

Chapter: Learning

Learning

Defining Learning

Elements of learning

Reinforcement

Classical Conditioning

Reinforcement in Classical Conditioning

USING PSYCHOLOGY: Breaking Bad Habits

Operant (Instrumental) Conditioning: A

Predecessor

Operant Conditioning: The Process and
EquipmentReinforcement and Punishment in Operant
Conditioning

USING PSYCHOLOGY: Breaking Bad Operants

Classical Versus Operant Conditioning

Observational Learning

Schedules of Reinforcement and Learning

Primary and secondary reinforcement

Timing of Reinforcements

Extinction

Generalization

Discrimination

Punishment

Factors which Influence Punishment

Punishment: Cues and Alternatives

REVIEW QUESTIONS

ACTIVITIES

INTERESTED IN MORE?

Learning

WHAT'S THE ANSWER?

"I was in fourth grade the first time I met up with that 1-inch rope dangling from the rafters of the school gym. The PE teacher shinnied right to the top and then shouted to me to grab the rope and climb it. I tried my best, but I barely got my feet off the floor. The next time I saw such a rope I was in 11th grade, and now the PE teacher was called 'Coach.' I was afraid I wouldn't be able to climb any better than I had in fourth grade. Imagine how surprised I was when I reached the rafters!" *How do you explain this?*

Do animals learn the same way humans do?

"My sister used to have a cute trick she pulled on my father all the time when she was about two years old. When he held her in his arms, she used to suck in a little breath and then blow just a small puff of air into the side of his eye. She always laughed when she saw him blink. A couple of days ago I made a bet with her that it wouldn't happen again. So, she cuddled up next to my father, and then she drew in some air and pretended she was going to blow air at him, and you know what? He blinked! My sister's 15 years old now, and she hasn't done that in a dozen years, I'll bet. He blinked!" *What happened here?*

"I've been working here at the Resorts National for nineteen years cleaning these ashtrays, sweeping this carpet, and fetching cigarettes for the customers. Some people think I'm in a rut, but I like what I do, and I enjoy seein' other people spend their money. You know who's in a rut? You see that old man in the tan slacks over there playing the slot machine? He comes in here at least once a month with a bag of coins and he keeps putting them into that machine until he runs out. He's been doing it for years . . . long as I can remember. It used to be nickels, but now it's quarters. He's the one's in a rut." *How do you explain this gambler's behavior?*



Learning is defined as a relatively permanent change in a behavioral tendency that occurs as a result of practice or observation. Learning cannot be directly observed. Psychologists observe an action and infer from that action that learning has occurred.

There are many processes and events that are usually present when learning occurs. These include (1) a stimulus and a response, (2) motivation (the goals -- internal and external to the organism -- that cause and support most behavior), (3) association between stimulus and response that develops if they occur close together, or contiguously, in time and space, and (4) the Law of Effect. There are two types of reinforcers, positive and negative, but whether they are necessary for learning to occur is under serious question.

Classical conditioning, as originally developed by Pavlov, involves four events: (1) an Unconditioned Stimulus (US) -- an event that always produces certain response; (2) an

Unconditioned Response (UR) -- a behavior brought about by a US; (3) a conditioned stimulus (CS -- any event which, prior to conditioning, does not produce any response resembling the UR; and finally, (4) a Conditioned Response (CR) -- the response produced by the CS, neutral before conditioning, but resembling the UR after conditioning. Classical conditioning is done by repeated pairing of the CS and US, using a positive or a negative reinforcer.

Trial and error learning involves gradual elimination of errors where cues for the correct solution of a problem are all present at once (puzzle box) or presented one at a time (maze). Operant conditioning today frequently involves use of an enclosed chamber called a Skinner box. The Method of Successive Approximations is an oft-used training technique. It is based on shaping, a process in which an organism is reinforced for performing responses more and more similar to the desired one. Operant conditioning itself involves increasing the frequency of a behavior, either by using a positive reinforcer, or by removing a negative reinforcer or decreasing its intensity.

Classical conditioning is based mainly on reinforcements that elicit behavior. Operant conditioning involves reinforcing emitted behaviors. By contrast, observational learning or imitation seems to occur mainly because of its efficiency. If an organism sees a behavior and can perform it, then the more traditional classical and operant procedures work much more quickly.

There continues to be some debate as to whether or not classical conditioning and operant conditioning are identical. However, there is no doubt that they share a number of effects. Reinforcement is one important process involved in many types of conditioning. Continuous reinforcement (giving a reinforcer for every correct response) occurs less often than various forms of partial reinforcement. Such partial reinforcement schedules include fixed ratio, variable ratio, fixed interval, and variable interval. Reinforcements are primary (essential to the well-being of the organism) or secondary (learned, and not necessary to survival).

