Print This Page

Module 6.3: Operant Conditioning

The previous section of this lesson focused on the type of associative learning known as classical conditioning. Remember that in classical conditioning, something in the environment triggers a reflex automatically, and researchers train the organism to react to a different stimulus. Now we turn to the second type of associative learning, **operant conditioning**. In operant conditioning, organisms learn to associate a behavior and its consequence (Table 6.2). A pleasant consequence makes that behavior more likely to be repeated in the future.

Table 6.2 Classical and Operant Conditioning Compared

Factors	Classical Conditioning	Operant Conditioning
Conditioning approach	An unconditioned stimulus (such as food) is paired with a neutral stimulus (such as a bell). The neutral stimulus eventually becomes the conditioned stimulus, which brings about the conditioned response (salivation).	The target behavior is followed by reinforcement or punishment to either strengthen or weaken it, so that the learner is more likely to exhibit the desired behavior in the future.
Stimulus timing	The stimulus occurs immediately before the response.	The stimulus (either reinforcement or punishment) occurs soon after the response.

Psychologist B. F. Skinner saw that classical conditioning is limited to existing behaviors that are reflexively elicited, and it doesn't account for new behaviors such as riding a bike. He proposed a theory about how such new behaviors come about. Skinner believed that behavior is motivated by the consequences we receive for the behavior: the reinforcements and punishments. His idea that learning is the result of consequences is based on the law of effect, which was first proposed by psychologist Edward Thorndike.

According to the **law of effect**, behaviors that are followed by consequences that are satisfying to the organism are more likely to be repeated, and behaviors that are followed by unpleasant consequences are less likely to be repeated. Essentially, if an organism does something that brings about a desired result, the organism is more likely to do it again. If an organism does something that does not bring about a desired result, the organism is less likely to do it again. An example of the law of effect is in employment. One of the reasons (and often the main reason) we show up for work is because we get paid to do so. If we stop getting paid, we will likely stop showing up—even if we love our job.

Link to Learning

Working with Thorndike's law of effect as his foundation, Skinner began conducting scientific experiments on animals (mainly rats and pigeons) to determine how organisms learn through operant conditioning (Skinner, 1938). He placed these animals inside an operant conditioning chamber, which has come to be known as a "Skinner box" (**Figure 6.9**). A Skinner box contains a lever (for rats) or disk (for pigeons) that the animal can press or peck for a food reward via the dispenser. Speakers and lights can be associated with certain behaviors. A recorder counts the number of responses made by the animal.

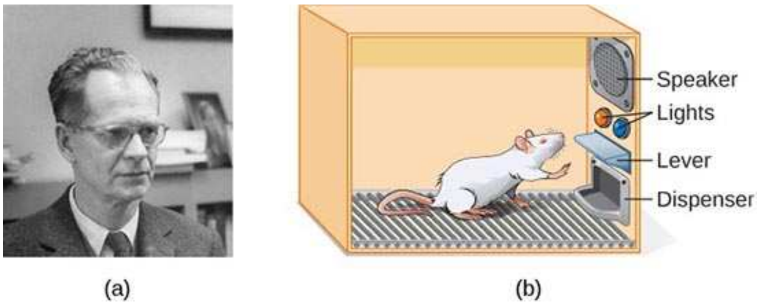

Figure 6.9 (a) B. F. Skinner developed operant conditioning for systematic study of how behaviors are strengthened or weakened according to their consequences. (b) In a Skinner box, a rat presses a lever in an operant conditioning chamber to receive a food reward. (credit a: modification of work by "Silly rabbit"/Wikimedia Commons)

Watch the brief video clip below to learn more about operant conditioning: Skinner is interviewed, and operant conditioning of pigeons is demonstrated.

