Chapter 7
Learning

The topic of this chapter is learning—*the relatively permanent change in knowledge or behavior that is the result of experience*. Although you might think of learning in terms of what you need to do before an upcoming exam, the knowledge that you take away from your classes, or new skills that you acquire through practice, these changes represent only one component of learning.

Learning is perhaps the most important human capacity. Learning allows us to create effective lives by being able to respond to changes. We learn to avoid touching hot stoves, to find our way home from school, and to remember which people have helped us in the past and which people have been unkind. Without the ability to learn from our experiences, our lives would be remarkably dangerous and inefficient. The principles of learning can also be used to explain a wide variety of social interactions, including social dilemmas in which people make important, and often selfish, decisions about how to behave by calculating the costs and benefits of different outcomes.

The study of learning is closely associated with the behaviorist school of psychology, in which it was seen as an alternative scientific perspective to the failure of introspection. The behaviorists, including John B. Watson and B. F. Skinner, focused their research entirely on behavior, to the exclusion of any kinds of mental processes. For behaviorists, the fundamental aspect of learning is the process of conditioning—*the ability to connect* stimuli *(the changes that occur in the environment) with* responses *(behaviors or other actions).*

But conditioning is just one type of learning. We will also consider other types, including learning through *insight,* as well as *observational learning* (also known as *modeling*). In each case we will see not only what psychologists have learned about the topics but also the important influence that learning has on many aspects of our everyday lives.

7.1 Learning by Association: Classical Conditioning

Pavlov Demonstrates Conditioning in Dogs

In the early part of the 20th century, Russian physiologist Ivan Pavlov (1849–1936) was studying the digestive system of dogs when he noticed an interesting behavioral phenomenon: The dogs began to salivate when the lab technicians who normally fed them entered the room, even though the dogs had not yet received any food. Pavlov realized that the dogs were salivating because they knew that they were about to be fed; the dogs had begun to associate the arrival of the technicians with the food that soon followed their appearance in the room.

With his team of researchers, Pavlov began studying this process in more detail. He conducted a series of experiments in which, over a number of trials, dogs were exposed to a sound immediately before receiving food. He systematically controlled the onset of the sound and the timing of the delivery of the food, and recorded the amount of the dogs' salivation. Initially the dogs salivated only when they saw or smelled the food, but after several pairings of the sound and the food, the dogs began to salivate as soon as they heard the sound. The animals had learned to associate the sound with the food that followed.

Pavlov had identified a fundamental associative learning process called *classical conditioning*. Classical conditioning refers to *learning that occurs when a neutral stimulus (e.g., a tone) becomes associated with a stimulus (e.g., food) that naturally produces a behavior*. After the association is learned, the previously neutral stimulus is sufficient to produce the behavior.

As you can see in Figure 7.3 "4-Panel Image of Whistle and Dog", psychologists use specific terms to identify the stimuli and the responses in classical conditioning.
The unconditioned stimulus (US) is *something (such as food) that triggers a natural occurring response*, and the unconditioned response (UR) is *the naturally occurring response (such as salivation) that follows the unconditioned stimulus*. The conditioned stimulus (CS) is *a neutral stimulus that, after being repeatedly presented prior to the unconditioned stimulus, evokes a similar response as the unconditioned stimulus*. In Pavlov's experiment, the sound of the tone served as the conditioned stimulus that, after learning, produced the conditioned response (CR), which is *the acquired response to the formerly neutral stimulus*. Note that the UR and the CR are

the same behavior—in this case salivation—but they are given different names because they are produced by different stimuli (the US and the CS, respectively).