Primary and secondary reinforcement are factors in both classical and operant conditioning. They have other factors in common, including (1) the need to time a reinforcement correctly in order to increase the frequency of a conditioned response; (2) extinction, or the fading away of a response that is no longer reinforced (the Partial Reinforcement Effect suggests that extinction occurs more slowly when reinforcement has not given each time the response is made); (3) generalization, or the response to stimuli that are similar, but not identical, to a previously conditioned stimulus; and (4) discrimination, or

the ability to respond to one stimulus and not to another that is similar.

The effectiveness of punishment is related to (1) its intensity, (2) the consistency with which it is applied, (3) how soon it is applied after the offensive behavior, and (4) the strength of the response to be eliminated. It can lose effectiveness if it becomes an acceptable price for the behavior or if the response history creates a special set of circumstances. A punishment works best when it offers or rewards an alternative behavior.

The Review Questions will help with mastery of the materials covered in this chapter. After reading the chapter you may be interested in trying some of the suggested ACTIVITIES. Further information about selected topics within this chapter is available in follow-up readings suggested in the INTERESTED IN MORE? section.

Defining Learning



Consider how important learning is to you. You learn to find shelter when you see lightning or hear thunder. You learn to swim. You learn how to ride a bicycle. You even learn how to talk with your friends. You gain a lot of skill in your alternating roles of speaker and listener. The

examples could go on and on, and learning is constituted of a number of elements. Some uses of learning are of life-and-death importance, such as the behavior of heading for the surface when you run short of breath under water. Others are trivial, such as the fact that a ring that fits on your little finger won't even slide onto your middle finger. All of these examples share one thing in common: they are learned bits of behavior. Examples of learning are present everywhere. Learning controls human and animal behavior.

Learning can be defined as a relatively permanent change in a behavioral tendency that occurs as a result of practice or observation. Learning does not mean just strengthening responses. It may or may not produce an immediate change in behavior, and it may or may not require reinforcement for learning to occur, whether we're talking about classical conditioning or operant conditioning.

We have defined learning as relatively permanent; it cannot be something that is "here today and gone tomorrow." It must have a continuing effect. In a marathon race you run much more easily at the start of the race than you do at the finish. You don't learn to run more poorly during the marathon. You simply grow tired, but your fatigue will go away with time. This change in your behavior is not learned.

We have also defined learning as involving a change: there must be something different about us after we have learned. Notice we have not said it must be a change for better or worse. Learning occurs whether the behavior change is good or bad.

In addition, we have said learning involves an altered behavioral tendency. Tendency is the important word here. In watching your girlfriend fix her bike you may notice that she leaves her pliers inside the house when you both go in for a Coke. What happens if you're outside later, and she can't find the pliers? You describe to her where they are, or get them for her. You learned earlier where they were, but you didn't behave as if you had learned until much later.

Finally, we have insisted that learning occurs as a result of practice or observation: As is true for a player on a school or major league athletic team, practice, or the opportunity to repeat a behavior, adds to a person's ability to perform it. Children learn how to handle eating utensils in part by observing their parents at meal time. That's not the only learning in childhood, however, as discussed further in the Think About It.

Think About It

The question: The 11th grader in the story at the start of this chapter found he could climb the rope without having practiced in seven years. Yet, the last time he tried in the 4th grade he couldn't. Did he learn to climb the rope?

The answer: There are situations in which a person's behavior may change even without practicing or observing the behavior of someone else. The 11th grader had simply matured, but he had not learned how to climb the rope. By growing up he had gained an ability he had not previously had; he was able to climb better simply because his arms grew stronger relative to his body weight.

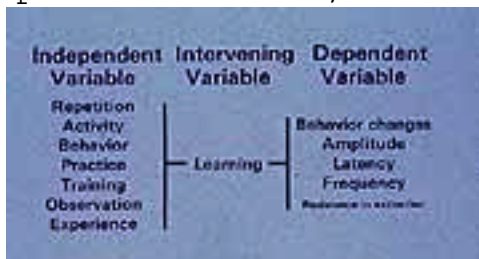
There is one more point you should keep in mind -- the importance of performance in demonstrating that learning has occurred. We may say, "He learned the poem" or "She learned how

to define 'independent variable'," but how do we know he or she learned? We find out by watching the person perform some act that shows the existence of learning. We talk about learning, but we never see it directly; we only see a performance. This is true whether we are dealing with classical conditioning, operant conditioning, or observational learning.

Elements of learning

There are many processes and events that are usually present when learning occurs. Some of them we'll deal with shortly; others we'll set aside for later chapters.

Stimuli and responses are among the most basic to psychology. A stimulus is anything that produces a change or sensation in a human or animal sense organ or causes a response. A response is any observable activity of a human or an animal. You can see in the Early Development Chapter how a tap on your knee (a stimulus) usually causes a small, involuntary kick by your foot (a response). The result of hearing a funny joke (a stimulus) is a laugh (a response). A stimulus is typically an independent variable, and a response is normally a dependent



variable. Learning, then, is an intervening variable. It summarizes what goes on between the time when we experience the stimulus and the time when we respond.