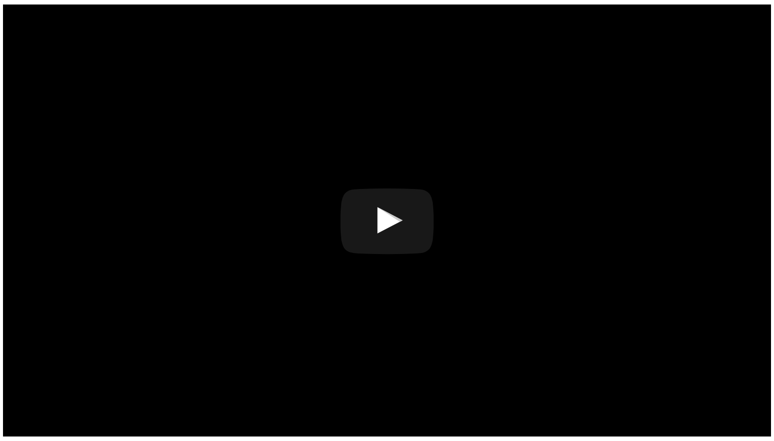

In discussing operant conditioning, we use several everyday words—positive, negative, reinforcement, and punishment—in a specialized manner. In operant conditioning, positive and negative do not mean good and bad. Instead, positive means you are adding something, and negative means you are taking something away. Reinforcementmeans you are increasing a behavior, and *punishment* means you are decreasing a behavior.

Reinforcement can be positive or negative, and punishment can also be positive or negative. All reinforcers (positive or negative) increase the likelihood of a behavioral response. All punishers (positive or negative) decrease the likelihood of a behavioral response.

Explore how these four terms are combined in operant conditioning by clicking on each of the images below. Before clicking, see if you can guess which term each image represents and why. Click the "Reset" button to clear the text and start over.

Access a printable copy of this activity

RESET

Dig Deeper

Physical Punishment

Punishment, especially when it is immediate, is one way to decrease undesirable behavior. For example, imagine your four-year-old son, Brandon, runs into the busy street to get his ball. You give him a time- out (positive punishment) and tell him never to go into the street again. Chances are he won't repeat this behavior. While strategies like time-outs are common today, in the past children were often subject to physical punishment, such as spanking. It's important to be aware of some of the drawbacks in using physical punishment on children.

First, punishment may teach fear. Brandon may become fearful of the street, but he also may become fearful of the person who delivered the punishment—you, his parent. Similarly, children who are punished by teachers may come to fear the teacher and try to avoid school. Consequently, most schools in the United States have banned corporal punishment.

Second, punishment may cause children to become more aggressive and prone to antisocial behavior and delinquency. They see their parents resort to spanking when they become angry and frustrated, so, in turn, they may act out this same behavior when they become angry and frustrated. For example, because you spank Brenda when you are angry with her for her misbehavior, she might start hitting her friends when they won't share their toys.

While positive punishment can be effective in some cases, Skinner suggested that the use of punishment should be weighed against the possible negative effects. Today's psychologists and parenting experts favor reinforcement over punishment—they recommend that you catch your child doing something good and reward her for it.

Shaping

In his operant conditioning experiments, Skinner often used an approach called shaping. Instead of rewarding only the target behavior, in **shaping**, we reward successive approximations of a target behavior. Why is shaping needed? Remember that in order for reinforcement to work, the organism must first display the behavior. Shaping is needed because it is extremely unlikely that an organism will display anything but the simplest of behaviors spontaneously. In shaping, behaviors are broken down into many small, achievable steps. The specific steps used in the process are the following:

1. Reinforce any response that resembles the desired behavior.
2. Then reinforce the response that more closely resembles the desired behavior. You will no longer reinforce the previously reinforced response.
3. Next, begin to reinforce the response that even more closely resembles the desired behavior.
4. Continue to reinforce closer and closer approximations of the desired behavior.
5. Finally, only reinforce the desired behavior.

Connect the Concepts

Shaping is often used in teaching a complex behavior or chain of behaviors. Skinner used shaping to teach pigeons not only such relatively simple behaviors as pecking a disk in a Skinner box, but also many unusual and entertaining behaviors, such as turning in circles, walking in figure eights, and even playing ping pong; the technique is commonly used by animal trainers today. Watch the brief video below of Skinner's pigeons playing ping pong.