Figure 7.3 4-Panel Image of Whistle and Dog

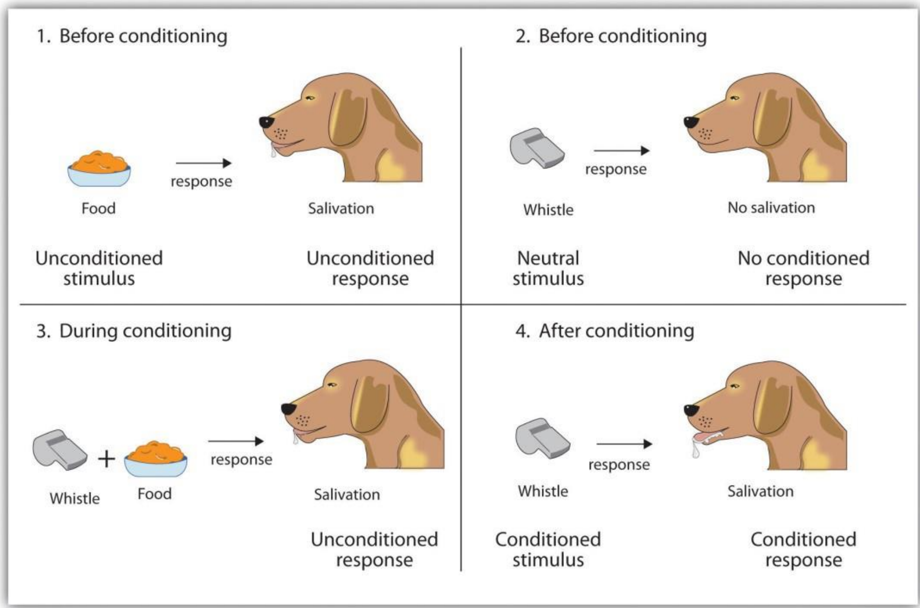

Top left: Before conditioning, the unconditioned stimulus (US) naturally produces the unconditioned response (UR). Top right: Before conditioning, the neutral stimulus (the whistle) does not produce the salivation response. Bottom left: The unconditioned stimulus (US), in this case the food, is repeatedly presented immediately after the neutral stimulus. Bottom right: After learning, the neutral stimulus (now known as the conditioned stimulus or CS), is sufficient to produce the conditioned responses (CR).

Conditioning is evolutionarily beneficial because it allows organisms to develop expectations that help them prepare for both good and bad events. Imagine, for instance, that an animal first smells a new food, eats it, and then gets sick. If the animal can learn to associate the smell (CS) with the food (US), then it will quickly learn that the food creates the negative outcome, and not eat it the next time.

The Persistence and Extinction of Conditioning

After he had demonstrated that learning could occur through association, Pavlov moved on to study the variables that influenced the strength and the persistence of conditioning. In some studies, after the conditioning had taken place, Pavlov presented the sound repeatedly but without presenting the food afterward. Figure 7.4 "Acquisition, Extinction, and Spontaneous Recovery"shows what happened. As you can see, after the intial acquisition (learning) phase in which the conditioning occurred, when the CS was then presented alone, the behavior rapidly decreased—the dogs salivated less and less to the sound, and eventually the sound did not elicit salivation at all. Extinction refers to the reduction in responding that occurs when the conditioned stimulus is presented repeatedly without the unconditioned stimulus.

Figure 7.4 Acquisition, Extinction, and Spontaneous Recovery

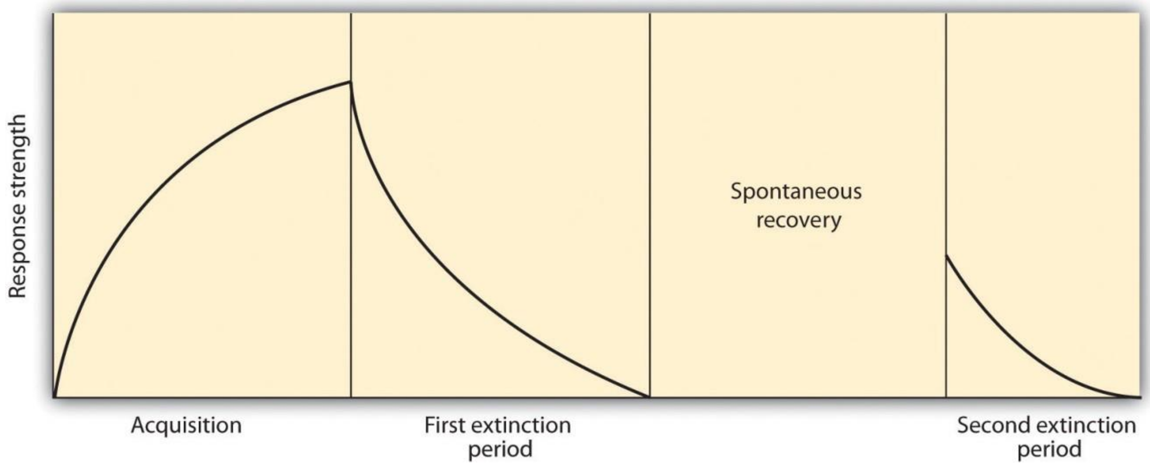

Acquisition: The CS and the US are repeatedly paired together and behavior increases. Extinction: The CS is repeatedly presented alone, and the behavior slowly decreases. Spontaneous recovery: After a pause, when the CS is again presented alone, the behavior may again occur and then again show extinction.