It's hard to talk about learning without mentioning motivation. Motivation concerns the goals, both inside and outside the organism, that cause and support any behavior. Thus, hunger is a motivating force. For now let's simply agree that the learning organism is motivated. We'll discuss motivation separately in the Motivation Chapter.

The principle of association (or contiguity) states that a stimulus and response become connected if they occur close together (contiguously) in time and space. It is a very simple concept, among the oldest and most interesting "laws" of psychology. Some say contiguity is all that's needed for learning to occur. Others say it is necessary, but not enough (alone), to assure learning. But almost everyone agrees that contiguity is important if learning is to occur.

Advertisers are constantly trying to get us to associate a product with a particular event or circumstance. Thus, we often see soft drinks linked with sports or outdoor work and play

activities. Dishwashing powders, on the other hand, are almost always advertised and displayed in a kitchen setting with a dishwasher or a set of dishes nearby. When the stimulus "All the News That's Fit to Print" makes you think of *The New York Times*, you can give credit to the principle of association.

Reinforcement



As a human or animal behaves, it gains reinforcers and punishments that strengthen or alter what it has learned. Robert Thorndike's Law of Effect states that when responses are reinforced, they will increase in frequency. Reinforcement is any event that increases the probability a response will occur.

For example, we are reinforced for putting money into a vending machine because of the food or drink we receive. If we like the food or drink we receive for our money, we are more likely to put money into the machine again later. Positive reinforcers are stimuli whose presentation after a response increases the probability of the response. Examples would include food for a hungry person, water for a thirsty person, or money for a poor person.

Effect of response	Type of reinforcer	
	Positive	Aversive
Reinforcer presented	Positive reinforcement	Punishment
Reinforcer withdrawn or withheld	Punishment	Negative reinforcement

A negative reinforcer is a stimulus whose removal after a response increases the probability that the response will occur again. If, as a child, you unintentionally touched a hot stove, your withdrawal was negatively reinforced by the reduction of

pain. The next time you were much more likely to withdraw -- far enough so you wouldn't be burned at all!


Is reinforcement necessary for learning to occur? Probably not, although there is some controversy on this point. Gaining a reinforcement is sufficient to encourage learning, but since we can learn by observing, reinforcement itself is apparently not necessary for the learning process.

Classical Conditioning



One of Ivan Pavlov's (1849 - 1936) most famous series of studies was an analysis of classical conditioning. Generally the term conditioning is applied to simple learning situations, whereas learning identifies the more complex forms.

Pavlov was studying salivary secretions and their role in the digestion of food when he noticed a very interesting reaction in his dogs. We know that the sight of food causes a hungry dog to salivate, but Pavlov was puzzled: his dogs were beginning to salivate as soon as he walked into the room. It happened before the food was even in sight. Pavlov decided to study this happening, shifting the emphasis of his research from physiological to psychological processes.



The basic procedure in classical conditioning involves presenting a neutral stimulus. This stimulus is followed shortly afterwards by a positive reinforcement, such as food, or a negative reinforcement, such as stopping a pain. Pavlov used the sound of a bell as the neutral

stimulus and food for the positive reinforcement. The illustration shows a set up built by one of Pavlov's students, though current evidence suggests Pavlov used a much simpler device to measure salivation. He simply operated on his dogs' salivary ducts, diverting them through the dogs' cheeks. Attached on the outside was a laboratory tube simply to collect the saliva each dog produced. There are four important elements in classical conditioning.

(1) An unconditioned stimulus (or US, as psychologists call it) is any event that reliably and naturally produces a response. Since food causes salivation, food works as an unconditioned stimulus.

(2) An unconditioned response (or UR) is any behavior predictably caused by an unconditioned stimulus. Salivation and

eating are unconditioned responses. They can reliably be expected to occur if a hungry animal is given food. Water as a US reliably leads to drinking (a UR) by a thirsty organism.

(3) A conditioned stimulus (CS) can be any event which is "neutral" at the start of conditioning. It can be anything that does not normally cause any response resembling the unconditioned response (UR). Obviously, in Pavlov's original experiments the bell could be the CS since a bell would not cause a dog to salivate prior to training.

(4) A conditioned response (or CR) is any behavior that is like the unconditioned response but is caused by the conditioned stimulus (CS) after conditioning. Conditioning is said to be complete when the conditioned stimulus (CS) will cause a conditioned response (CR) to occur before the unconditioned stimulus (US) is presented. This is discussed further in the Think About It.

Think About It

The question: Do animals learn the same way humans do?

