the sound part wasn't important. The basic principle of operant conditioning is that some behaviors can be modified by their consequences. Behaviors that have positive consequences tend to occur more often. Those with negative consequences tend to occur less often. The child in the story had learned that looking hurt got him attention and comfort from his parents. He didn't cry out loud because, for him, the sound of crying had no consequences. Operant behaviors develop and change because of what happens *after* they occur. They depend on their consequences. In contrast, classical conditioning involves putting together two stimuli *before* the behavior occurs. To use operant conditioning to teach your dog to sit, you might praise the dog every time it sits on command. Rather than pairing two stimuli as we did in classical conditioning, you are making sure that there are positive consequences for sitting on command. The praise comes *after* the response. We'll discuss the details of operant conditioning in more detail but first, let's look at its history.

Thorndike's Puzzle Box—The Law of Effect

In 1898, **E. L. Thorndike** performed a series of classic experiments with cats. These experiments laid the groundwork for later research in operant conditioning. Thorndike put the cats in a "puzzle box" and placed some food just outside the box. In order to escape from the box and get the food, the cats had to perform some simple act such as pressing a lever or pulling a string. When a cat was first placed in the box, it tried to escape by clawing or biting various parts of the cage in a random fashion. Eventually, the cat would stumble onto the correct response to release the door. The next time the cat was placed in the box, it might escape a little sooner. After a number of trials, the cat would make the response as soon as it was placed in the cage. Thorndike made a graph (called a **learning curve**) of the time it took each cat to escape from the box. He noticed that the time dropped gradually and that often the cat took longer to escape than it had on the previous trial.

Q: Did Thorndike believe that the cats "figured out" how to escape from the box?

No, Thorndike was a behaviorist. He argued that if the cats "understood" how to get out of the box, the time it took them to escape would be high at first and then drop suddenly when they figured out how to escape. The typical learning curve doesn't show this kind of pattern at all. Thorndike believed that when the cat made the correct response, it was rewarded with escape and food. This reward, according to Thorndike, simply made the correct response a little more likely. In other words, it raised the probability of the correct response. Over time, the consequences of the response made it more and more likely until, finally, it occurred very rapidly. Thorndike believed that, over time, positive consequences "stamped in" the correct response and negative consequences "stamped out" all the possible incorrect responses. He called this the **law of effect.** Thorndike didn't believe that the cats needed to "think" or "understand" to escape from the box. Their correct responses just became more likely over time as a result of the law of effect. Thorndike called this instrumental conditioning because that cat's correct responses were "instrumental" in getting them out of the box. B. F. Skinner later called this same kind of learning **operant conditioning** (Skinner, 1938). The two terms mean the same thing but most modern psychologists use the term operant conditioning.

law of effect
Thorndike's term for the principle that responses that are followed by desirable consequences occur more often and responses that are not followed by desirable consequences occur less often

operant conditioning
A process of conditioning in which behavior change is caused by the consequences of the behavior

B. F. Skinner

Noted behaviorist **B. F. Skinner** spent most of his life studying learning. Skinner wanted to find scientific principles that would explain the behavior of humans and other animals (Skinner, 1956, 1967). Skinner agreed with Watson and Pavlov that some behaviors were elicited in response to certain stimuli. He also agreed that these behaviors could be conditioned using classical conditioning. Skinner disagreed, however with idea that this kind of conditioning explained all human learning. Skinner began studying the large number of behaviors that animals emitted on their own like the ones Thorndike's cats used to escape from the puzzle box. Skinner did not call these behaviors "voluntary" because he believed that the behaviors were a response to the animal's environment. He did not believe that the behaviors were intentional or the result of a conscious choice on the part of the animal. In looking for scientific principles that could explain how these behaviors were learned, Skinner developed his model of *operant conditioning.* In operant conditioning, behavior develops and changes because of its consequences. By consequences, Skinner meant reinforcement and punishment. A reinforcer is something like a reward. We'll discuss the difference between a reinforcer and a reward and define both reinforcement and punishment later in this chapter. Although Skinner's theory was based on both reinforcement and punishment, his research focused almost entirely on reinforcement. Skinner considered punishment a relatively ineffective technique with many negative side effects. He believed that in practical applications of operant conditioning, punishment should be avoided whenever possible. (Skinner, 1953).

B. F. Skinner (1904–1990), noted behaviorist and developer of operant conditioning

Skinner developed an experimental chamber to study operant conditioning in rats (see Figure 7.4). We now call this kind of experimental chamber a "Skinner box" in his honor. Later, he developed a similar chamber to study operant conditioning in pigeons. In the original Skinner box, every time the rat pressed down on a brass lever, a food pellet dropped down a small tube into a food cup. Skinner soon found that reinforcing the rat with food for pressing the bar increased the rate of bar pressing. As a firm behaviorist, Skinner avoided making any statements about the rats' mental states. He based his explanation of the behavior completely on observable stimuli and responses. If you had asked Skinner "is the rat pressing the bar because it wants food?" Skinner might have replied, "We can't know what the rat wants. It is pressing the bar because it been reinforced for pressing the bar." Notice that both the bar pressing and the reinforcement are easily observed.

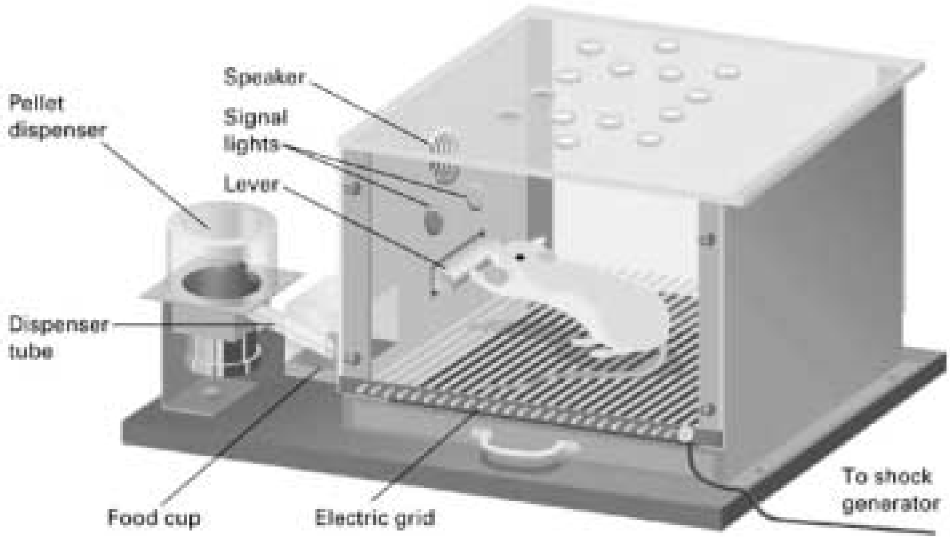

FIGURE 7.4

A TYPICAL SKINNER BOX

Although this drawing shows a shock grid on the floor of the box, B. F. Skinner never shocked any rats (MacCorquodale, 1971)

Skinner and his associates went on to develop many practical applications of operant conditioning. Skinner taught pigeons to play ping-pong, to remove defective capsules from the assembly line in a drug factory, and to control guided missiles. He created programmed instruction in schools that allowed students to work at their own pace and receive immediate reinforcement for their correct answers. Later, operant conditioning was used to create incentive programs in businesses, systems of reinforcers for mental patients in hospitals, and deposit legislation that reinforces people for returning bottles and cans.

Because of Skinner's fame and his theories about controlling behavior, there are many misconceptions and false rumors about his life and work. Many critics have suggested that Skinner proposed the cold-hearted manipulation of human behavior. To the contrary, Skinner argued that human behavior is already controlled by its consequences. He believed that it would be unethical not to try to arrange those consequences so that people could lead more satisfying lives. He also pointed out that we could help solve many of the problems we face such as poverty, crime, illiteracy, and environmental destruction by the intelligent application of operant principles.

B. F. Skinner's daughter Deborah in the air crib attended by her mom

One of the most misleading myths about Skinner's life comes from his invention of a very high-tech crib for his daughter, Deborah. The Air-Crib, as he called it, was nothing more than an extremely comfortable and convenient crib and playpen. It had temperature and humidity control and was almost completely dust-proof. According to the rumors, Skinner's daughter was put in a large Skinner box, forced to press a lever to get food, never taken out of the box, and later became mentally ill and/or committed suicide. None of these rumors are even slightly true. Skinner used the Air-Crib with both his daughters, Deborah and Julie. According to all reports, the children were very comfortable in the crib and almost never cried. Both daughters were quite normal and Julie later used the Air-Crib with both her daughters.

Reinforcement and Punishment

Q: Is a Reinforcer the Same as a Reward?

Rewards are often reinforcers but there is an important difference in how they are defined. Reinforcers are defined strictly according to their effect on behavior. A reinforcer (sometimes called a reinforcing stimulus) is defined as *anything that increases the frequency of the behavior it follows.* If a particular behavior is followed by a stimulus every time it occurs and if, over time, the stimulus occurs

reinforcer
A stimulus that increases the frequency of the behavior it follows

more and more often, the stimulus is a reinforcer. Suppose that every time you say the word "toast" I give you a dollar. It is likely that you would develop a rapid interest in toast as a topic of conversation. You would probably try to work the word into every sentence. You might even end up just repeating the word "toast" over and over while I handed you dollars. Giving you the dollar has increased the frequency of the behavior it follows. Therefore, it's a reinforcer. Now suppose that every time you say the word "banana" I rap your knuckles with a ruler. It seems unlikely, but if you start saying the word "banana" more and more often, we would have to conclude that hitting you on the knuckles is a reinforcer. Few people would consider it a reward. Rewards are usually defined by the person who gives them. A teacher might say "Lucinda, you got a perfect score on the spelling test; as a reward, I'm going to let you wear this gold star on your forehead for the rest of the day." The star is a reward because the teacher says it is. Is it a reinforcer? If Lucinda starts getting more perfect scores on spelling tests, it's a reinforcer. If Lucinda never gets another perfect test score, however, it's not a reinforcer.

punishment
A stimulus that decreases the frequency of the behavior it follows

We can change our definition of reward to the definition of punishment by changing a single word. Punishment is defined as *anything that decreases the frequency of the behavior it follows*. Notice that punishment, like reinforcement, is defined strictly in terms of its effect on behavior. Parents sometimes say, "I don't know what's the matter with my kids; the more I punish them, the worse they behave." Knowing the scientific definition of punishment, we can be sure that the parent is actually using reinforcement, not punishment. Suppose that a teacher decides to punish a child for standing up during class. Every time the child gets up, the teacher yells, "Get back in your seat." The teacher hopes that this punishment will cut down the frequency of the child getting up. What if the result of this "punishment" is that the child gets up more and more often? We can see that yelling at this child is not a punishment but a reinforcer. The teacher might say, "I don't know what's the matter with this kid; the more I yell, the worse things get. This child must like punishment." What the child is getting, of course, is attention. Attention is a powerful reinforcer for both children and adults. A better question would be, "Teacher, you've just discovered an effective reinforcer. Why are you using it to reinforce a behavior you don't like?" It is a good idea to be careful what behaviors you pay attention to. Often, paying attention to a behavior will make it happen more frequently. Reinforcement and punishment can be powerful scientific tools in changing behavior.