Although at the end of the first extinction period the CS was no longer producing salivation, the effects of conditioning had not entirely disappeared. Pavlov found that, after a pause, sounding the tone again elicited salivation, although to a lesser extent than before extinction took place. *The increase in responding to the CS following a pause after extinction* is known as spontaneous recovery. When Pavlov again presented the CS alone, the behavior again showed extinction until it disappeared again.

Although the behavior has disappeared, extinction is never complete. If conditioning is again attempted, the animal will learn the new associations much faster than it did the first time.

Pavlov also experimented with presenting new stimuli that were similar, but not identical to, the original conditioned stimulus. For instance, if the dog had been conditioned to being scratched before the food arrived, the stimulus would be changed to being rubbed rather than scratched. He found that the dogs also salivated upon experiencing the similar stimulus, a process known as *generalization*. Generalization refers to *the tendency to respond to stimuli that resemble the original conditioned stimulus*. The ability to generalize has important evolutionary significance. If we eat some red berries and they make us sick, it would be a good idea to think twice before we eat some purple berries. Although the berries are not exactly the same, they nevertheless are similar and may have the same negative properties.

The flip side of generalization is discrimination—*the tendency to respond differently to stimuli that are similar but not identical*. Pavlov's dogs quickly learned, for example, to salivate when they heard the specific tone that had preceded food, but not upon hearing similar tones that had never been associated with food. Discrimination is also useful—if we do try the purple berries, and if they do not make us sick, we will be able to make the distinction in the future. And we can learn that although the two people in our class, Courtney and Sarah, may look a lot alike, they are nevertheless different people with different personalities.

In some cases, *an existing conditioned stimulus can serve as an unconditioned stimulus for a pairing with a new conditioned stimulus*—a process known as second-order conditioning. In one of Pavlov's studies, for instance, he first conditioned the dogs to salivate to a sound, and then repeatedly paired a new CS, a black square, with the sound. Eventually he found that the dogs would salivate at the sight of the black square alone, even though it had never been directly associated with the food. Secondary conditioners in everyday life include our attractions to things that stand for or remind us of something else, such as when we feel good on a Friday because it has become associated with the paycheck that we receive on that day, which itself is a conditioned stimulus for the pleasures that the paycheck buys us.

The Role of Nature in Classical Conditioning

As we have seen in Chapter 1 "Introducing Psychology", scientists associated with the behavioralist school argued that all learning is driven by experience, and that nature plays no role. Classical conditioning, which is based on learning through experience, represents an example of the importance of the environment. But classical conditioning cannot be understood entirely in terms of experience. Nature also plays a part, as our evolutionary history has made us better able to learn some associations than others.

Clinical psychologists make use of classical conditioning to explain the learning of a phobia—*a strong and irrational fear of a specific object, activity, or situation*. For example, driving a car is a neutral event that would not normally elicit a fear response in most people. But if a person were to experience a panic attack in which he suddenly experienced strong negative emotions while driving, he may learn to associate driving with the panic response. The driving has become the CS that now creates the fear response.

Psychologists have also discovered that people do not develop phobias to just anything. Although people may in some cases develop a driving phobia, they are more likely to develop phobias toward objects (such as snakes, spiders, heights, and open spaces) that have been dangerous to people in the past. In modern life, it is rare for humans to be bitten by spiders or snakes, to fall from trees or buildings, or to be attacked by a predator in an open area. Being injured while riding in a car or being cut by a knife are much more likely. But in our evolutionary past, the potential of being bitten by snakes or spiders, falling out of a tree, or being trapped in an open space were important evolutionary concerns, and therefore humans are still evolutionarily prepared to learn these associations over others (Öhman & Mineka, 2001; LoBue & DeLoache, 2010).[1]

Another evolutionarily important type of conditioning is conditioning related to food. In his important research on food conditioning, John Garcia and his colleagues (Garcia, Kimeldorf, & Koelling, 1955; Garcia, Ervin, & Koelling, 1966)[2] attempted to condition rats by presenting either a taste, a sight, or a sound as a neutral stimulus before the rats were given drugs (the US) that made them nauseous. Garcia discovered that taste conditioning was extremely powerful—the rat learned to avoid the taste associated with illness, even if the illness occurred several hours later. But conditioning the behavioral response of nausea to a sight or a sound was much more difficult. These results contradicted the idea that conditioning occurs entirely as a result of environmental events, such that it would occur equally for any kind of unconditioned stimulus that followed any kind of conditioned stimulus. Rather, Garcia's research showed that genetics matters—organisms are evolutionarily prepared to learn some associations more easily than others. You can see that the ability to associate smells with illness is an important survival mechanism, allowing the organism to quickly learn to avoid foods that are poisonous.