An important part of shaping is stimulus discrimination. Recall Pavlov's dogs—he trained them to respond to the tone of a bell, and not to similar tones or sounds. This discrimination is also important in operant conditioning and in shaping behavior.

It's easy to see how shaping is effective in teaching behaviors to animals, but how does shaping work with humans? Let's consider parents whose goal is to have their child learn to clean his room. They use shaping to help him master steps toward the goal. Instead of performing the entire task, they set up these steps and reinforce each step. First, he cleans up one toy. Second, he cleans up five toys. Third, he chooses whether to pick up ten toys or put his books and clothes away. Fourth, he cleans up everything except two toys. Finally, he cleans his entire room.

Primary and Secondary Reinforcers

Rewards such as stickers, praise, money, toys, and more can be used to reinforce learning. Let's go back to Skinner's rats again. How did the rats learn to press the lever in the Skinner box? They were rewarded with food each time they pressed the lever. For animals, food would be an obvious reinforcer.

What would be a good reinforce for humans? If you gave Joaquin a piece of candy every time he made a goal, you would be using a primary reinforcer. **Primary reinforcers** are reinforcers that have innate reinforcing qualities. These kinds of reinforcers are not learned. Water, food, sleep, shelter, sex, and touch, among others, are primary reinforcers. Pleasure is also a primary reinforcer. Organisms do not lose their drive for these things.

A **secondary reinforcer** has no inherent value and only has reinforcing qualities when linked with a primary reinforcer. Praise, linked to affection, is one example of a secondary reinforcer, as when you called out "Great shot!" every time Joaquin made a goal. Another example, money, is only worth something when you can use it to buy other things—either things that satisfy basic needs (food, water, shelter—all primary reinforcers) or other secondary reinforcers. If you were on a remote island in the middle of the Pacific Ocean and you had stacks of money, the money would not be useful if you could not spend it.

What about the stickers on a behavior chart? They also are secondary reinforcers. Sometimes, instead of stickers on a sticker chart, a token is used. Tokens, which are also secondary reinforcers, can then be traded in for rewards and prizes. Entire behavior management systems, known as token economies, are built around the use of these kinds of token reinforcers. Token economies have been found to be very effective at modifying behavior in a variety of settings such as schools, prisons, and mental hospitals.

Everyday Connection

Behavior Modification in Children

Parents and teachers often use behavior modification to change a child's behavior. Behavior modification uses the principles of operant conditioning to accomplish behavior change so that undesirable behaviors are switched for more socially acceptable ones. Some teachers and parents create a sticker chart, in which several behaviors are listed (**Figure 6.10**). Sticker charts are a form of token economies, as described above. Each time children perform the behavior, they get a sticker, and after a certain number of stickers, they get a prize, or reinforcer. The goal is to increase acceptable behaviors and decrease misbehavior.

Figure 6.10 Sticker charts are a form of positive

Remember, it is best to reinforce desired behaviors, rather than to use punishment. In the classroom, the teacher can reinforce a range of behaviors, from students raising their hands, to walking quietly in the hall, to turning in their homework. At home, parents might create a behavior chart that rewards children for things such as putting away toys, brushing their teeth, and helping with dinner. In order for behavior modification to be effective, the reinforcement needs to be connected with the behavior; the reinforcement must matter to the child and be done consistently.

reinforcement and a tool for behavior modification. Once this little girl earns a certain number of stickers for demonstrating a desired behavior, she will be rewarded with a trip to the ice cream parlor. (credit: Abigail Batchelder)

Time-out is another popular technique used in behavior modification with children. It operates on the principle of negative punishment. When a child demonstrates an undesirable behavior, she is removed from the desirable activity at hand (**Figure 6.11**).