Comparing Classical and Operant Conditioning We will discuss the principles and techniques of operant conditioning in more detail later in the chapter. First, though, let's look more closely at the differences between classical and operant conditioning. Classical and operant conditioning share many terms such as stimulus, response, extinction, generalization, discrimination, etc. They differ, however, in the kinds of behaviors they work on. They also differ in the conditioning procedure itself (see table 8.1 earlier in the chapter). Classical conditioning works well with reflexes and emotional behaviors. Operant conditioning is used with most other behaviors. If you want to condition an animal to blink, choke, salivate, or be afraid, classical conditioning is the right method. To condition other behaviors that are not reflexes or emotional behaviors, however, we must use operant conditioning. The conditioning procedure is also very different. In classical conditioning, we pair two stimuli, the US and CS. Nature takes care of the rest. In operant conditioning, what happens after the response is most important. In operant conditioning, we follow the response with reinforcement or punishment to make it occur more or less often.

Locking It In

1. Anything that increases the frequency of the behavior it follows is ________________.
2. Anything that decreases the frequency of the behavior if follows is ________________.
3. ____________ put cats in puzzle boxes.
4. ____________ reinforced rats with food for pressing a bar.
5. In ________________ conditioning, we put two stimuli together.
6. In ________________ conditioning, we follow a response with reinforcement or punishment.
7. Critical Thinking: If we use the scientific definitions of reinforcement and punishment, can there be such a thing as a reinforcement or punishment that "doesn't work" to change behavior?

Answers: 1) reinforcement, 2) punishment, 3) E. L. Thorndike, 4) B. F. Skinner, 5) classical, 6) operant

Negative Reinforcement—It's Not Punishment You may have noticed that a pleasant stimulus like food can be used either for punishment or reinforcement. Presenting food is reinforcing; taking it away can be punishing. Similarly, presenting a painful electric shock can be punishing; taking it away can be reinforcing. This means that we really have four kinds of reinforcement and punishment (see Table 7.2)—positive and negative reinforcement, and positive and negative punishment. These are technical terms and students often have trouble keeping them straight. The key is to remember that to a behavioral scientist, "positive" means that we are presenting something and "negative" means that we are taking something away. This is not the way we use these terms in everyday life. It also helps to remember that reinforcement always increases a behavior and punishment always decreases it.

Positive reinforcement means increasing a behavior by presenting something (something pleasant). **Negative reinforcement** means increasing a behavior by taking something away (something unpleasant). If we give food as a reinforcer, we are using positive reinforcement. We can expect the behavior to increase in frequency. Negative reinforcement means increasing a behavior by taking away something unpleasant. If we turn off a painful electric shock every time a rat presses a bar, we are using negative reinforcement to increase the

positive reinforcement Increasing the frequency of a response by presenting a reinforcing stimulus after the response occurs

negative reinforcement Increasing the frequency of a response by removing an aversive stimulus after the response occurs

	STIMULUS IS PRESENTED	STIMULUS IS REMOVED
STIMULUS IS REINFORCING	Positive reinforcement Response increases	Negative Punishment Response decreases
STIMULUS IS AVERSIVE	Positive punishment Response decreases	Negative reinforcement Response increases.

TABLE 7.2

TYPES OF REINFORCEMENT AND PUNISHMENT

Remember that reinforcement always increases the frequency of the response it follows. Punishment always decreases the frequency of the response it follows. Remember also that if we present a reinforcing or punishing stimulus, we call it *positive* reinforcement or punishment. If we take something away, we call it *negative* reinforcement or punishment rate of bar pressing. Psychology professors are often frustrated by the fact that students always seem to think negative reinforcement is the same as punishment. If you think about this for a moment, you can see that it is impossible. Reinforcement increases the probability of a behavior. Punishment decreases the probability of a behavior. That means that it is impossible for reinforcement (positive or negative) to be punishment.

Q: Is there such a thing as positive and negative punishment?

positive punishment
Decreasing the frequency of a response by presenting an aversive stimulus after the response occurs

negative punishment
Decreasing the frequency of a response by removing a reinforcing stimulus after the response occurs

As with reinforcement, punishment can also be positive or negative. In positive punishment, we present something. In negative punishment, we remove something. A parent might stop a child from using objectionable language by washing out the child's mouth with soap whenever they say one of the forbidden words. In doing this, they are **positive punishment** by presenting something unpleasant. In positive punishment, an *aversive* stimulus follows the behavior. An aversive stimulus is any stimulus that the animal will work to avoid such as electric shock or a rap on the knuckles. Instead of *positive punishment*, the parent might charge the child a quarter every time the child says an objectionable word. Now they are using **negative punishment** by taking away something pleasant. Both can be effective. To get these terms correct, just remember the following two steps. First, "positive" means something is presented and "negative" means something is taken away. Second, reinforcement is meant to increase a behavior, punishment is meant decrease a behavior. Let's look at one more example. Suppose that a parent gets a child to clean his or her room by nagging. The nagging continues until the behavior occurs, and then it stops. When the behavior occurs, is something presented or taken away? Because the nagging stops, something is being taken away when the behavior occurs. That means whatever is being used is "negative." Is this procedure supposed to increase of decrease the behavior? Because the procedure is designed to increase the rate of room cleaning, it is reinforcement. Therefore, the parent is using negative reinforcement (not punishment!).

Locking It In

1. Negative reinforcement is the same as punishment. (T/F)
2. Presenting food as a reinforcer would be an example of _________ ________.
3. A child stops whining when given a piece of candy. The child is using _________ _________.
4. If you give your dog a treat when it sits on command, you are using _________ _________.
5. If you fine a friend a dollar whenever she interrupts you, you are using _________ _________.
6. If you spank a child for lying, you are using _________ _________.
7. Critical Thinking: Why is it seldom effective when parents use punishment to stop their children from lying to them?

Answers: 1) F 2) positive reinforcement, 3) negative reinforcement, 4) positive reinforcement, 5) negative punishment, 6) positive punishment

Shaping Behavior

Q: How can you use reinforcement on a behavior that never occurs?

Suppose that you'd like a child to do a good job of cleaning his or her room. You'd be more than happy to reinforce the behavior but, sadly, it never occurs.

How can you use operant conditioning to condition this behavior? What if the behavior you want to condition does happen, but it happens so rarely that the training might take a very long time? The answer to both of these questions lies in a scientific technique called **shaping** (Skinner, 1937). Shaping is defined as *reinforcing successive approximations to the desired behavior.* In shaping, we reinforce any behavior that is a little more like the behavior we want to see. Let's look at a simple example. When I let my dog Penny into the house, she would come racing in at full speed, running over my feet, and bumping into the furniture. Yelling at her only reinforced this behavior with attention and made it worse. I couldn't reinforce her for coming through the door slowly because it never happened. I finally realized that shaping was the best way to change this behavior. I began saying "good dog" whenever she came in just a little slower than usual. Over the next few weeks, she began coming through the door slower and slower. After three weeks, she was coming through the door in slow motion. I had reinforced successive approximations to entering the house slowly.

shaping
reinforcing successive approximations to the desired behavior

Shaping is a powerful and underused scientific technique. It can be used to change both desirable and undesirable behaviors in others and in ourselves. It has several important advantages. As we mentioned above, it allows us to reinforce behaviors that otherwise would never happen. It also works much faster than simply reinforcing the desired behavior. Shaping also can be a powerful alternative to punishment.

Animals can be taught some amazing behaviors using operant conditioning techniques

To make shaping effective, it is important to remember two practical facts. First, the steps must be small. Shaping works by breaking up the progress toward the desired behavior into small steps. If the steps are too big, the shaping will not be successful. Say, for example, that you would like to exercise more. Currently you don't exercise at all. You'd like to exercise for at least 20 minutes per day. You could start by giving yourself a reinforcer every time you exercised for 20 minutes or more. It is unlikely, however, that you will be successful in going from no exercise to 20 minutes in one jump. You are likely to be more successful if at first you reinforce yourself for 1 minute of exercise. Once you are exercising regularly, increase this to 2 minutes, then 3 minutes, and so on. The best way to tell if the shaping steps are small enough is to see if reinforcers are being delivered fairly often. If you are not delivering reinforcers regularly, the steps are probably too big. If you haven't delivered a reinforcer for significant amount of time, and the person or animal you are training is showing signs of frustration or if the behavior is not occurring as often as it was before, you should probably make the steps smaller.

The second thing to remember with shaping is that at first, you will be reinforcing relatively poor performance. Sometimes this is hard to do, especially for parents, bosses, and people trying to change their own behavior. If you want your children to learn to do a good job of cleaning their room, you will have more success if you first reinforce any attempt to clean the room at all. In fact, the first step in shaping room cleaning should probably be just going into the room. Gradually, you can make the reinforcement depend on better and better performance.

Q: Doesn't this teach the child that doing a poor job is acceptable?