[1] Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: Toward an evolved module of fear and fear learning. *Psychological Review, 108*(3), 483–522; LoBue, V., & DeLoache, J. S. (2010). Superior detection of threat-relevant stimuli in infancy.*Developmental Science, 13*(1), 221–228.

[2] Garcia, J., Kimeldorf, D. J., & Koelling, R. A. (1955). Conditioned aversion to saccharin resulting from exposure to gamma radiation. *Science, 122*, 157–158; Garcia, J., Ervin, F. R., & Koelling, R. A. (1966). Learning with prolonged delay of reinforcement. *Psychonomic Science, 5*(3), 121–122.

7.2 Changing Behavior Through Reinforcement and Punishment: Operant Conditioning

LEARNING OBJECTIVES

1. Outline the principles of operant conditioning.
2. Explain how learning can be shaped through the use of reinforcement schedules and secondary reinforcers.

In classical conditioning the organism learns to associate new stimuli with natural, biological responses such as salivation or fear. The organism does not learn something new but rather begins to perform in an existing behavior in the presence of a new signal. Operant conditioning, on the other hand, is *learning that occurs based on the consequences of behavior* and can involve the learning of new actions. Operant conditioning occurs when a dog rolls over on command because it has been praised for doing so in the past, when a schoolroom bully threatens his classmates because doing so allows him to get his way, and when a child gets good grades because her parents threaten to punish her if she doesn't. In operant conditioning the organism learns from the consequences of its own actions.

How Reinforcement and Punishment Influence Behavior: The Research of Thorndike and Skinner

Psychologist Edward L. Thorndike (1874–1949) was the first scientist to systematically study operant conditioning. In his research Thorndike (1898) [1]observed cats who had been placed in a “puzzle box” from which they tried to escape (Note 7.21 "Video Clip: Thorndike’s Puzzle Box"). At first the cats scratched, bit, and swatted haphazardly, without any idea of how to get out. But eventually, and accidentally, they pressed the lever that opened the door and exited to their prize, a scrap of fish. The next time the cat was constrained within the box it attempted fewer of the ineffective responses before carrying out the successful escape, and after several trials the cat learned to almost immediately make the correct response.

Observing these changes in the cats’ behavior led Thorndike to develop his law of effect, *the principle that responses that create a typically pleasant outcome in a particular situation are more likely to occur again in a similar situation, whereas responses that produce a typically unpleasant outcome are less likely to occur again in the situation* (Thorndike, 1911).[2] The essence of the law of effect is that successful responses, because they are pleasurable, are “stamped in” by experience and thus occur more frequently. Unsuccessful responses, which produce unpleasant experiences, are “stamped out” and subsequently occur less frequently.

Video Clip: Thorndike’s Puzzle Box

When Thorndike placed his cats in a puzzle box, he found that they learned to engage in the important escape behavior faster after each trial. Thorndike described the learning that follows reinforcement in terms of the law of effect.

The influential behavioral psychologist B. F. Skinner (1904–1990) expanded on Thorndike’s ideas to develop a more complete set of principles to explain operant conditioning. Skinner created specially designed environments known as *operant chambers* (usually called *Skinner boxes*) to systemically study learning. A Skinner box (operant chamber) is *a structure that is big enough to fit a rodent or bird and that contains a bar or key that the organism can press or peck to release food or water. It also contains a device to record the animal’s responses*.

The most basic of Skinner’s experiments was quite similar to Thorndike’s research with cats. A rat placed in the chamber reacted as one might expect, scurrying about the box and sniffing and clawing at the floor and walls. Eventually the rat chanced upon a lever, which it pressed to release pellets of food. The next time around, the rat took a little less time to press the lever, and on successive trials, the time it took to press the lever became shorter and shorter. Soon the rat was pressing the lever as fast as it could eat the food that appeared. As predicted by the law of effect, the rat had learned to repeat the action that brought about the food and cease the actions that did not.