The answer: Pavlov used animals to study conditioning. Yet, they were studied in the hope of better understanding humans. The principles developed in studying simple processes in animals could often be applied to humans. Many of the basic learning processes are very much alike. However, this is not so true of the complex processes we'll be studying in Chapter 7. As the complexity of the learning processes grows, it becomes more difficult to apply knowledge gained from studying animals directly to the performance of humans.



Reinforcement in Classical Conditioning

Remember, when Pavlov started feeding his dog, it salivated automatically. The US food reliably produced the unconditioned response (UR). However, the CS bell was not

classically conditioned to produce the response CR salivation. There was no connection between the UR and the CR. But what happened later in training? If the CS bell was presented, the dog was more and more likely to salivate. The dog "knew" that



the USfood was coming. Once that relationship had been formed and performed, we say that the dog had been conditioned. The reinforcement for the conditioned stimulus and the conditioned response it caused was the USfood-URsalivation, eating that followed it.

Pavlov's original experiment used a positive reinforcer. Yet, some of our learning is the result of negative reinforcers. The same basic order of stimuli and response still operates. For instance, our fear of a buzzing bee is a result of classical conditioning. Based on a negative reinforcement (avoidance of pain by escaping the bee), these fears result from only a few childhood experiences.

Classical conditioning in Pavlov's dogs is perhaps the simplest type of learning to be studied in the laboratory. We can also see this type of learning in our everyday lives. One example is provided by the ice cream man who drives around many city neighborhoods ringing a bell to attract attention. He is simply applying the principles of classical conditioning. Ring the bell and feed the children. Pretty soon the children will come running as soon as they hear the bell. This is classical conditioning using a positive reinforcement.

USING PSYCHOLOGY: Breaking Classically Bad Habits

If most of what we do is learned, we ought to be able to change our behavior if we want to. Yet some of us still smoke, even though we know we shouldn't. Others of us bite our fingernails -- only occasionally, of course. In short, few of us are perfect. How should we go about changing previously learned responses? It is, in a sense, the hardest test of all since we must learn new responses to old stimuli. How can it be done?


Take drinking liquor as an example. A person may pour a drink in response to a variety of stimuli, including social stress or any situation where he or she doesn't know what else to do. The key in counter conditioning is to develop new

responses to the old stimuli. One way you do this is by associating the stimuli (such as stress) that usually bring about the undesired response (such as excessive drinking) with a new, unpleasant response (such as nausea).

For example, the alcoholic who is addicted to alcohol can be injected with a drug called Antabuse. The drug causes no ill effects at all -- you can't even tell that it's been injected. That is, you can't until you take drink of alcohol of any kind. Then a kind of internal horror is unleashed. You feel nauseous; you want to (and usually do) vomit all that's in your stomach. Generally, the more alcohol you've taken in, the worse you will feel.

Having gone through this suffering once, the alcoholic who continues taking Antabuse very seldom wants to repeat the drinking experience. Thus, whatever were the stimuli that formerly elicited drinking, those same stimuli now begin to elicit thoughts of nausea and sickness instead. The result is that behavior changes -- the alcoholic learns new responses to the old stimuli as a result of counter conditioning.

Operant (Instrumental) Conditioning: A Predecessor



Now let's look at a more recently identified form of learning -- operant, or instrumental, conditioning. There are many research techniques that involve operant conditioning. We'll briefly examine one such technique from which our modern procedures for operant conditioning have developed: Trial and error learning.

Pavlov was working out the details of classical conditioning in Russia at the turn of the century. Meanwhile, the American psychologist Robert L. Thorndike was studying trial and error learning.

In his experiments, Thorndike put a cat in a puzzle box with a trick door. To escape, the cat had to open the door by pulling down a string, pushing a bar, or emitting some other response. The cat usually tried a variety of wrong responses in attempting to get out of the box. Thus, it was called trial and error learning. By accident, most cats finally opened the door.

After more practice the cat could reduce its errors until only the correct response for opening the door remained. The animal's behavior operated to or was instrumental in obtaining its release.

In the puzzle box all of the equipment to solve the problem was in sight of the animal. In a maze, which is another form of puzzle box, the decisions must be made one at a time. A maze usually offers only one path as the correct route from a starting point to freedom, food, or some other reinforcer. The study of trial and error learning was a predecessor to the more detailed studies that have been directed at operant conditioning and the roles of reinforcement and punishment in such conditioning.

Operant Conditioning: The Process and Equipment



Operant, or instrumental, conditioning is a learning process in which the frequency or probability of a response is increased by giving a reinforcement whenever that response takes place. The basic process is simply a matter of timing the

reinforcement so that it immediately follows the response to be learned. To aid in studying the processes involved in operant conditioning, B. F. Skinner developed a piece of apparatus which is now called the Skinner box. How is it used? One way is simply to place a hungry rat in the box. Set up the box so that every time a lever is pressed a single small pellet of food (about the size of a BB-pellet) will drop into the food cup. Then you can demonstrate trial and error learning in a much more sophisticated piece of equipment -- but you really haven't moved past Thorndike.