Figure 6.11 Time-out is a popular form of negative punishment used by caregivers. When a child misbehaves, he or she is removed from a desirable activity in an effort to decrease the unwanted behavior. For example, (a) a child might be playing on the playground with friends and push another child; (b) the child who misbehaved would then be removed from the activity for a short period of time. (credit a: modification of work by Simone Ramella; credit b: modification of work by "JefferyTurner"/Flickr)

There are several important points that you should know if you plan to implement time-out as a behavior modification technique:

- Make sure the child is being removed from a desirable activity and placed in a less desirable location. If the activity is something undesirable for the child, this technique will backfire because it is more enjoyable for the child to be removed from the activity.
- The length of the time-out is important. The general rule of thumb is one minute for each year of the child's age. Sophia is five; therefore, she sits in a time-out for five minutes. Setting a timer helps children know how long they have

to sit in time-out.
- As a caregiver, keep several guidelines in mind over the course of a time-out: remain calm when directing your child to time-out; ignore your child during time-out (because caregiver attention may reinforce misbehavior); and give the child a hug or a kind word when time-out is over.

Reinforcement Schedules

Remember, the best way to teach a person or animal a behavior is to use positive reinforcement. When an organism receives a reinforcer each time it displays a behavior, it is called **continuous reinforcement**. This reinforcement schedule is the quickest way to teach someone a behavior, and it is especially effective in training a new behavior. When teaching your dog to sit, you might give him a treat each time he sits on command. Timing is important here: you will be most successful if you present the reinforcer immediately after he sits, so that he can make an association between the target behavior (sitting) and the consequence (getting a treat).

Link to Learning

Watch this video clip where veterinarian Dr. Sophia Yin shapes a dog's behavior using the steps outlined above.

Once a behavior is trained, researchers and trainers often turn to another type of reinforcement schedule—partial reinforcement. In **partial reinforcement**, also referred to as intermittent reinforcement, the person or animal does not get reinforced every time they perform the desired behavior. There are several different types of partial reinforcement schedules (**Table 6.3**). These schedules are described as either fixed or variable, and as either interval or ratio. Fixed refers to the number of responses between reinforcements, or the amount of time between reinforcements, which is set and unchanging. Variable refers to the number of responses or amount of time between reinforcements, which varies or changes. Interval means the schedule is based on the time between reinforcements, and ratio means the schedule is based on the number of responses between reinforcements.

Table 6.3 Reinforcement Schedules

Reinforcement Schedule	Description	Result	Example

Fixed interval	Reinforcement is delivered at predictable time intervals (e.g., after 5, 10, 15, and 20 minutes).	Moderate response rate with significant pauses after reinforcement	Hospital patient uses patient-controlled, doctor-timed pain relief
Variable interval	Reinforcement is delivered at unpredictable time intervals (e.g., after 5, 7, 10, and 20 minutes).	Moderate yet steady response rate	Checking Facebook
Fixed ratio	Reinforcement is delivered after a predictable number of responses (e.g., after 2, 4, 6, and 8 responses).	High response rate with pauses after reinforcement	Piecework—factory worker getting paid for every x number of items manufactured
Variable ratio	Reinforcement is delivered after an unpredictable number of responses (e.g., after 1, 4, 5, and 9 responses).	High and steady response rate	Gambling

Now let's combine these four terms:

A **fixed interval reinforcement schedule** is when behavior is rewarded after a set amount of time. A patient recovering from surgery in a hospital is expected to experience pain and will require prescription medications for pain relief. She is given an IV drip with a patient-controlled painkiller that can only be used once per hour. She pushes a button when pain becomes difficult, and she receives a dose of medication. Since the reward (pain relief) only occurs on a fixed interval, there is no point in exhibiting the behavior when it will not be rewarded.

With a **variable interval reinforcement schedule**, the person or animal gets the reinforcement based on varying amounts of time, which are unpredictable. In some fast food chains, every once in a while, someone from the quality control division will visit a restaurant and everyone on the shift earns a $20 bonus if the restaurant is clean and the service is fast. Because the restaurant manager never knows when the quality control person will show up, he makes sure that things are clean and that his employees provide prompt and courteous service. His productivity regarding prompt service and keeping a clean restaurant are steady because he wants his crew to earn the bonus.

With a **fixed ratio reinforcement schedule**, there are a set number of responses that must occur before the behavior is rewarded. A salesperson may earn a commission every time she sells a pair of glasses. She always tries to sell people extra pairs of glasses so she can increase her commission. She does not care if the extra classes are necessary, because she just wants the bonus. The quality of what she sells does not matter because her commission is only based on the number of pairs sold.