Skinner studied, in detail, how animals changed their behavior through reinforcement and punishment, and he developed terms that explained the processes of operant learning (Table 7.1 "How Positive and Negative Reinforcement and Punishment Influence Behavior"). Skinner used the term reinforcer to refer to *any event that strengthens or increases the likelihood of a behavior* and the term punisher to refer to *any event that weakens or decreases the likelihood of a behavior*. And he used the terms *positive* and *negative* to refer to whether a reinforcement was presented or removed, respectively. Thus positive reinforcement *strengthens a response by presenting something pleasant after the response* and negative reinforcement *strengthens a response by reducing or removing something unpleasant*. For example, giving a child praise for completing his homework represents positive reinforcement, whereas taking aspirin to reduced the pain of a headache represents negative reinforcement. In both cases, the reinforcement makes it more likely that behavior will occur again in the future.

Table 7.1 How Positive and Negative Reinforcement and Punishment Influence Behavior

Operant conditioning term	**Description**	**Outcome**	**Example**
Positive reinforcement	Add or increase a pleasant stimulus	Behavior is strengthened	Giving a student a prize after he gets an A on a test
Negative reinforcement	Reduce or remove an unpleasant stimulus	Behavior is strengthened	Taking painkillers that eliminate pain increases the likelihood that you will take painkillers again
Positive punishment	Present or add an unpleasant stimulus	Behavior is weakened	Giving a student extra homework after she misbehaves in class
Negative punishment	Reduce or remove a pleasant stimulus	Behavior is weakened	Taking away a teen's computer after he misses curfew

Reinforcement, either positive or negative, works by increasing the likelihood of a behavior. Punishment, on the other hand, refers to *any event that weakens or reduces the likelihood of a* *behavior*. Positive punishment *weakens a response by presenting something unpleasant after the response*, whereas negative punishment *weakens a response by reducing or removing something pleasant*. A child who is grounded after fighting with a sibling (positive punishment) or who loses out on the opportunity to go to recess after getting a poor grade (negative punishment) is less likely to repeat these behaviors.

Although the distinction between reinforcement (which increases behavior) and punishment (which decreases it) is usually clear, in some cases it is difficult to determine whether a reinforcer is positive or negative. On a hot day a cool breeze could be seen as a positive reinforcer (because it brings in cool air) or a negative reinforcer (because it removes hot air). In other cases, reinforcement can be both positive and negative. One may smoke a cigarette both because it brings pleasure (positive reinforcement) and because it eliminates the craving for nicotine (negative reinforcement).

It is also important to note that reinforcement and punishment are not simply opposites. The use of positive reinforcement in changing behavior is almost always more effective than using punishment. This is because positive reinforcement makes the person or animal feel better, helping create a positive relationship with the person providing the reinforcement. Types of positive reinforcement that are effective in everyday life include verbal praise or approval, the awarding of status or prestige, and direct financial payment. Punishment, on the other hand, is more likely to create only temporary changes in behavior because it is based on coercion and typically creates a negative and adversarial relationship with the person providing the reinforcement. When the person who provides the punishment leaves the situation, the unwanted behavior is likely to return.

Creating Complex Behaviors Through Operant Conditioning

Perhaps you remember watching a movie or being at a show in which an animal—maybe a dog, a horse, or a dolphin—did some pretty amazing things. The trainer gave a command and the dolphin swam to the bottom of the pool, picked up a ring on its nose, jumped out of the water through a hoop in the air, dived again to the bottom of the pool, picked up another ring, and then took both of the rings to the trainer at the edge of the pool. The animal was trained to do the trick, and the principles of operant conditioning were used to train it. But these complex behaviors are a far cry from the simple stimulus-response relationships that we have considered thus far. How can reinforcement be used to create complex behaviors such as these?

One way to expand the use of operant learning is to modify the schedule on which the reinforcement is applied. To this point we have only discussed a continuous reinforcement schedule, in which *the desired response is reinforced every time it occurs*; whenever the dog rolls over, for instance, it gets a biscuit. Continuous reinforcement results in relatively fast learning but also rapid extinction of the desired behavior once the reinforcer disappears. The problem is that because the organism is used to receiving the reinforcement after every behavior, the responder may give up quickly when it doesn't appear.

Most real-world reinforcers are not continuous; they occur on a partial (or intermittent) reinforcement schedule—*a schedule in which the responses are sometimes reinforced, and sometimes not*. In comparison to continuous reinforcement, partial reinforcement schedules lead to slower initial learning, but they also lead to greater resistance to extinction. Because the reinforcement does not appear after every behavior, it takes longer for the learner to determine that the reward is no longer coming, and thus extinction is slower. The four types of partial reinforcement schedules are summarized in Table 7.2 "Reinforcement Schedules".