More fun, however, is to use a process called the Method of Successive Approximations. This method involves a series of steps. Step One is to train the rat to eat from the food cup -- a process called magazine training. If the animal has been deprived of food for 23 hours, it will

explore the chamber actively when placed in it. If the food delivery apparatus is repeatedly operated, the rat will soon start eating when it first hears the sound of the apparatus being activated. This portion of the training is pure classical conditioning where the CSfood apparatus sound is conditioned to yield a CR which will serve as a secondary reinforcer for the following steps.

Step Two involves delivering a pellet of food only when the animal is in the general neighborhood of the cup.

Step Three is yet more demanding. Now a food pellet may be delivered only when the rat is close enough to the bar to touch it, even though we are not yet demanding that it do so.

Step Four means that food is delivered to the animal only when its foot is raised or touching the bar.

Finally, in Step Five, food is delivered only when the rat presses the bar itself. That action will close a switch to operate the food delivery mechanism. Now the animal has fed itself and operant conditioning is complete. Steps two to five are called shaping.

Feature 1

"DON'T WALK UNDER THAT LADDER!"

Reinforcement leads to an increased frequency of the responses that occurred just before the reinforcer. What happens to irrelevant, or random, responses that just happened to occur prior to a reinforcement? Do they increase in frequency? Are these irrelevant responses learned?

For example, suppose you are using operant conditioning to train your dog. Suppose as you reward her for the first time for the complete act of sitting up and barking ("speaking"), your dog touches her nose to the ground at your feet. And suppose that just by accident your dog happens to touch her nose to the ground just before barking correctly and being reinforced the second time.

By now your dog is beginning to learn that sitting up and barking leads to a small piece of food. However, she is also touching her nose to the ground just as often and just as regularly. Yet the nose touching the ground has nothing to do with getting the food. Your dog has been rewarded for sitting up and barking, but she's also been rewarded for touching her nose to the ground. In short, your dog has accidentally conditioned herself in error.

The late Vince Lombardi was a highly successful coach of the Green Bay Packers (and later the Washington Redskins). He always wore a brown suit to the Sunday professional football games he coached. Neither that suit nor its color had anything to do with winning, yet he continued to wear it. Knocking on wood or throwing salt over your left shoulder for good luck are other examples. They are behaviors that are not related to reinforcement, yet they continue to occur anyway. It's probably due to an accidental earlier reinforcement. They are, in short, superstitions—bits of behavior with no demonstrated impact on daily events.

Reinforcement and Punishment in Operant Conditioning



Operant conditioning is accomplished through the systematic reinforcement of an organism's responses. For example, a circus trainer often gives an animal a positive reinforcement (small pieces of food) after the animal correctly performs a

trick. A positive reinforcer is any event which increases the frequency of the behavior being emitted just prior to the event. The trainer is using the reinforcer and principles of operant conditioning to increase or maintain the frequency of the animal's doing the trick correctly. This is the Law of Effect in action.

In addition to the use of positive reinforcement, operant conditioning can also be achieved using negative reinforcement. The procedures are almost the same as those with positive reinforcers discussed in presenting the processes and equipment utilized in operant conditioning. Of course, with negative reinforcers there is an obvious change in the nature of the reinforcement. Now the organism is trying to get away from or reduce an undesirable state. Thus, a negative reinforcement is any event the removal or reduction of which leads to an increase in the behavior being emitted just prior to the event.

One example would be taking an aspirin to reduce a headache. The seek-find-and-swallow aspirin response is negatively reinforced by the later reduction in pain. Another example: When walking barefoot in the summertime, you are not

likely to stay very long on a heated black highway surface. It's hot -- very hot. It is negatively reinforcing to step off the highway onto cool grass.

But this last example raises one other point that needs to be emphasized. A negative reinforcer is any event the decrease or removal of which leads to an increase in the frequency of the behavior preceding it. Contrast negative reinforcers with punishers. A punishment causes a decrease in the frequency of the behavior preceding it, as summarized in Table 1.

Table 1

Reinforcers and Punishers

	Left to its own, the organism _____ seek the goal	
	Would	Would Not
The goal is:		
Presented	Positive Reinforcer	Punisher
Reduced, removed	Punisher	Negative Reinforcer

USING PSYCHOLOGY: Breaking Bad Operants


As noted in our discussion of classical conditioning most of what we do is learned, so we ought to be able to change our behavior if we want to, especially using the principles of operant conditioning. The question remains: How should we go about changing previously learned responses? It is a difficult learning situation because we must learn new responses to old stimuli.