This distinction in the quality of performance can help determine which reinforcement method is most appropriate for a particular situation. Fixed ratios are better suited to optimize the quantity of output, whereas a fixed interval, in which the reward is not quantity based, can lead to a higher quality of output.

In a **variable ratio reinforcement schedule**, the number of responses needed for a reward varies. This is the most powerful partial reinforcement schedule. An example of the variable ratio reinforcement schedule is gambling. Imagine that a woman who is generally smart and thrifty visits Las Vegas for the first time. She is not a gambler, but out of curiosity she puts a quarter into the slot machine, and then another, and another. Nothing happens. Two dollars in quarters later, her curiosity is fading, and she is just about to quit. But then, the machine lights up, bells go off, and she gets 50 quarters back. She starts putting quarters into the machine with renewed interest, and a few minutes later she has used up all her gains. Now might be a sensible time to quit, but she keeps putting money into the slot machine because she never knows when the next reinforcement is coming. She keeps thinking that with the next quarter she could win $50, or $100, or even more. Because the reinforcement schedule in most types of gambling has a variable ratio schedule, people keep trying and hoping that the next time they will win big. This is one of the reasons that gambling is so addictive—and so resistant to extinction.

In operant conditioning, extinction of a reinforced behavior occurs at some point after reinforcement stops, and the speed at which this happens depends on the reinforcement schedule. In a variable ratio schedule, the point of extinction comes very slowly, as described above. But in the other reinforcement schedules, extinction may come quickly, as in the example of the patient in the hospital. If she presses the button for the pain relief medication before the allotted time her doctor has approved, no medication is administered. She is on a fixed interval reinforcement schedule (dosed hourly), so extinction occurs quickly when reinforcement doesn't come at the expected time. Among the reinforcement schedules, variable ratio is the most productive and the most resistant to extinction. Fixed interval is the least productive and the easiest to extinguish.

Connect the Concepts

Gambling and the Brain

Skinner (1953) stated, "If the gambling establishment cannot persuade a patron to turn over money with no return, it may achieve the same effect by returning part of the patron's money on a variable-ratio schedule" (p. 397). Skinner uses gambling as an example of the power and effectiveness of conditioning behavior based on a variable ratio reinforcement schedule. In fact, Skinner was so confident in his knowledge of gambling addiction that he even claimed he could turn a pigeon into a pathological gambler. Beyond the power of variable ratio reinforcement, gambling seems to work on the brain in the same way as some addictive drugs. The Illinois Institute for Addiction Recovery reports evidence suggesting that pathological gambling is an addiction similar to a chemical addiction (Figure 6.12).

Figure 6.12 Some research suggests that pathological gamblers use gambling to compensate for abnormally low levels of the hormone norepinephrine, which is associated with stress and is secreted in moments of arousal and thrill. (credit: Ted Murphy)

Specifically, gambling may activate the reward centers of the brain, much like cocaine does. Research has shown that some pathological gamblers have lower levels of the neurotransmitter (brain chemical) known as norepinephrine than do normal gamblers. According to a study conducted by Alec Roy and colleagues, norepinephrine is secreted when a person feels stress, arousal, or thrill; pathological gamblers use gambling to increase their levels of this neurotransmitter. Another researcher, neuroscientist Hans Breiter, has done extensive research on gambling and its effects on the brain. Breiter reports that "Monetary reward in a gambling-like experiment produces brain activation very similar to that observed in a cocaine addict receiving an infusion of cocaine" (para. 1). Deficiencies in serotonin (another neurotransmitter) might also contribute to compulsive behavior, including a gambling addiction.

It may be that pathological gamblers' brains are different than those of other people, and perhaps this difference may somehow have led to their gambling addiction, as these studies seem to suggest. However, it is very difficult to ascertain the cause because it is impossible to conduct a true experiment (it would be unethical to try to turn randomly assigned participants into problem gamblers). Therefore, it may be that causation actually moves in the opposite direction—perhaps the act of gambling somehow changes neurotransmitter levels in some gamblers' brains. It also is possible that some overlooked factor, or confounding variable, played a role in both the gambling addiction and the differences in brain chemistry.