Acceptable performance almost never comes right away. If you are taking tennis lessons, a good instructor won't criticize your poor early performance. Instead, he or she will reinforce you for the little improvements you make on the way to becoming a good tennis player. The instructor is not teaching you that poor performance is acceptable. Instead, he or she is reinforcing you for improving your performance. The fact that the final behavior has been described in detail won't usually result in perfect performance right away. Explaining to a child how we want the room cleaned probably won't mean that the child will do a perfect job the first time. Reinforcing successive approximations to a good job, however, will lead to good performance in the long run and make the learning experience more pleasant for both you and the child. In life, there are many behaviors that we may want to change. Our own behavior and the behavior of our pets and the people we deal with on a daily basis are seldom perfect. Shaping will often change the behavior faster and with fewer unpleasant side effects than any other technique.

Operant Extinction

Q: How can we use operant conditioning to get rid of unwanted behaviors?

In classical extinction, we stop presenting the unconditioned stimulus. In operant extinction, we no longer present the reinforcer. Without the reinforcer to prop it up, the behavior gradually grows less and less frequent. Eventually, it will disappear. Just as in classical conditioning, we can expect spontaneous recovery. The response will still reappear after a rest period. In our traditional example, the rat presses a bar to get food. To extinguish this behavior, we need to make sure that we no longer deliver food after a bar press. Over time, the rat will press the bar less and less and will eventually quit. Because of spontaneous recovery, the rat will still press the bar a few times after a rest period.

In operant extinction, two other things can be expected that don't occur in classical extinction. First, the rate of the behavior often goes up before it goes down. When we first stop the food delivery, the rat will actually increase its rate of bar pressing. The rate will go down but we can count on it going up first. The second thing we can expect in operant extinction is emotional behavior. Rats will shake, defecate, and sometimes bite the bar. This is the rat's version of a temper tantrum. It must be expected in operant extinction. This behavior may seem silly but humans behave exactly the same way.

Imagine that every day you put money in a candy machine, press a button, and receive a candy reinforcer. One day, when you press the button, no candy comes out. The machine has put you on extinction. What's the first thing you do? If you said, "kick the machine," you haven't been paying attention. Usually, the first thing people do in this situation is to press the button several times. Notice that your usual rate for this behavior is once a day. In extinction, the rate of response goes up at first. If the machine continues to keep you on this extinction schedule, your response will eventually stop altogether. The *second* thing you will probably do is to kick the machine, curse, or mutter something insulting to the machine. This is emotional behavior and it is also a standard response to extinction in operant conditioning. You may be thinking that pushing the button is just a reasonable way of checking to make sure that the machine isn't working. If so, imagine this example. Have you ever gone up to the door of a store just after closing and found it locked? Even after seeing the closed sign and the bolt locking the door, have you ever tried opening it? I have. In fact, I sometimes pull several times on the door. Then I grumble some and leave. This is the standard pattern of operant extinction—first the rate goes up, then the emotional behavior appears. If you plan to use operant extinction to eliminate some undesirable response, be sure that you are prepared to hold out through these two effects. People sometimes give up too soon. When the response goes up, they think that what they are doing is making things worse. They give in and deliver the reinforcer. When this happens, they end up actually reinforcing the behavior they are trying to eliminate. Other times, they may give up too soon because they are surprised by the emotional behavior. This reinforces the emotional behavior and can teach your subject to throw tantrums to get what he or she wants.

Q: What if I want to eliminate a behavior but don't know what the reinforcer is?

Since scientific research tells us that reinforcement works best when the reinforcer comes immediately after the response, the place to look for the reinforcer is right after the behavior. If your child interrupts you while you are on the phone, whatever you do in response to the interruption is probably a reinforcer. If you get into lots of arguments with a particular friend, whatever you do when the friend disagrees with you is probably a reinforcer.

Locking It In

1. Reinforcing successive approximations to some desired behavior is called ________________.
2. When a behavior is no longer followed by reinforcement, it occurs less often. This is called __________________.
3. Presenting a stimulus that increases the probability of the behavior it follows is called ___________________.
4. Presenting a stimulus that decreases the probability of the behavior it follows is called ___________________.
5. Critical Thinking: Why does shaping work faster than simple positive reinforcement?

Answers: 1) shaping, 2) extinction, 3) reinforcement, 4) punishment

Kinds of Reinforcers

Q: What kinds of reinforcers are there?

When Thorndike first put cats in his puzzle box, he put food outside the box as a reinforcer. He soon learned, however, that the food was not necessary. The cats would learn the response just to get out of the box. Apparently, escaping from the box was also a reinforcer. Reinforcers that don't need to be learned such as food, water, escape from pain, and the chance to engage in sex, are called **primary reinforcers.** These reinforcers are obviously important for the survival of the species and are built in by nature. Primary reinforcers are often limited by the condition of the animal. Food, for example, will not work well as a reinforcer unless the animal is hungry. Psychologists agree that the biological necessities listed above are all primary reinforcers. There may be other primary reinforcers as well. Scientific researchers in the 1950s identified a group of reinforcers they called **sensory reinforcers.** These reinforcers involve the animals' senses and, like primary reinforcers, are probably unlearned. Monkeys, for example, can be reinforced with the opportunity to look out a window (Butler, 1954). Humans will press a button in order to see a display of lights in a dark room (Jones, Wilkinson, & Braden, 1961). The popularity of many computer games probably depends on sensory reinforcers. This is especially true of games that involve a lot of exploration. Humans and animals can be reinforced by the opportunity to explore their environment, physical contact, the opportunity to hold something in their hands, and a number of other sensory reinforcers. These sensory reinforcers have obvious survival value for the species. The desire to explore and understand your environment and to have contact with your own species are valuable traits for any animal.

primary reinforcer An innate (unlearned) reinforcer such as food or water

sensory reinforcer a reinforcer that has the stimulation of the senses as its only reinforcing property

A friend of mine once served as a playground monitor for an elementary school. He developed a plan to get the children to stop pushing and shoving on the playground. He drew a chalk square in one corner of the playground and declared it the "jail." He announced to the children that anyone pushing or shoving on the playground would have to spend fifteen minutes in jail. Within ten minutes, over half the children were in jail. The ones still outside were madly pushing and shoving each other and begging to be arrested. My friend was a clever enough scientist to realize that he had discovered a powerful *social reinforcer.* He kicked everyone out of the jail and announced that the children who behaved the best and who did not push or shove would get to go to jail. For the rest of the recess, the children behaved perfectly. Humans and some animals can be reinforced with praise, attention, and the chance to spend time with members of their own species.

Social reinforcers can be very powerful. Humans are very social animals and spend much of their time with others. Social reinforcement probably explains why many people are willing to spend more for a drink in a bar than the same drink would cost them at home. It also explains why the same people may leave the bar if there is no one there. One of the more severe punishments in prisons around the world is solitary confinement. Some psychologists consider social and sensory reinforcers to be learned. It seems likely, however, that they are built in by nature. They have survival value for the species for one thing. For another, it is easy to think of cases where these reinforcers work without any apparent learning on the part of the animal. Praise, for example, is a powerful reinforcer for many dogs. Many dog owners do not bother to pair praise with food because it already works for their dogs. Dogs, like humans, are very social animals and responding to praise and attention are probably part of their biological programming. For Thorndike's cats, escaping from the puzzle box was an effective reinforcer. For cats (and for other animals as well) escaping from a confining environment may be a primary reinforcer.

Q: Are there reinforcers that are not built in by nature?

secondary reinforcer
A learned reinforcer that gets its power by being paired with a primary reinforcer

Some reinforcers are definitely learned. These reinforcers are called **secondary reinforcers.** Secondary reinforcers have no value of their own. They get their power by being associated with primary reinforcers. If a tone sounds every time a rat gets a food pellet, the tone will soon become a secondary reinforcer. Secondary reinforcers will lose their effectiveness if they are not paired occasionally with a primary reinforcer. Secondary reinforcers tend to be limited in effectiveness in the same way that primary reinforcers are. We said earlier that we can't use food to reinforce a rat that isn't hungry. Similarly, a tone that has been associated with food will not be effective if the rat is full. Some learned reinforcers, however, are good almost any time. This is because they have been paired with a number of different primary reinforcers. We call these all-purpose reinforcers **generalized reinforcers.** Money, in many societies, is a powerful generalized reinforcer. In these societies, money can be an effective reinforcer even if the person is not hungry or thirsty. If a person uses money to buy food when they are hungry and something to drink when they are thirsty, money will become a generalized reinforcer. In some institutions, generalized reinforcers called *tokens* are created for use in changing behavior. Schools, mental hospitals, and group homes often create a *token economy* using artificial reinforcers such as poker chips or points to reinforce desirable behaviors. In a mental hospital, for example, patients can use tokens to buy things like dessert, field trips, or access to the television. The tokens become generalized reinforcers. This allows the staff to use them to reinforce behaviors like tooth brushing and room cleaning. They can also use the tokens as part of a shaping program to reinforce the patients for improving their behavior (Kazdin, 1977; Pitts, 1976).

generalized reinforcer
A learned reinforcer that is effective at any time because it has been paired with several different primary reinforcers

Psychologist **David Premack** discovered another kind of reinforcer in 1962. Premack (pronounced "PREE-mack") noticed that some behaviors occur more often than others do. Being a good scientist, he wondered if the opportunity to engage in a high-frequency behavior could be used to reinforce a low-frequency behavior. His experiments showed that this was the case. We now call this effect the **Premack principle** in his honor. Consider a child who spends a lot of time playing video games and very little time reading. According to the Premack principle, we could increase this child's reading by using video game playing as a reinforcer. On the other hand, with a child who reads a lot and seldom plays video games, we could do the opposite. We could increase the rate of game playing by reinforcing game playing with the opportunity to read.