Smoking is a behavior that many people say they would like to give up. Rather than changing the response directly by counter conditioning, the response can be elicited so many times as to exhaust the person literally. Many nonsmoking adults can recall the time when they were caught smoking as a child by a parent who showed no anger -- in fact, just the opposite: "You want to smoke? Well, here, have one of mine. . . . Done with

that one, well, have another. In fact, why don't we have a cigar? Let's really smoke!"

What occurs is easy to predict. The child gets so sick of smoke, cigarettes, and smoking that the next time the thought of smoking enters his or her head it is followed immediately by memories of how bad it was the last time. Flooding, in a sense, created counter conditioning by using the offending response itself to generate nauseous feelings because of concentrated overexposure.

Classical Versus Operant Conditioning



You have now considered two types of learning: classical conditioning and operant, both of which use positive or negative reinforcement. Although the procedures and reinforcements may differ, the two types of learning are very similar.

There is one point of difference, however, and this is in the responses.

In classical conditioning procedures the response is always elicited, or forced, from the subject. The subject has little or no choice regarding whether or not to respond. This is because of the nature of the US-UR pairings that are selected. For the hungry animal food will always elicit saliva; for the thirsty animal water will always elicit drinking. The animal simply has few other choices.

By contrast, in operant conditioning the response must always be emitted, or given, before a reinforcement can be gained. If a rat chooses just to sit in a Skinner box, no responses will be recorded or reinforced. Not until the organism emits a response can anything be accomplished.

Observational Learning

For many years learning was analyzed only in terms of classical and operant conditioning. Most of our explanations so far have placed a lot of stress on the role of reinforcement. Gain a reinforcer and learn the behavior. Yet, think for a moment about your everyday life. Do your parents play tennis? If you've started practicing the sport too, it's unlikely you began by holding the racket like a golf club! In short, you had already learned a lot just by watching them play.

Since the 1960's more and more psychologists have begun to recognize the obvious: we learn just by watching each other or reading about or seeing pictures of someone making a skilled response. Then -- sometimes much later -- the observer can try to imitate the model's response. It's called observational learning, social learning, or imitation.

Albert Bandura suggests that observational learning involves three elements: a modeling stimulus, a matching response performed by the observer (learner), and a reinforcer for the matching responses. There are four processes that control your observational learning: you must pay attention to the model, you must show retention of the model's behavior, you must be able to perform the behavior, and you must be motivated to perform. Motivation can be provided by external rewards you hope to gain, by the pleasure of seeing someone else reinforced, or by your own self-generated reinforcers.

Many of the factors influencing operant and classical conditioning also affect learning by imitation. If the model (the person being observed) is successful, we are likely to follow his or her example. The more similar models are to us, the more likely we are to imitate them. The more similar the environment and stimuli we share with our models, the more likely we are to imitate their actions. But the main advantage in imitation is its efficiency. Once any of us has learned how to perform an important behavior, it can often be passed on best through this method of observational learning.

USING PSYCHOLOGY: Replacing Bad Habits

One of the problems causing great difficulty in unlearning or changing (bad) habits is referred to in our discussions of using psychological principles to change habits through classical conditioning and operant conditioning. It concerns the same situation that occurs when you and a significant other (boyfriend; girlfriend) break up -- especially if he/she precipitated the break-up. The problem is that the same old stimulus (boyfriend; girlfriend) suddenly expects new or different responses (frowning, avoidance).

Both counter conditioning and flooding involve direct attempts to change the responses elicited by a particular stimulus. Another obvious approach is to prevent the stimulus from occurring. It is very similar to the advice given by a doctor to one of her patients who complained, "Well, doctor, my arm has recovered pretty well. It only hurts now when I twist my wrist around this way and try to reach way up behind my back to scratch." Replied the doctor, "Why not just stop twisting

your wrist and reaching way up behind your back that way?" In other words, if it doesn't hurt when you don't do it, why do it? The essential point is a valid one: simply don't present the stimuli that elicit the undesired response.

This is one of several principles underlying the concept of "half-way houses," in which convicts about to be released from prison are aided in readjusting themselves to the outside world. Help is made available to convicts trying to rearrange their style of life. This discourages contact between the convicts and the persons or environment that caused them to get into trouble in the past. Halfway Houses give participants exposure to others with similar problems developing new adaptive responses in life. Through modeling of responses to new stimuli and reducing the likelihood of experiencing old stimuli and the associated problematic responses, new responses are learned.

Schedules of Reinforcement and Learning

At this point we've looked at the most basic aspects of both classical and operant conditioning and observational learning. It may seem that they are quite different types of learning, but in many ways they are, in fact, very similar. There are a number of different learning events, many of which are shared or occur not only in classical and operant conditioning but also in observational learning.

In conditioning, reinforcing every correct response is called continuous reinforcement. Almost nothing we humans do is reinforced every single time we do it. Continuous reinforcement (or CRF, as the operant conditioners call it) does not occur in the real world very often. Usually, reinforcements occur, or are given, only some of the times a correct response occurs. This is called partial reinforcement. It is used in both classical and operant conditioning procedures.