Cognition and Latent Learning

Although strict behaviorists such as Skinner and Watson refused to believe that cognition (such as thoughts and expectations) plays a role in learning, another behaviorist, Edward C. Tolman, had a different opinion. Tolman's experiments with rats demonstrated that organisms can learn even if they do not receive immediate reinforcement. This finding was in conflict with the prevailing idea at the time that reinforcement must be immediate in order for learning to occur, thus suggesting a cognitive aspect to learning.

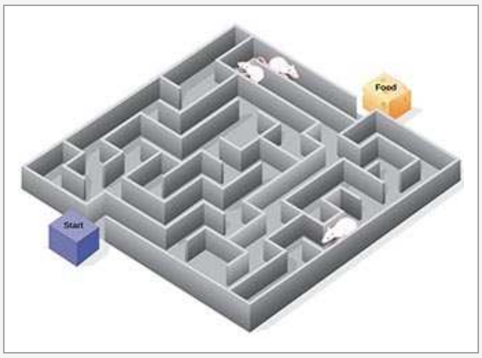

Figure 6.13 Psychologist Edward Tolman found that rats use cognitive maps to navigate through a maze. Have you ever worked your way through various levels on a video game? You learned when to turn left or right, move up or down. In that case you were relying on a cognitive map, just like the rats in a maze. (credit: modification of work by "FutUndBeidl"/Flickr)

In the experiments, Tolman placed hungry rats in a maze with no reward for finding their way through it. He also studied a comparison group that was rewarded with food at the end of the maze. As the unreinforced rats explored the maze, they developed a **cognitive map**: a mental picture of the layout of the maze (Figure 6.13). After 10 sessions in the maze without reinforcement, food was placed in a goal box at the end of the maze. As soon as the rats became aware of the food, they were able to find their way through the maze quickly, just as quickly as the comparison group, which had been rewarded with food all along. This is known as **latent learning**: learning that occurs but is not observable in behavior until there is a reason to demonstrate it.

Latent learning also occurs in humans. Children may learn by watching the actions of their parents but only demonstrate it at a later date, when the learned material is needed. For example, suppose that Ravi's dad drives him to school every day. In this way, Ravi learns the route from his house to his school, but he's never driven there himself, so he has not had a chance to demonstrate that he's learned the way. One morning Ravi's dad has to leave early for a meeting, so he can't drive Ravi to school. Instead, Ravi follows the same route on his bike that his dad would have taken in the car. This demonstrates latent learning. Ravi had learned the route to school, but had no need to demonstrate this knowledge earlier.

Everyday Connection

This Place Is Like a Maze

Have you ever gotten lost in a building and couldn't find your way back out? While that can be frustrating, you're not alone. At one time or another we've all gotten lost in places like a museum, hospital, or university library. Whenever we go someplace new, we build a mental representation—or cognitive map—of the location, as Tolman's rats built a cognitive map of their maze. However, some buildings are confusing because they include many areas that look alike or have short lines of sight. Because of this, it's often difficult to predict what's around a corner or decide whether to turn left or right to get out of a building.

Psychologist Laura Carlson (2010) suggests that what we place in our cognitive map can impact our success in navigating through the environment. She suggests that paying attention to specific features upon entering a building, such as a picture on the wall, a fountain, a statue, or an escalator, adds information to our cognitive map that can be used later to help find our way out of the building.

Watch this video to learn more about Carlson's studies on cognitive maps and navigation in buildings.

Check Your Knowledge

As you read your assignment for this lesson, pay close attention to the key terms and phrases listed throughout the chapter. These terms and concepts are important to your understanding of the information provided in the lesson.

Module 6.3 Flashcards

Click on the notecard to reveal the definition.

RESET

Operant conditioning

◀ Previous Card | Term 1 of 11 | Next Card ▶

←Module 6.2 Module 6.4→