Premack principle
David Premack's idea that the opportunity to engage in a high-frequency behavior can be used to reinforce a low-frequency behavior

One problem with the Premack principle as a scientific explanation of behavior is that sometimes low-frequency events can be very reinforcing. The opportunity to watch the Winter Olympics, for example, happens only once every four years. This makes it a very low-frequency behavior. In spite of this, people often give up high-frequency behaviors to watch the Olympics. Another way of looking at the relative power of various reinforcers is called the *disequilibrium principle* (Timberlake & Farmer-Dougan, 1991). According to the disequilibrium principle, stimuli can be reinforcing some times and not others. Whether a stimulus is reinforcing or not depends on the state of the animal. The most obvious example is food. If a person is hungry, food is a very powerful reinforcer. If a person is full, however, food may not be an effective reinforcer. The disequilibrium principle suggests that all other reinforcing stimuli work this way as well. If you have been alone for a long time, the opportunity to see other people will be a powerful reinforcer. If you have been on a crowded bus for many hours, however, the opportunity to be by yourself may be reinforcing. The disequilibrium principle suggests that each person has a preferred range for every stimulus. When the frequency of the stimulus gets outside of the preferred range, the person experiences disequilibrium. When this happens, the stimulus becomes an effective reinforcer. The further outside the preferred range the reinforcer is, the more effective it is as a reinforcer. One person, for example, might prefer getting a haircut about once a week. Another person might prefer getting a haircut once a month. The first person might experience haircut disequilibrium after two weeks without a haircut. For this person, two weeks after the last haircut, the opportunity to get a haircut would be an effective reinforcer. The second person would not experience disequilibrium until much later. For the second person, the opportunity to get a haircut two weeks after the last one might not be reinforcing at all.

Locking It In

1. Reinforcers such as food and water that don't need to be learned are called _________ reinforcers.
2. ___________ reinforcers get their power by being presented with primary reinforcers.
3. ___________ are good at any time because they have been paired with several primary reinforcers.
4. Schools and mental hospitals sometimes use artificial generalized reinforcers called __________.
5. When we reinforce a low-frequency behavior with the opportunity to engage in a high-frequency behavior, we are using the ____________ principle.
6. Critical Thinking: Why is money an ineffective reinforcer with young children?

Answers: 1) primary, 2) secondary, 3) generalized, 4) tokens, 5) Premack

superstitious behavior Behavior that is reinforced by accidental (non-contingent) reinforcement

Superstitious Behavior

Q: What about accidental reinforcers, are they effective?

My family sometimes gets together on Mother's Day at a local racetrack. One Mother's Day at the track, my wife and son both bet on a particular horse, a long shot, to win. They had never watched a race from the finish line and decided that this would be a good time to try it. They went down to the track and stood at the rail next to the finish line for the race. Their horse won. Both the race and the payoff were very exciting for them. That was a number of years ago but, to this day, they stand at the finish line for every race they bet on. How would a scientist explain this behavior? We all know that standing at the finish line didn't actually make their horse run any faster but they were reinforced for it. As scientists, we know that behaviors that are reinforced tend to happen more often. This happens even if the behavior doesn't lead directly to the reinforcer. Behaviors that develop or persist because of accidental reinforcement are called superstitious behaviors. The scientific term for these accidental reinforcers is *non-contingent reinforcement.* Non-contingent reinforcers are not directly related to the behavior that comes before them, they just happen to follow it. Suppose we put a pigeon in a cage and give it access to food every thirty or forty seconds. The pigeon will get the food no matter what it does. The food is a non-contingent reinforcer. That means that it is not dependent on any behavior. Although the pigeon can receive the same number of reinforcers just by standing still, we can expect the pigeon to develop some superstitious behavior. It might begin pecking the wall, turning in circles, or flapping its wings. Whatever the pigeon happens to be doing when the food appears will be reinforced. The pigeon will perform this behavior more often and it will be reinforced again the next time the food appears. Soon, the pigeon will be performing its act continuously even though it has nothing to do with the reinforcement (Skinner, 1948).

In human behavior, superstitions are common. They are particularly common, however, among athletes and gamblers. Can you think of why this would be? Superstitious behavior is most likely when there are many reinforcers and their timing is somewhat unpredictable. A person working for a salary gets paid every two weeks. For that person, the reinforcer is predictable and not very frequent. For an athlete or gambler, on the other hand, the reinforcers are much more frequent and much less predictable. A gambler never knows whether the next hand of cards or roll of the dice will be successful. A baseball pitcher never knows whether the next pitch will be a strike or result in a hit. We can also guess which kinds of athletes will be most likely to engage in superstitious behaviors. The ones who receive the most separate reinforcers and the ones for whom the reinforcers are the most unpredictable should be the most superstitious. Think about marathon runners, tennis players, ski racers, and baseball pitchers. What we know about the nature of reinforcers lets us make a scientific prediction that tennis players and pitchers should show more superstitious behaviors than marathon runners or ski racers do. The next time you see a tennis match or a baseball game, watch for superstitious behaviors. Carefully observe the tennis players' behavior just before they serve and observe the brief rituals that the baseball players go through at the plate or on the pitchers mound. Much of what they do is probably superstitious behavior.

Operant Generalization and Discrimination

stimulus generalization The tendency to respond to a stimulus that is similar to the training stimulus

Earlier, we talked about the scientific principle of generalization. Responses in classical conditioning will generalize to other stimuli if they are similar to the training stimulus. This is also true in operant conditioning. If you teach your dog to sit to the command "sit," it is likely that the dog will also sit when you say "spit," or even "fit." This is called **stimulus generalization.** Sometimes, stimulus generalization can lead to incorrect or annoying responses. I had trained my dog Penny to come when she heard her name. When I got a cat named Barney, I found that Penny also came when I called the cat. Her response to "Penny" generalized to the stimulus "Barney" because they were similar sounds. As in classical conditioning, discrimination training can change this behavior. When I called "Penny" and the dog came, I praised her. When I called "Barney" and the dog came, I ignored her. After about a week, the dog came immediately when I called "Penny." When I called the cat, Penny didn't even look up. She had learned to discriminate between the two stimuli. The stimulus "Penny" had become a **discriminative stimulus.** A discriminative stimulus is a signal that a response will be followed by reinforcement. After successful discrimination training, we say that the response has come under **stimulus control.** If you look closely at the illustration of the original Skinner box (Figure 7.4), you can see that there is a light near the top of the box. Skinner found that he could use discrimination training to condition the rats to press the bar only when the light was on. If the rat is reinforced only when the light is on, pressing the bar when the light is off soon goes through extinction. The light becomes a discriminative stimulus and bar pressing comes under stimulus control. There are many discriminative stimuli in our daily lives. As scientists, we can predict that most people will put money in vending machines only if the lights in the machine are on and dial a phone number only if they hear a dial tone. Discriminative stimuli tell us that a response is likely to be followed by a reinforcer. If the discriminative stimulus is present, we are much more likely to make the response.

discriminative stimulus
A stimulus that serves as a signal that a response will be followed by a reinforcer

stimulus control
A response is said to be under stimulus control when a particular stimulus controls the occurrence or form of the response

Locking It In

1. Behavior that is caused by non-contingent or accidental reinforcement is called _______________ behavior.
2. If an animal responds to a stimulus similar to the training stimulus, it is showing ____________.
3. If an animal responds to one stimulus but not to another, it is showing _________________.
4. Conditioning an animal to respond to one stimulus and not to another is called _____________ _____________.
5. Critical Thinking: Why are baseball players likely to be more superstitious than long-distance runners?

Answers: 1) superstitious, 2) generalization, 3) discrimination, 4) discrimination training

Schedules of Reinforcement

Continuous and Intermittent Reinforcement When I was about eight years old, my mother bought a ceramic cookie jar in the shape of a pig. When it was new, she often filled the pig with cookies. I soon learned to check inside. I was often reinforced with a cookie. Over time, my mother made cookies less often but I continued to check the jar. Every so often, I'd still be rewarded with a cookie. Now, many years later, I don't have a cookie jar in the house and seldom see one. Every once in a while, I see a cookie jar like my mother's at an antique shop or in someone's home. To this day, whenever this happens, I look inside. The jar is always empty. I don't think I've found a cookie in a cookie jar in over 20 years but I still look. How can we give a scientific explanation of why this behavior has persisted for so long without being reinforced?

If behaviors stopped whenever they were not reinforced, many of our everyday behaviors would disappear. Operant behavior does not need to be reinforced every time to continue. For operant conditioning to be effective, reinforcement should be delivered consistently at first. It should follow every time the behavior occurs. Later, though, the reinforcer can be delivered less and less often without the behavior going through extinction. When a response is reinforced every time it happens, the animal is receiving **continuous reinforcement.** If the reinforcer only follows the response some of the time, the animal is receiving **intermittent reinforcement** (Skinner, 1933). B. F. Skinner, in an early experiment, had a group of rats pressing a bar to receive food pellets. In those days, the food pellets had to be made by hand in a press. One Friday, Skinner noticed that he was running short of pellets. There were probably not enough left to last through the weekend. He wasn't in the mood to make more food pellets so he tried to think of a way to make the ones he had last through the weekend. He hit on the idea of rigging the food delivery machine so that the rats had to press the bar twice to get one food pellet. He wasn't sure what would happen but, as a scientist, he thought that it would be interesting to find out. When he returned on Monday, he found that the rats were still pressing the lever for food. He also noticed that their rate of bar pressing had gone up. He started experimenting with different ways to present the food and discovered what he called *schedules of reinforcement* (Skinner, 1938).

continuous reinforcement
Reinforcing a response every time it occurs

intermittent reinforcement
Reinforcing a response only some of the times it occurs

This slot-machine player is being reinforced on a variable interval schedule

Ratio Schedules The original Skinner box had a wheel with a series of holes around the edge. The wheel sat above the tube that delivered a food pellet into a cup where the rat could eat it. Every time the rat pressed the bar, the wheel turned so that the next hole lined up with the tube and allowed one food pellet to drop into the cup. To make the rats press the bar twice for one food pellet, Skinner just put tape over every other hole on the wheel. Once he discovered that the rats continued to work under these conditions, he tried putting tape over more and more holes of the wheel. This kind of schedule of reinforcement is called a ratio schedule. In a ratio schedule there is always some ratio between the number of responses and the number of reinforcers. The more responses the animal makes, the more reinforcers it gets. On a **fixed-ratio schedule,** the rat must make exactly the same number of responses for each reinforcer. During Skinner's original intermittent schedule, the rat was reinforced for every other response. This is a fixed ratio of two responses per reinforcer (see Figure 7.5).