Table 7.2 Reinforcement Schedules

Reinforcement schedule	**Explanation**	**Real-world example**
Fixed-ratio	Behavior is reinforced after a specific number of responses	Factory workers who are paid according to the number of products they produce
Variable-ratio	Behavior is reinforced after an average, but unpredictable, number of responses	Payoffs from slot machines and other games of chance
Fixed-interval	Behavior is reinforced for the first response after a specific amount of time has passed	People who earn a monthly salary
Variable-interval	Behavior is reinforced for the first response after an average, but unpredictable, amount of time has passed	Person who checks voice mail for messages

Partial reinforcement schedules are determined by whether the reinforcement is presented on the basis of the time that elapses between reinforcement (interval) or on the basis of the number of responses that the organism engages in (ratio), and by whether the reinforcement occurs on a regular (fixed) or unpredictable (variable) schedule. In a fixed-interval schedule,*reinforcement occurs for the first response made after a specific amount of time has passed*. For instance, on a one-minute fixed-interval schedule the animal receives a reinforcement every minute, assuming it engages in the behavior at least once during the minute. As you can see in Figure 7.7 "Examples of Response Patterns by Animals Trained Under Different Partial Reinforcement Schedules", animals under fixed-interval schedules tend to slow down their responding immediately after the reinforcement but then increase the behavior again as the time of the next reinforcement gets closer. (Most students study for exams the same way.) In a variable-interval schedule,*the reinforcers appear on an interval schedule, but the timing is varied around the average interval, making the actual appearance of the reinforcer unpredictable*. An example might be checking your e-mail: You are reinforced by receiving messages that come, on average, say every 30 minutes, but the reinforcement occurs only at random times. Interval reinforcement schedules tend to produce slow and steady rates of responding.

Figure 7.7 Examples of Response Patterns by Animals Trained Under Different Partial Reinforcement Schedules

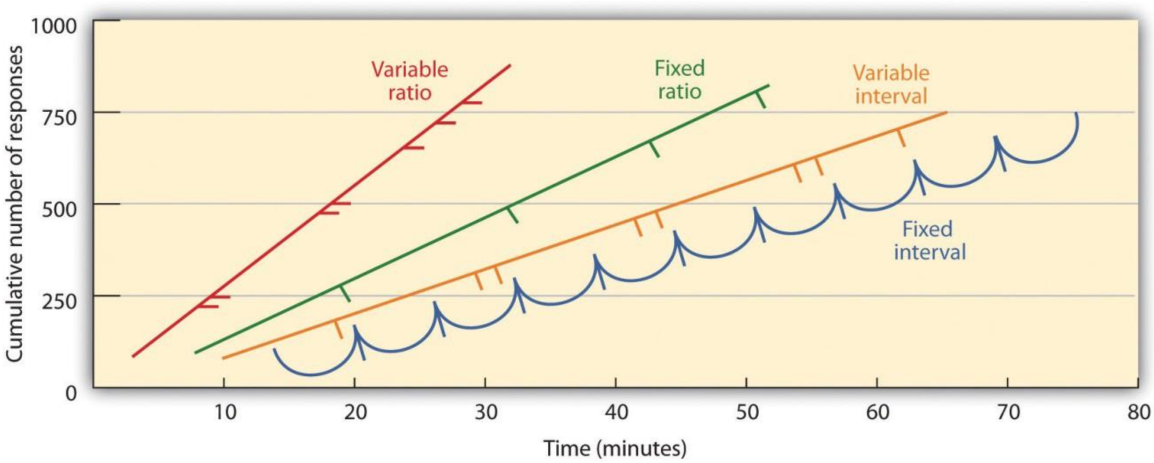

Schedules based on the number of responses (ratio types) induce greater response rate than do schedules based on elapsed time (interval types). Also, unpredictable schedules (variable types) produce stronger responses than do predictable schedules (fixed types).

Source: Adapted from Kassin, S. (2003). Essentials of psychology. Upper Saddle River, NJ: Prentice Hall. Retrieved from Essentials of PsychologyPrentice Hall Companion Website:http://wps.prenhall.com/hss_kassin_essentials_1/15/3933/1006917.cw/index.html.