Basis for Reinforcement	Schedule	
	Ratio	Interval
Rigid (fixed)	FR	FI
Variable	VR	VI

Partial reinforcement can be accomplished in one of two ways. The schedule of reinforcement is a statement of the rules governing how and when reinforcements are being given. One partial reinforcement schedule gives reinforcements for either a

fixed or average number of correct responses. These are the ratio schedules. The other kind of reinforcement schedule gives reinforcements for the first correct response after a certain

fixed or average amount of time has passed. These are the interval schedules. Let's look briefly at each of the four possible schedules of partial reinforcement in somewhat more detail.

In a fixed ratio schedule, the experimenter gives a reinforcer for a fixed number of responses. Every second, third, or fourth correct response might be reinforced. It is quite rigid and predictable, and it produces very stable response rates. The migrant workers who follow the crops up and down the West Coast of the United States each year are often paid on a fixed ratio schedule. They get a certain amount of cash for every basket they harvest.



With a variable ratio schedule the organism will be reinforced on the average, or "in the long run," every second, third, or fourth (or more) time it responds. However, for any particular response there is no guarantee there will be a reinforcement. The schedule is based on

performance averages rather than a specific number of responses. It is on this schedule that gambling machines such as the one-armed bandits are operating. The "gamble" is that you never know when a reinforcer in the form of a payoff will be won as suggested in the Think About It.

Think About It

The question: Remember the old man in the tan slacks mentioned earlier? The janitor at a Las Vegas nightclub-hotel said the man arrived once a month or so with a bag of coins. He played the slot machine each time until he ran out of coins. What kind of reinforcement is involved here? What keeps the old man coming back so often to "play the slots"?

The answer: It is known, of course, that in the long run the customer will lose money playing the Las Vegas slot machines. The machine must be set to give the "house" a small profit to make it worthwhile to have the machine installed in the first place. However, the man in the tan slacks couldn't care less. He's putting a coin in and pulling the handle. Every now and then he gets a lot of coins back. He is "hooked" on a variable ratio schedule of partial reinforcement.

On a fixed Interval schedule an organism is reinforced for the first correct response after a set interval of time has

lapsed since the last reinforcement. It doesn't matter how often the organism responds during the interval following the reinforcement. It won't be reinforced until the first correct response following that interval occurs. An organism on a fixed interval schedule for a long time yields very stable performance reflected in what is called a scalloped performance curve. The reason? Organisms learn that they are never reinforced immediately after gaining a reinforcement -- so they stop responding!

If your instructor gives you a quiz at the end of each week, your studying is being "reinforced" on a fixed interval schedule. The interesting thing about fixed interval reinforcement schedules is that they allow us to "predict" when a reinforcement is due. The response rate picks up just before the end of the interval. And isn't that pattern of behavior exactly the pattern you follow in studying this material? Most likely you study psychology more the night before a quiz than any other night of the week.

With a variable interval schedule the organism is reinforced for the first response after an interval of time has elapsed since the last reinforcement -- but the precise interval varies. It may average 5, 10, 30 seconds, or whatever. It is unknown until a correct response yields reinforcement.

Does your instructor give quizzes without announcing beforehand when they'll occur? If so, then your studying behavior is being maintained or reinforced on a variable interval schedule. You never know when you'll get a reinforcer. As a result you tend to respond (that is, study) more regularly, but less than you would on a fixed interval schedule. You never know when that next quiz is coming. The same schedule describes the situation when you are trying to reach someone whose phone is busy. Your desperateness to talk with him or her will determine how often you dial again and again. . . .

Primary and Secondary Reinforcement

Reinforcers can be divided into two broad categories, depending on how necessary they are for the organism's survival.

What happens if you deprive an animal of food or water or expose it to some kind of pain? The animal (or human, for that matter) usually learns very quickly to perform whatever actions are required to gain the required food, water, or reduction of the pain. These are all examples of primary reinforcers. They gain their ability to reinforce without the need for much prior learning. For all of these reinforcers there is a significant contribution of heredity to the organism's sensitivity to the

goal. There are not very many primary reinforcers, and each is specifically required/sought when a specific need such as hunger or thirst is being experienced.



By contrast there is a much larger group of stimuli that may become reinforcers. These secondary reinforcers gain their ability to reinforce because learning has gone on before. Remember how much you salivated before your last holiday dinner was

finally put on the family table? Most of the stimuli causing your mouth to water were learned secondary reinforcers. Many times before you've associated a certain time of day, the dishes on the table, the smells of the cooking food, and all the kitchen preparations with eating. These stimuli become classically conditioned, secondary cues that reinforce you. Your response is based on the many times in the past when these cues have been associated with the food you were seeking. The linkage to and the importance of the cues is learned, not inherited, and based on experience (learning), many such events can become reinforcers, as discussed the Think About It.