Imagine that we put tape over many of the holes in the food-delivery wheel of a Skinner box but we don't space the untaped holes evenly around the wheel. The rat still has to make a certain number of responses to get a reinforcer. Now, however, the number of responses varies. This is called a **variable-ratio schedule.** A slot machine reinforces humans using a variable-ratio schedule. The slot player has to make a certain number of responses to get their next reinforcer. The number, however, is not fixed—it varies.

fixed-ratio schedule
A schedule of reinforcement in which the animal must make an exact number of responses to receive a reinforcer

variable-ratio schedule
A schedule of reinforcement in which the animal is reinforced after a certain number of responses but the number varies

fixed-interval schedule
A schedule of reinforcement in which the animal is not reinforced for a fixed time interval after receiving a reinforcer

variable-interval schedule
A schedule of reinforcement in which the animal is not reinforced for a certain time interval after receiving a reinforcer but the interval varies

Interval Schedules Suppose that each time the rat gets a reinforcer, we shut off the food-delivery system for one minute. Now the rat can only get one reinforcer per minute no matter how many responses it makes. This is not a ratio schedule because there is no ratio between responses and reinforcers. Now reinforcement depends on the passage of time. Because the pattern of reinforcement is based on a time interval, this kind of schedule is called an *interval schedule.* If the time-out is always the same length, we call this a **fixed-interval schedule.** If it varies, we call it a **variable-interval schedule.** Interval schedules tend to produce slower rates of responding than ratio schedules. This makes sense since, with a ratio schedule, the more responses the animal makes, the more reinforcers it gets. With an interval schedule, the animal can only get so many reinforcers in a given time period no matter how many responses it makes. On a fixed-interval schedule, animals tend to make more responses just before it is time for the reinforcer to be available. For example, most of us check the mail more and more often as the usual delivery time approaches. On a variable-ratio schedule, animals tend to

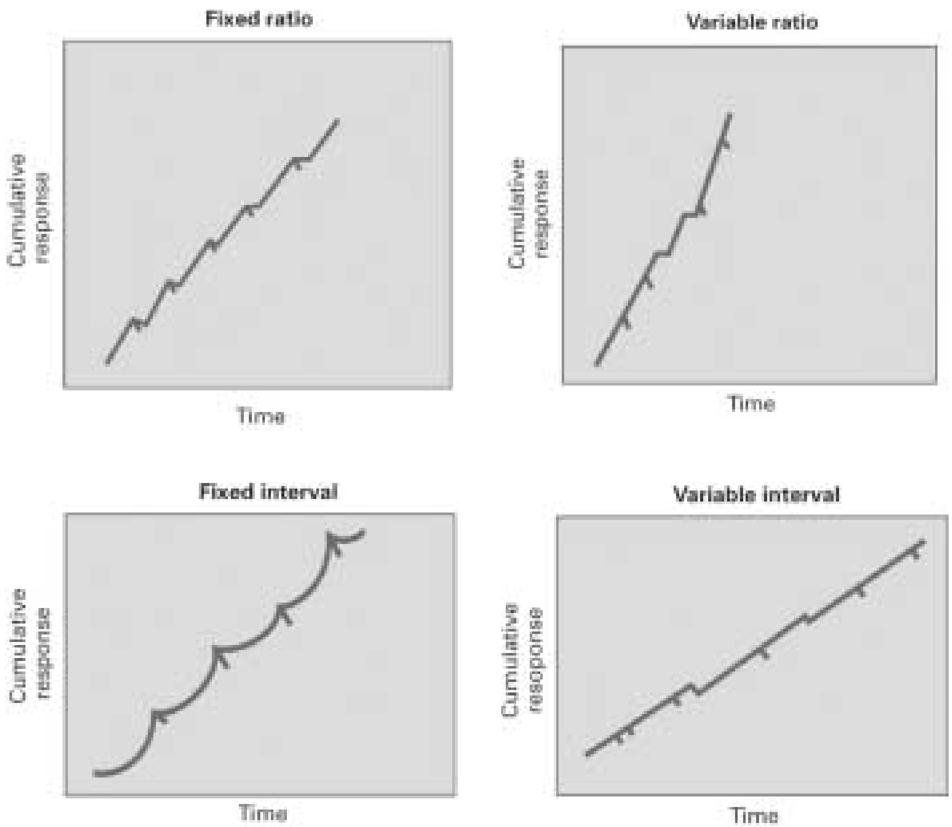

FIGURE 7.5

CUMULATIVE RECORD FOR FR, VR, FI, VI SCHEDULES OF REINFORCEMENT

maintain a slow but steady rate of response. This is what most of us do when redialing a busy phone number. We will get through after a certain amount of time but the time interval is unpredictable (see Figure 7.5). Generally, ratio schedules produce higher rates of response than interval schedules. This is not surprising since, on a ratio schedule, the faster the animal responds, the more reinforcers it will get. On an interval schedule, the number of reinforcers is limited by the schedule. You may work, for example, at a job where you get paid every two weeks. You can only receive one paycheck every two weeks no matter how often you make the response of visiting the payroll office.

Locking It In

1. Responses are reinforced every time they occur during ________________ reinforcement.
2. Responses are not reinforced every time they occur during ______________ reinforcement.
3. Reinforcing every other response is a _____________ schedule.
4. A slot machine reinforced gambles on a _________ schedule.
5. Most salaried workers are reinforced on a ___________ schedule.
6. People redialing a busy phone number are reinforced on a _____________ schedule.
7. Critical Thinking: Why do blackjack players prefer a fast dealer even though faster dealing means they lose their money sooner?

Answers: 1) continuous, 2) intermittent, 3) fixed interval, 4) variable ratio, 5) fixed interval, 6) variable interval

Making Punishment Effective

In the laboratory, behavioral scientists have shown that punishment can sometimes suppress undesirable behaviors. Outside the laboratory, however, it can be difficult to use effectively. In real life, it often does not work. Worse yet, even when it is effective, it often has disturbing side effects.

Q: Under what conditions is punishment effective?

As with reinforcement, punishment works best when it is immediate. Delayed punishment is often not effective at all. Anyone with a full-time job who has tried to use punishment to housebreak a dog is aware of how ineffective delayed punishment is. If the dog makes a mess at ten in the morning, punishing the dog at five in the afternoon will have little effect on the behavior. Delayed punishment is also ineffective with young children. Punishment must also be consistent if it is to work well. The punishment must follow the behavior every time it occurs. For many undesirable behaviors, this is impossible. For example, it is very unlikely that you will catch every lie that a child tells. Many behaviors that we might like to eliminate with punishment just don't happen when we are around. Parents often have little control over a child's behavior at school. Any punishment they deliver will be long after the behavior has occurred.

Because punishment works poorly when it is not immediate, there is little point in extended punishments. If a child is grounded for bad behavior, much of the punishment will be very delayed. Worse yet, the child may be behaving very well while grounded. Now the child is being punished for good behavior. If you use punishment to reduce the frequency of a behavior, try to make sure that the punishment is brief and follows immediately after the behavior.

It may seem obvious to say so, but punishment must also be unpleasant in order to be effective. My grandmother's idea of punishment was to threaten to spank me with a yardstick. Since I knew that she didn't actually own a yardstick, this punishment wasn't very effective. Sometimes, parents who feel guilty when their children are angry with them will follow punishment immediately with love, affection, and even treats. Since the punishment and the reinforcement both follow closely after the behavior, the punishment is not likely to be very effective. Because attention is often a powerful reinforcer, mild punishments are often ineffective. The reinforcing effects of the attention can outweigh the punishment. To be most effective, punishment should be delivered with as little attention as possible. To complicate things, experts on punishment recommend that punishment be no more severe than necessary to stop the behavior. This is because of the many negative side effects of punishment discussed below. So punishment should be severe enough to work, but no more severe than necessary. It can be very hard to guess ahead of time just how severe a punishment should be. In addition, parents who are angry often use more severe punishment than they need to.

To show how difficult it is to use punishment effectively, I ask my students two questions. First, I ask how many of them were punished for lying to their parents. Most of the hands in the class go up. Then, I ask how may of them lied to their parents anyway. Almost all the hands stay up. Their punishment was not immediate or consistent so it wasn't effective. Since they were more likely to be punished when their lies were poor ones, the main effect of this punishment was probably to make them more skillful liars. This is probably not what their parents had in mind.

Q: Can punishment ever be effective?

Sometime punishment can be effective. If a behavior can be punished immediately, every time it happens, and the punishment is severe enough, the behavior can be suppressed. Take, for example, a child who interrupts a parent while the parent is on the phone. The parent is always there when the behavior occurs. If the parent delivers punishment immediately after the behavior and the punishment is severe enough, the behavior will be suppressed. The fact that such punishment is effective, however, doesn't necessarily mean that it is a good choice for modifying behavior.

Side Effects and Problems with Punishment

Punishment often produces unwanted side effects in addition to suppressing the undesirable behavior. It also can lead to a variety of problems. These side effects and problems can make it a poor choice for modifying behavior in real life. Even when a method of changing behavior is effective, it still may not be a good idea. Let's take a look at some of these common problems with punishment.

Punishment Can Cause Anxiety and Emotional Behavior and Can Lower Self-esteem Punishment tends to make animals anxious and emotional (Deater-Deckard & Dodge, 1997; Skinner, 1953). For children, it can lower self-esteem (Larzelere et al., 1989). Few parents want to raise children with low self-esteem. Nor do they want children who are nervous and emotional. Because it increases emotional behavior, punishment is an especially bad choice for reducing emotional behaviors like crying. People commonly try to change behavior in ways that they think should work. No one can expect every method they use to be effective. As scientists, we should be willing to give up methods that don't work and try other methods that may be more effective.

Punishment Suppresses Behaviors but Doesn't Eliminate Them Established behaviors that have been punished will go away temporarily but tend to show up again later (Azrin & Holz, 1966). If the behavior is not punished when it reappears, it will be back at full strength. If a rat has been trained to press a bar to get food, we can suppress bar pressing by giving the rat a painful electric shock every time it presses the bar. We can expect, however, that the rat will eventually press the bar again. If the rat is not shocked after pressing the bar, it will go right back to its original rate of bar pressing. The response is still "in there" it has just been suppressed. Imagine a burglar who has been reinforced for breaking into houses. Punishment may suppress this response but won't really eliminate it. Eventually, the burglar will break into a house again. If there is no punishment for the behavior, we can expect the burglar to go right back to his or her old rate of burglary.