In a fixed-ratio schedule, *a behavior is reinforced after a specific number of responses*. For instance, a rat's behavior may be reinforced after it has pressed a key 20 times, or a salesperson may receive a bonus after she has sold 10 products. As you can see in Figure 7.7 "Examples of Response Patterns by Animals Trained Under Different Partial Reinforcement Schedules", once the organism has learned to act in accordance with the fixed-reinforcement schedule, it will pause only briefly when reinforcement occurs before returning to a high level of responsiveness. A variable-ratio schedule *provides reinforcers after a specific but average number of responses*. Winning money from slot machines or on a lottery ticket are examples of reinforcement that occur on a variable-ratio schedule. For instance, a slot machine may be programmed to provide a win every 20 times the user pulls the handle, on average. As you can see in Figure 7.8 "Slot Machine", ratio schedules tend to produce high rates of responding because reinforcement increases as the number of responses increase.

Complex behaviors are also created through shaping, *the process of guiding an organism's behavior to the desired outcome through the use of successive approximation to a final desired behavior*. Skinner made extensive use of this procedure in his boxes. For instance, he could train a rat to press a bar two times to receive food, by first providing food when the animal moved near the bar. Then when that behavior had been learned he would begin to provide food only when the rat touched the bar. Further shaping limited the reinforcement to only when the rat pressed the bar, to when it pressed the bar and touched it a second time, and finally, to only when it pressed the bar twice. Although it can take a long time, in this way operant conditioning can create chains of behaviors that are reinforced only when they are completed.

Reinforcing animals if they correctly discriminate between similar stimuli allows scientists to test the animals' ability to learn, and the discriminations that they can make are sometimes quite remarkable. Pigeons have been trained to distinguish between images of Charlie Brown and the other Peanuts characters (Cerella, 1980), [3] and between different styles of music and art (Porter & Neuringer, 1984; Watanabe, Sakamoto & Wakita, 1995). [4]

Behaviors can also be trained through the use of *secondary reinforcers*. Whereas a primary reinforcer includes *stimuli that are naturally preferred or enjoyed by the organism, such as food, water, and relief from pain*, a secondary reinforcer (sometimes called *conditioned reinforcer*) is a *neutral event that has become associated with a primary reinforcer through classical conditioning*. An example of a secondary reinforcer would be the whistle given by an animal trainer, which has been associated over time with the primary reinforcer, food. An example of an everyday secondary reinforcer is money. We enjoy having money, not so much for the stimulus itself, but rather for the primary reinforcers (the things that money can buy) with which it is associated.

[1] Thorndike, E. L. (1898). *Animal intelligence: An experimental study of the associative processes in animals.* Washington, DC: American Psychological Association.

[2] Thorndike, E. L. (1911). *Animal intelligence: Experimental studies.* New York, NY: Macmillan. Retrieved from http://www.archive.org/details/animalintelligen00thor

[3] Cerella, J. (1980). The pigeon's analysis of pictures. *Pattern Recognition, 12*, 1–6.

[4] Porter, D., & Neuringer, A. (1984). Music discriminations by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 10*(2), 138–148; Watanabe, S., Sakamoto, J., & Wakita, M. (1995). Pigeons' discrimination of painting by Monet and Picasso. *Journal of the Experimental Analysis of Behavior, 63*(2), 165–174.

7.3 Learning by Insight and Observation

John B. Watson and B. F. Skinner were behaviorists who believed that all learning could be explained by the processes of conditioning—that is, that associations, and associations alone, influence learning. But some kinds of learning are very difficult to explain using only conditioning. Thus, although classical and operant conditioning play a key role in learning, they constitute only a part of the total picture.

One type of learning that is not determined only by conditioning occurs when we suddenly find the solution to a problem, as if the idea just popped into our head. This type of learning is known as insight, *the sudden understanding of a solution to a problem*. The German psychologist Wolfgang Köhler (1925) [1]carefully observed what happened when he presented chimpanzees with a problem that was not easy for them to solve, such as placing food in an area that was too high in the cage to be reached. He found that the chimps first engaged in trial-and-error attempts at solving the problem, but when these failed they seemed to stop and contemplate for a while. Then, after this period of contemplation, they would suddenly seem to know how to solve the problem, for instance by using a stick to knock the food down or by standing on a chair to reach it. Köhler argued that it was this flash of insight, not the prior trial-and-error approaches, which were so important for conditioning theories, that allowed the animals to solve the problem.