Punishment Leads to Escape and Avoidance Behaviors One of the great myths of punishment is that children learn to tell right from wrong by being punished for doing wrong. There is little or no scientific evidence to support this idea. In fact, prisons are full of people who were punished every day for their bad behavior. Many adults who were seldom punished as children have a strong sense of right and wrong. As scientists, we have to give more weight to the scientific evidence than to common sense ideas no matter how often we hear them.

Children seem to learn their moral values by having a warm, affectionate relationship with their parents (Mussen, Conger, & Kagan, 1969; Sears, Maccoby, & Levin, 1957). They *model* their moral behaviors after those of their parents. If anything, punishment actually interferes with this process (McCord, 1991). Punishment seems to weaken the relationship between a parent and child. It makes the child less likely to model the parent's behavior. Escape and *avoidance* behaviors also interfere with this modeling process. During punishment, most animals will try to escape. Rats in a cage who are getting electric shock will try desperately to get out of the cage. Children who are being spanked will try to get away from whoever is spanking them. Every year, children run away from homes where they are punished. Animals will also try to avoid the situation in which they were punished. This is called avoidance. Avoidance makes punishment effective if the animal avoids the behavior that resulted in punishment. Unfortunately, the animal also avoids the situation where it was punished and the person who did the punishing (Azrin & Holz, 1966). Children who receive a lot of punishment will spend much of their time away from home. They can be found at the park, at a friend's house, at the mall—anywhere but the place they are punished. If they are home, they will try to stay away from the person who punishes them. Children can't learn to model a parent's behavior if they constantly avoid the parent.

Punishment Provides Poor Feedback Punishment, even when effective, usually gives information about what you shouldn't do. Unfortunately, it doesn't do a very good job of telling you what you *should* do. A child who is punished for leaving a jacket on the floor by the front door may learn not to leave the jacket in that particular spot. The punishment, however, doesn't reinforce them for putting the jacket where it belongs. Since there are many places where the jacket doesn't belong, it may be quite a while before the correct response is learned. Using shaping to teach the correct response is almost always faster than using punishment.

Punishment Can Lead to Aggressive Responses One of the most troubling side effects of punishment is its tendency to create aggressive responses. Azrin and Holz (1966) found that if two rats in the same cage are given electric shock, their first response is to try to escape. After the escape behaviors are extinguished, however, the rats soon turn on each other. The punishment appears to make them aggressive and violent. Children who are punished can turn on their parents too. When the punishment is mild, the children may "get even" with their parents by being irresponsible or by doing things to annoy their parents. In extreme cases, children sometimes injure or even kill their parents.

Q: Is punishment a direct cause of this kind of behavior?

We can't perform the scientific research that would give us a clear answer to this question. It would be unethical to punish children intentionally just to observe the results. Instead, we have to observe existing families. As scientists, we try to find patterns of behavior that will help explain the effects of punishment on children.

Punishment Can Model Violent Behavior When a parent uses physical punishment, he or she is providing a violent model for the child. The parent has a problem with the child's behavior and is trying to solve that problem with physical violence. We know that children imitate the behaviors of their parents. If a parent solves problems with punishment, the child is likely to do the same (Patterson, 2002; Reid et al., 2002). If a child often receives physical punishment, it is likely that they will grow up to be a more violent person. This kind of imitation is called *modeling.* In one study, trained observers recorded violent behavior at playgrounds in Germany, Denmark, and Italy. They recorded aggressive acts between adults and children, and between children. They found that where adults were aggressive toward children, children were more violent to other children (Hyman, 1997). We will be discuss modeling in more detail in the section on observational learning later in this chapter.

Q: Is punishment ever a good idea?

Psychologists don't agree on whether punishment is ever an acceptable way of changing behavior. We have seen that punishment is difficult to use effectively. We have also seen that even when punishment is used effectively, it can lead to emotional behavior, low self-esteem, escape and avoidance behavior, and aggression. It can suppress behavior without actually getting rid of it. It serves as poor feedback about what the person *should* do. It can also model undesirable behaviors and lead to violence and revenge. Most psychologists agree that punishment should be avoided whenever possible. Some psychologists, however, believe that there are situations where punishment is appropriate. Some behaviors like playing with weapons, or running into traffic, put the child or others in serious danger. If the parent is present and can deliver punishment immediately, punishment can suppress these behaviors. Sometimes a single punishment can be effective if it follows immediately after the behavior. Some psychologists believe that in these cases punishment is justified. Others believe that the negative effects of punishment are bad enough that it should never be used. These psychologists argue that parents should modify the environment so that the punishment isn't necessary. Although psychologists don't agree on whether punishment is ever necessary, they generally agree that the scientific evidence strongly supports two basic statements about punishment. First, much of the punishment that occurs in everyday life is not effective. Second, even when punishment is effective, it is often a bad idea. Many undesirable behaviors can be eliminated either by using extinction, or by reinforcing an alternative behavior. These techniques work as well or better than punishment in most situations and do not have punishment's undesirable side effects.

Going Beyond the Data

Can Prisons Work?

There are several reasons for having prisons. One obvious one is to keep dangerous people from harming the rest of us. Putting a person in prison can also satisfy society's desire to punish wrongdoers for their antisocial acts. Crime victims and their families are also thought to receive satisfaction by the imprisonment of the criminal. Let's look, however, at one of the primary functions of prison in modern society. Prisons are supposed to reduce the crime rate through the use of punishment. The threat of prison is meant to deter people from committing crimes. Common sense tells us that this should work. The traditional idea is that if we punish antisocial behaviors, they will become less common. Prison is the primary punishment for criminal behaviors in our society. Earlier, we said that to be effective punishment should be immediate, consistent, brief, and unpleasant. Does a prison sentence meet these requirements? Is it immediate or consistent? Many crimes go unpunished. Criminals are often arrested and not charged. They are often tried but not convicted. Even if they are convicted, the prison sentence usually comes months or even years after the crime. Is the punishment brief? Even the shortest prison sentences usually involve months of punishment. Is prison unpleasant? For many prisoners, prison is a very unpleasant punishment. Other criminals, though, have lives filled with danger, uncertainty, and deprivation. For these criminals, prison may actually be less unpleasant than their lives outside the prison. Mass murderer Charles Manson was in prison before the sensational murders that led to his life sentence. During his earlier parole hearing, Manson said that he was uncomfortable outside of prison. He told the parole board that, if he was released from prison, he would probably do something to get back inside. Even when prison is unpleasant, we can't expect it to be an effective punishment if it is not immediate or consistent. In the 1980s, the prison population in the United States doubled without any significant reduction in the crime rate (Lashmar, 1994). In fact, the United States has a greater proportion of its citizens in prison than any other industrialized nation in the world. If tough prison sentences lowered the crime rate, we would also expect to have the lowest crime rate in the world. Instead, we have one of the highest.

Another problem with prison is that in most prisons there is no program to actually eliminate the criminal's antisocial behavior. Consider the career burglar. For some time, the person has been reinforced regularly for breaking into houses. The reinforcement is immediate and consistent. Eventually, the person is caught and sentenced to prison. The punishment, if it occurs at all, comes a long time after the behavior. It may suppress the behavior temporarily. Sooner or later, though, the person will be released from prison and the odds are

(continued)

good that the behavior will resume. Worse yet, the person may have learned how to avoid getting caught. From other prisoners, they may learn how to defeat alarm systems or how to select houses to burglarize. The net result of this very expensive procedure may be to actually make the person a better burglar.

There is no question that there are dangerous people in the world. Such people need to be isolated from society to prevent them from hurting the rest of us. We will probably always need prisons for this purpose. Based on what we know about how punishment works, however, we shouldn't expect that the threat of going to prison will have much effect on the crime rate. Prison is a particularly ineffective form of punishment. There is no way of making the punishment of criminals immediate or consistent without taking away the civil rights of citizens. Therefore, if we want to reduce the crime rate, we have to use other methods. We can identify children who are at risk of becoming criminals. Intervening early in their lives and finding ways to prevent or extinguish their antisocial behaviors is likely to be a much more effective way of preventing criminal behavior (Parker, 1997). In the long run, it is also much less expensive. As scientists, we should look to the research evidence to find out what works rather than relying on what common sense tells us *should* work.

The Evolution of Conditioning—Biological Predispositions

The process of *natural selection* is expressed by Charles Darwin's theory of evolution. Characteristics that help a species survive are more likely to persist in that species. Behaviors like eating, drinking, and breathing, have obvious survival value. Other behaviors like imitating successful adults, exploring the environment, and protecting our children are more complex, but are also clearly valuable for the survival of our species. B. F. Skinner (1984, 1990) pointed out an important limitation of natural selection, however. Natural selection only prepares you for an environment that is like the environment of your ancestors. If your ancestors lived in a relatively mild environment, shivering may have evolved as a way to keep warm on cool days. If you live in northern Canada, however, shivering is not a very effective method for dealing with the cold. Skinner argues that another form of selection evolved to make animals better able to cope with their immediate environment. He called this process *selection by consequences* and it is the basis for operant conditioning. In operant conditioning, behaviors that result in favorable consequences are *selected* to occur more often in the future. If a behavior results in finding food, we can expect that behavior to become more common. This is true even if your ancestors didn't find food that way. If putting on a warm coat helps warm you up, we can expect that people in cold climates will often put on coats. Selection by consequences helps us adapt to an ever-changing environment.

Earlier in this chapter, we saw how an animal's biological programming can affect the learning of classically conditioned responses. We learned, for example, that it is easier to teach someone to be afraid of a rat than a beach ball. This is because nature has prepared us to be afraid of certain things like rats, snakes, and heights. In the work of Rescorla, and Garcia, we saw how some stimulus-response connections are more easily learned than others are. A number of researchers have observed something similar in operant conditioning. Because animals are prepared by nature to learn some things more easily than others, biological factors can play a role in operant conditioning too. Keller Breland and Marian Breland (1961) were associates of B. F. Skinner's who left the academic world to become world-famous animal trainers. They used the scientific principles of operant conditioning to train over 6000 animals for television shows, circuses, movies, and animal parks. From their work, the Brelands concluded that *biological predispositions* were important in determining what an animal could learn. Raccoons, for example, can be taught to perform tricks to receive wooden tokens they can use to "buy" food. It was difficult, however, to get the raccoons to stop washing the tokens in water. They were not reinforced for this behavior and it delayed them from receiving the food. Raccoons in the wild often wash their food in water. It appears, then, that raccoons have a biological predisposition to wash reinforcers that have been associated with food.

Many species have natural behaviors (like the raccoon's tendency to wash food in water) that can influence what they can, and can't be trained to do

In a more formal scientific test of the idea of biological predispositions, Foree & LoLordo (1973) taught pigeons to get food or avoid shock by pecking or flapping their wings. The researchers found that the pigeons easily learned to peck for food and flap their wings to avoid shock. These were natural behaviors for the pigeons. Pigeons naturally use their beaks to get food and flee from danger by flapping their wings. The animals had a very hard time, however, learning the behaviors the other way around. Teaching them to flap their wings to get food or peck to avoid shock was much more difficult. Biological programming makes it much easier for animals to learn certain behaviors. The behaviors they learn easily seem to be related to natural responses that have evolved for that species. This effect can be found in both classical and operant conditioning.

COGNITIVE LEARNING

Cognitive learning is a general term and one that is not always used in the same way by all psychologists. Some psychologists (strict behaviorists in particular) don't accept the principle of cognitive learning at all. Others insist that cognitive learning is common and that much of what we know is the result of cognitive processes. Cognitive learning usually refers to all learning that does not occur through traditional classical or operant conditioning. Usually, what is learned is more complex than a simple response to a stimulus. In cognitive learning, we often learn by observing, by reading, by imitating others, or by reasoning. As we discussed in Chapter one, cognitive psychology developed when psychologists tried to give scientific explanations of complex human behaviors like language, memory, and problem solving. These, and many other common human behaviors, are very hard to explain scientifically using just classical and operant conditioning.

Vicarious Conditioning

If a scientist were to give you a painful electric shock every time a small red light came on, you would soon have an emotional reaction to the light. When the light came on, your heart would speed up and your breathing would become more rapid. You might begin to sweat. In fact, it is likely that you would experience a mild version of the "fight or flight" response—the body's normal reaction to an emergency. Psychologists have found that this reaction can be produced without ever shocking the subject. Bandura and Rosenthal (1966) gave subjects an electric shock whenever a light came on. The subjects soon learned to respond emotionally to the light itself. They also found, though, that other subjects who merely watched this process also became aroused when the light came on. This is called **vicarious conditioning.** In vicarious conditioning, people learn emotional responses by watching what happens to someone else. Some of the classical conditioning used in advertising depends on vicarious conditioning. We often see someone being embarrassed, frightened, or happy in an advertisement. The advertisers hope that we will later experience the same emotion in similar situations. They hope that our emotional response will make us more likely to buy their product. Vicarious conditioning also explains the fears that people can develop after watching a horror movie. In a very frightening scene in Alfred Hitchcock's Movie, *Psycho*, actor

vicarious conditioning
Conditioning that occurs when an animal observes another animal being conditioned

Actress Janet Leigh developed a conditioned fear of taking showers from watching herself in this scene from the movie *Psycho*

Janet Leigh is stabbed to death while taking a shower in a motel. Years later, Leigh confessed that after appearing in the movie she never took another shower. I saw the movie only once, many years ago. To this day, I'm afraid to take a shower in a motel. Many of my friends confess that they have the same fear. I have never been punished in any way for taking a shower in a motel so my fear can't be the result of simple classical conditioning. Janet Leigh revealed in 1995 that her fear was not a result of filming the famous shower scene. She acquired her fear of showers along with many of the rest of us, by watching the movie. A cognitive psychologist would argue that "higher mental processes" like thinking and imagining are necessary to explain these examples of vicarious conditioning.

Cognitive Maps and Latent Learning

Imagine that you are on the way to a friend's house. As you come around a corner, you find that a tree has fallen across the road. Your usual route to your friend's house is completely blocked. What do you do? Most likely, you quickly find another way to get to your friend's house. How could we give a scientific explanation of this behavior? Your method for getting to your friend's house is probably not just a series of linked right and left turns. More likely, it is a path across a map you keep in your head. This map is a kind of mental picture of how your town is laid out. When you find your usual route blocked, you consult this mental map and find a new way to get to the same place. Sometimes, the map is wrong. Most of the time, however, a blocked street is just a minor inconvenience. This kind of internal representation of the real world is called a **cognitive map.**

cognitive map
A mental image of an animal's environment that the animal can use to guide its movements

Q: Since we can't see a person's cognitive map, how do we know they really have one?

As scientists, we must be somewhat skeptical about maps that no one can see. How could we produce scientific evidence that such maps exist? Suppose that you move to a new town. After you've been in the new town for a month, we ask you do draw a map of the town. Your map might be incomplete but it would probably be fairly accurate in showing the routes you take often. What if we asked you to draw a new map every few months? Each new map would probably be more accurate and complete than your earlier maps. Notice that the maps are not getting better because your map-making skills are improving. The maps show your increasing knowledge about the town. We could prove this by making you produce a second map each time of a town you had only visited once. These second maps wouldn't change much over time. The maps of your new hometown, however, would improve as your cognitive map of the town became more accurate. Notice that no one is reinforcing you for the accuracy of the maps you draw. Why are they improving? The odds are that, in your everyday life, having an accurate map of the town lets you reach your destinations more quickly. The sooner you get where you are going, the sooner you can receive whatever reinforcers are waiting there for you. This reinforces you for improving your cognitive map.

Researchers in 1930 (Tolman & Honzik) did an experiment designed to test **latent learning** in rats. Two groups of rats were placed in a maze for little while each day (see Figure 7.6). One group of rats received a food reward in a particular part of the maze. The other group received no food at all in the maze. The food-rewarded rats soon learned to run quickly through the maze to the food box. The rats that were not given food showed no signs of learning. Later, when the second group of rats discovered food in the food box, they began running directly to the food. They quickly became just as fast as the other group. Tolman and Honzik concluded that even though they were not rewarded with food during the first part of the experiment, they had developed a cognitive

latent learning
Learning that occurs in the absence of obvious reinforcers and only appears after reinforcement is introduced

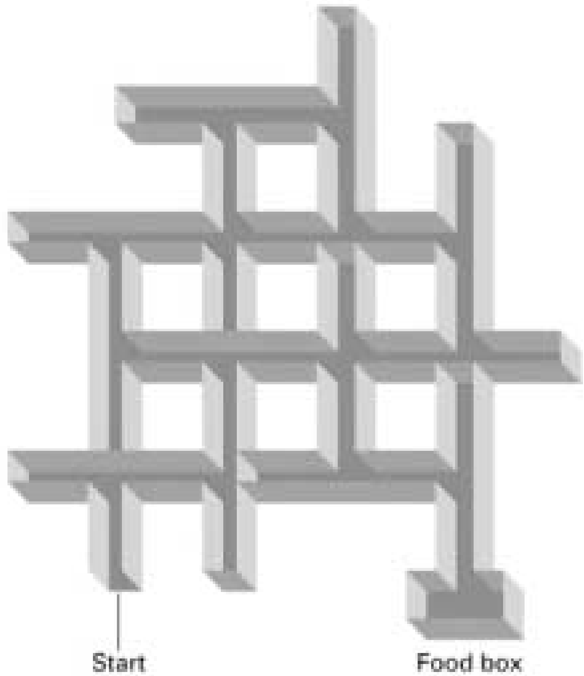

FIGURE 7.6

COGNITIVE MAPS

E. C. Tolman used a maze like this one to investigate latent learning in rats

map of the maze. When the food became available, they used their cognitive map to find it quickly. What they had learned while wandering through the maze earlier was not a simple response. They had learned the layout of the maze much as you and I have learned the way around the town we live in. If we hear from a friend that someone is giving away twenty-dollar bills at Fifth and Main, we can use our knowledge to get there as quickly as possible.

Q: Did these rats learn without any reinforcement at all?

The rats learned to navigate the maze without being reinforced by food. It seems likely, however, that there were other reinforcers at work. Animals are often reinforced for developing and maintaining accurate cognitive maps of their environments. Hunters remember where they have found game and are often reinforced for doing so. I have a friend who is an enthusiastic golfer. He can draw you very accurate map of his favorite golf course. He claims that when he has trouble sleeping, he plays the course in his head. Because his knowledge of the course helps him get lower scores, my friend is often reinforced for maintaining this map of the course. The rats in the experiment may have been reinforced in the past for having accurate cognitive maps of the world around them. If this behavior was reinforced in a variety of settings, generalization would take effect. We would expect the rats to construct cognitive maps in any new environment, including the maze. As scientists, we can't say that the rats learned without being reinforced. We can only say that food was not the reinforcer. We should also consider the possibility that making accurate cognitive maps has survival value. Perhaps making cognitive maps has evolved because it helps us survive long enough to reproduce. This means that making cognitive maps may not be a learned behavior at all.

Modeling—Imitating the Behavior of Others

I mentioned my friend the golfer in the previous section. One day, after we had finished a round of golf together, I asked him if he had noticed anything about my golf swing that needed improvement. To my surprise, he answered that he

modeling
The tendency to imitate the behavior of another person

had never seen it. He said that if he watches someone else's swing, he begins to imitate it. Sometimes the results are disastrous. **Modeling** is a powerful force in shaping human behavior. We imitate our parents, our siblings, people we see in films and television shows, even cartoon characters. An important area of study for psychologists is the question of whether watching violent behavior in others makes us more violent. How could we, as scientists, test the theory that we imitate violent behavior when we see it?