Edward Tolman (Tolman & Honzik, 1930) [2] studied the behavior of three groups of rats that were learning to navigate through mazes. The first group always received a reward of food at the end of the maze. The second group never received any reward, and the third group received a reward, but only beginning on the 11th day of the experimental period. As you might expect when considering the principles of conditioning, the rats in the first group quickly learned to negotiate the maze, while the rats of the second group seemed to wander aimlessly through it. The rats in the third group, however, although they wandered aimlessly for the first 10 days, quickly learned to navigate to the end of the maze as soon as they received food on day 11. By the next day, the rats in the third group had caught up in their learning to the rats that had been rewarded from the beginning.

It was clear to Tolman that the rats that had been allowed to experience the maze, even without any reinforcement, had nevertheless learned something, and Tolman called this *latent learning*. Latent learning refers to *learning that is not reinforced and not demonstrated until there is motivation to do so*. Tolman argued that the rats had formed a "cognitive map" of the maze but did not demonstrate this knowledge until they received reinforcement.

Observational Learning: Learning by Watching

The idea of latent learning suggests that animals, and people, may learn simply by experiencing or watching. Observational learning (modeling) is*learning by observing the behavior of others*. To demonstrate the importance of observational learning in children, Bandura, Ross, and Ross

(1963) [3]showed children a live image of either a man or a woman interacting with a Bobo doll, a filmed version of the same events, or a cartoon version of the events. The Bobo doll is an inflatable balloon with a weight in the bottom that makes it bob back up when you knock it down. In all three conditions, the model violently punched the clown, kicked the doll, sat on it, and hit it with a hammer.

The researchers first let the children view one of the three types of modeling, and then let them play in a room in which there were some really fun toys. To create some frustration in the children, Bandura let the children play with the fun toys for only a couple of minutes before taking them away. Then Bandura gave the children a chance to play with the Bobo doll.

If you guessed that most of the children imitated the model, you would be correct. Regardless of which type of modeling the children had seen, and regardless of the sex of the model or the child, the children who had seen the model behaved aggressively—just as the model had done. They also punched, kicked, sat on the doll, and hit it with the hammer. Bandura and his colleagues had demonstrated that these children had learned new behaviors, simply by observing and imitating others.

Observational learning is useful for animals and for people because it allows us to learn without having to actually engage in what might be a risky behavior. Monkeys that see other monkeys respond with fear to the sight of a snake learn to fear the snake themselves, even if they have been raised in a laboratory and have never actually seen a snake (Cook & Mineka, 1990). [4] As Bandura put it,

the prospects for [human] survival would be slim indeed if one could learn only by suffering the consequences of trial and error. For this reason, one does not teach children to swim, adolescents to drive automobiles, and novice medical students to perform surgery by having them discover the appropriate behavior through the consequences of their successes and *failures. The more costly and hazardous the possible mistakes, the heavier is the reliance on observational learning from competent learners. (Bandura, 1977, p. 212)* [5]

Although modeling is normally adaptive, it can be problematic for children who grow up in violent families. These children are not only the victims of aggression, but they also see it happening to their parents and siblings. Because children learn how to be parents in large part by modeling the actions of their own parents, it is no surprise that there is a strong correlation between family violence in childhood and violence as an adult. Children who witness their parents being violent or who are themselves abused are more likely as adults to inflict abuse on intimate partners or their children, and to be victims of intimate violence (Heyman & Slep, 2002). [6] In turn, their children are more likely to interact violently with each other and to aggress against their parents (Patterson, Dishion, & Bank, 1984). [7]

[1] Köhler, W. (1925). *The mentality of apes* (E. Winter, Trans.). New York, NY: Harcourt Brace Jovanovich.

[2] Tolman, E. C., & Honzik, C. H. (1930). Introduction and removal of reward, and maze performance in rats. *University of California Publications in Psychology, 4*, 257–275.

[3] Bandura, A., Ross, D., & Ross, S. A. (1963). Imitation of film-mediated aggressive models. *The Journal of Abnormal and Social Psychology, 66*(1), 3–11.

[4] Cook, M., & Mineka, S. (1990). Selective associations in the observational conditioning of fear in rhesus monkeys. *Journal of Experimental Psychology: Animal Behavior Processes, 16*(4), 372–389.

[5] Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavior change.*Psychological Review, 84*, 191–215.

[6] Heyman, R. E., & Slep, A. M. S. (2002). Do child abuse and interparental violence lead to adulthood family violence? *Journal of Marriage and Family, 64*(4), 864–870.

[7] Patterson, G. R., Dishion, T. J., & Bank, L. (1984). Family interaction: A process model of deviancy training. *Aggressive Behavior, 10*(3), 253–267.