By far, the best-known scientific experiments on modeling are those performed by **Albert Bandura** and his associates (Bandura et al., 1961, 1963). In these experiments some children observed adults attacking an inflatable Bobo doll. The Bobo doll is an inflatable toy with weight in the bottom. It pops back up when it is knocked down. In a typical experiment, children in one group were playing quietly when an adult in the room went over to the Bobo doll and began attacking it. The adult beat on the doll for about ten minutes, kicking it, hitting it with a hammer, knocking it down, and sitting on it. During this performance, the adult also spoke out loud saying things like "Sock him in the nose," and "Kick him." The other group of children saw no aggressive behavior at all. One at a time, the children were then taken into another room. This new room contained a number of attractive toys and another Bobo doll. In the new room, each child was told that the toys were being saved for other children and that he or she was not allowed to play with them. The child was then left alone and watched through hidden windows. The children who had seen the aggressive model were much more likely to attack the Bobo doll in the second room. Often, they attacked it in the same way as the adult had. They even repeated the phrases they had heard the adult use. They appeared to be modeling their behavior closely after the behavior of the adult. Bandura and his associates found that the children imitated the violent behavior of adults even when the adult violence they saw was on film. Bandura's research raises many questions about showing children violent images on television or in the movies. We will discuss the effects of violence in the media further in Chapter 14.

Q: Does reinforcement or punishment of the model play a role in modeling?

In Bandura's early studies, the adult models received no punishment or reinforcement for their violent behavior. Mary Rosekrans and Willard Hartup (1967) designed a scientific experiment designed to find out if what happened to the adult model after their violent outburst would make a difference in the child's imitation of the behavior. In this chapter's *Science of Psychology Journal*, we take a look at their experiment.

Science of Psychology Journal

Modeling Violent Behavior and its Consequences

Background

Bandura and his associates demonstrated that children would model violent behavior even when no reinforcement was given either to the adult or the child. Rosekrans and Hartup designed an experiment to test the effects of observed reinforcement or punishment on the child's behavior.

The Experiment

Rosekrans and Hartup had preschool children watch an adult female model playing with a Bobo doll and some clay figures. She displayed violence toward the toys. She stuck a fork in the clay figures and hit the Bobo doll with a hammer. While attacking the toys, she made hostile comments such as "Punch him, punch him," and "I'll knock your head off." Another adult was present during the attack and had one of two reactions. In the reinforcement condition, the other adult made positive comments about the aggressive behavior such as, "Good for you! I guess you really fixed him that time." In the punishment condition, the second adult made negative comments about the aggressive behavior such as, "Now look what you've done, you've ruined it." Children in the reinforcement group saw the violent model being reinforced. Children in the punishment group saw the violent model being punished. Children in a third group (the "no model" group) saw no violent behavior at all. Children in all three groups were then allowed to play alone with the toys. Hidden observers watched to see if they hit, stabbed, or were otherwise violent toward the toys.

Results

The children who saw the adult model reinforced for violent behavior were much more likely to be violent themselves. The children who saw the adult model punished actually showed slightly less violent behavior than the "no model" group.

Discussion

It appears that children do not just blindly model any behavior they see an adult perform. They are much more likely to imitate a behavior if they have seen someone reinforced for performing the same behavior. In the same way, if they have seen someone punished for a behavior, they are much less likely to model that behavior. If we want to prevent violence, we should not show children violent models who are reinforced following their violent behavior. Unfortunately, we often do just that. In many movies and television shows, the main characters are frustrated in reaching some goal. Toward the end of the show, they engage in some spectacular violence. Once the violence is over, they have reached their goal and are often reinforced with praise and admiration. Sometimes they are reunited with loved ones; other times they experience romantic success. Rosekrans and Hartup's results suggest that showing this sequence of events to children over and over will make them more violent in real life. Other studies have supported this hypothesis (Huesmann et al., 2003). See Chapter 14 for more information on media violence and its effect on violence in the real world. Each year, the average American child spends more time in front of the television than they do in school. Limiting the amount of violence that children see might help to reduce the level of violent behavior that we see all around us.

Chapter Summary

What are the three basic kinds of learning?

- Learning is defined as a relatively permanent change in behavior due to experience. The three basic kinds of learning are classical conditioning, operant conditioning, and observational learning.

How did Ivan Pavlov discover classical conditioning?

- While studying digestion in dogs, Ivan Pavlov learned that if a neutral stimulus (a bell) is repeatedly paired with an unconditioned stimulus (food) that already elicits the unconditioned response (salivation), the neutral stimulus will come to elicit the same response.

Who made Little Albert afraid of the rat and how was it done?

- J. B. Watson made Little Albert afraid of the rat by using classical conditioning. He presented the rat together with a loud noise.

How are behaviors learned and unlearned in classical conditioning?

- During the acquisition phase of classical conditioning, a neutral stimulus (CS) is repeatedly paired with an unconditioned stimulus (US) that already elicits the unconditioned response (UR). After conditioning, the CS will elicit the same response. The response to the CS is called a conditioned response (CR).

- Once an animal is conditioned to respond to a particular stimulus, it will make the same response to other similar stimuli. This is called generalization. The animal can be taught discrimination, (responding differently to different stimuli) through a process called discrimination training.
- In classical extinction, the CS is presented over and over without the US. During extinction, the response to the CS grows weaker at weaker and finally disappears. It may occur again after a rest period. This return of the response is called spontaneous recovery.

What did Robert Rescorla, Leon Kamin, and John Garcia, add to our understanding of classical conditioning?

- Rescorla, Kamin and Garcia showed that the process of classical conditioning is more complicated than Pavlov or Watson thought.
- Robert Rescorla found that to produce conditioning, the CS must be a reliable predictor of the US.
- Leon Kamin demonstrated that if a classically conditioned response already exists, learning that response to a new stimulus may be "blocked." He called this process "blocking."
- John Garcia's research on taste aversions showed that classical conditioning can occur even when the US does not follow immediately after the CS.

Does classical conditioning explain all human learning?

- Classical conditioning only explains the learning of conditioned reflexes and emotional behaviors.
- E. L. Thorndike and B. F. Skinner extended our understanding of conditioning by investigating responses that are learned or modified depending on their consequences.
- E. L. Thorndike found that a cat in a puzzle box would learn behaviors that allowed it to escape from the box. He believed that escaping from the box, a positive consequence, "stamped in" their correct responses. He called this the "law of effect."
- B. F. Skinner studied "operant" behaviors that were not reflexes or emotional behaviors. Operant behaviors are made freely by the animal rather than being "elicited" by a particular stimulus. Skinner argued that following an operant behavior with reinforcement or punishment could make it occur more or less often.

How are behaviors learned and unlearned in operant conditioning?

- In operant conditioning, behaviors are learned when they are followed by reinforcing stimuli. A reinforcer is defined as anything that increases the frequency of the behavior it follows.
- Primary reinforcers such as food and water are built in by nature and do not need to be learned. Secondary reinforcers get their power by being associated with primary reinforcers. Generalized reinforcers are good any time because they have been associated with several different primary reinforcers.
- As in classical conditioning, once an animal is conditioned to respond to a particular stimulus, it makes the same response to other similar stimuli. This is called generalization. The animal can be taught discrimination, (responding differently to different stimuli) through a process called discrimination training.

Eliminating responses in operant conditioning can be done with extinction or punishment.

- In extinction, the response is no longer followed by reinforcement. As with classical conditioning, operant extinction is a gradual process and the response will usually show spontaneous recovery and reappear after a rest period.

- In punishment, the response is followed by an aversive stimulus. Over time, the frequency of the response decreases.

How does the schedule of reinforcement affect operant conditioning?

- With continuous reinforcement, the response is reinforced every time it occurs. During intermittent reinforcement, the response is not reinforced every time it occurs. There are two kinds of intermittent schedules of reinforcement: ratio schedules and interval schedules.
- On ratio schedules, the animal must make a certain number of responses to receive a reinforcer. The number may be fixed (fixed-ratio schedule) or it may vary (variable-ratio schedule). Ratio schedules generally produce higher rates of responding than interval schedules.
- On interval schedules, the animal will not be reinforced for a certain period of time after receiving a reinforcer. The interval may be fixed (fixed-interval schedule) or it may vary (variable-interval schedule).

Can punishment be an effective method of eliminating undesirable behaviors?

- Punishment can suppress undesirable behaviors if it is immediate, consistent, brief, and aversive. However, in everyday life, it is often difficult to meet these standards. Using prison to punish criminals, for example, is unlikely to be effective in suppressing their criminal behavior because the punishment is not immediate, consistent, or brief, and may not be aversive for some criminals.

Are there problems or negative side effects associated with the use of punishment?

- Punishment has a number of problems and negative side effects. It can cause anxiety and emotional behaviors. It can suppress behaviors without really eliminating them. It can lead to escape and avoidance behaviors. Punishment provides poor feedback because it doesn't give any information about what behavior is desired. It can also produce aggression and hostility. Adults who use physical punishment are modeling violent behavior.

Can we learn without undergoing classical or operant conditioning?

- In vicarious conditioning, conditioned responses can be learned by observing another animal being conditioned.
- The research of Albert Bandura and his associates demonstrated that children imitate the behavior of adults through a process called modeling.

Does watching violent shows make children more violent?

- A number of researchers have found that children imitate the behavior of adult models. This is true even when they view films or videos of adult behavior. It the adults behave violently, it appears that the children will become more violent. This is especially true when the adult model it reinforced for violent behavior—a common feature of much media violence.

Important Terms and Concepts

learning
stimulus
response
habituation a decrease in response to a repeated stimulus
classical conditioning
unconditioned stimulus (US)
unconditioned response (UR)
conditioned stimulus (CS)
unconditioned response (UR)
extinction
spontaneous recovery
generalization
discrimination
discrimination training

- conditioned emotional response
- systematic desensitization
- biological preparedness
- blocking
- taste aversion
- law of effect
- operant conditioning
- reinforcer
- punishment
- positive reinforcement
- negative reinforcement
- positive punishment
- negative punishment
- shaping
- sensory reinforcer
- primary reinforcer
- secondary reinforcer
- generalized reinforcer
- Premack principle
- superstitious behavior
- stimulus generalization
- discriminative stimulus
- stimulus control
- continuous reinforcement
- intermittent reinforcement
- fixed-ratio schedule
- variable-ratio schedule
- fixed-interval schedule
- variable-interval schedule
- vicarious conditioning
- cognitive map
- latent learning
- modeling

Important Names

- Ivan Pavlov
- J. B. Watson
- Robert Rescorla
- Leon Kamin
- John Garcia
- E. L. Thorndike
- B. F. Skinner
- David Premack
- Albert Bandura