Think About It


The question: Why is money such a powerful reinforcer for humans? What kind of conditioning is involved in making money into an effective secondary reinforcer?

The answer: In terms of classical conditioning, the money is the conditioned stimulus. What accounts for the powerful nature of the conditioning is that the same CS-US pairing occurs many times. That same Cs (money) is paired just before people acquire food -- a cupcake, a milkshake, or whatever -- or when they buy merchandise -- a radio or any number of other things people normally desire. Money (the CS) is present time after time just before we get things in which we are interested. So time and time again money is conditioned as a positive secondary reinforcer. It always precedes the reinforcements we need or want.

Timing of Reinforcements

In classical conditioning the Unconditioned Stimulus is the reinforcer. In successful operant conditioning the

reinforcement is whatever follows the emitted response. If there is no increase in the frequency of the preceding response in operant conditioning, the following event is -- by definition -- not a reinforcer.



In both classical and operant conditioning timing is very important and often determines whether a reinforcement has any effect at all on conditioning. In classical conditioning, the optimal interval between the CS and US is about 1/2 second, though under certain conditions -- for instance, conditioned food aversion -- the interval may be several hours. The young girl in the "Think about it" inhaled and puckered her lips, acting as if she was going to blow a brief burst of air into her father's eye. Without knowing it she timed her inhale-followed-by-pucker almost perfectly. It would have been difficult for her to hold her breath for very many seconds. If she had, her father would have had no reason to blink at the sight/sound of the pucker. Conditioning would not have worked.

Think About It

The question: Remember the story about the little girl who blew a small puff of air into her father's eye? Her brother was surprised that his father blinked when his sister inhaled and puckered her lips 13 years later as if she were going to blow into his eye again. What kind of conditioning was involved? What was the reinforcer?

The answer: What she had done, of course, without being aware of it, was to condition her father, using a negative reinforcer. The puff of air was the US, and the eyeblink was the UR. Once the conditioning of her father was complete, the inhaling of air and puckering of her lips served as the CS. It produced his anticipatory eyeblink as a CR -- much to the teenager's surprise and delight.

Operant conditioning is also more effective if you can give reinforcement right after the response. Suppose you're reinforcing someone every time he or she uses the pronoun "I." Any time "I" is said, you smile and nod in agreement. It will increase the frequency of "I . . ." statements, but not if you delay the smile and nod. Five seconds after starting an "I" sentence your friend may have spoken two dozen words. How could he or she possibly guess which word(s) caused you to smile? Again, timing is critical.

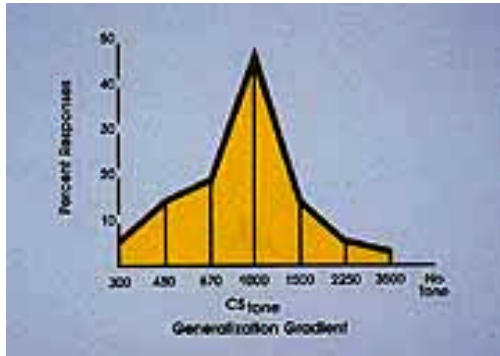
Extinction

Elsewhere in the chapter we discussed Thorndike's Law of Effect: Responses that are reinforced increase in frequency -- that is, they are learned. But, what happens to any response that is no longer followed by the reinforcement that led to its being learned in the first place? If the reinforcer is no longer given, the frequency of the learned response will decrease. This decrease is called extinction. Consider the animal in the Skinner box who has been getting a food pellet each and every time it hits the bar. What happens if the animal is suddenly placed on an extinction schedule and never again reinforced? The once-reinforced response is no longer producing the food the animal wants. It makes sense for that animal to quit pressing the bar and to try other responses. Extinction, therefore, comes easily.

On the other hand, think of an animal that is being trained on a partial reinforcement schedule. Neither the precise number of responses nor the exact moment of reinforcement is known to the animal on any schedule of partial reinforcement. This animal goes for long stretches before it is reinforced. When placed on an extinction schedule (no more reinforcers are to be given), the animal won't become aware of this for some time. Only slowly will the animal begin to show any decrease in its level of performance. This resistance to extinction is called the Partial Reinforcement Effect. It apparently applies to most organisms at all levels of the animal hierarchy.

Generalization

There are two processes that are among the most important aspects of learning. They occur in both classical and operant conditioning. One is generalization, and the other is discrimination.



To illustrate generalization let's look at an event that has occurred at some time to almost all of us. Through years of training, most drivers have learned how to respond when a car with a flashing red light on its roof appears behind them. They pull over to the side of the road

and stop driving. It is a highly predictable response for most drivers.


What would happen if a car with a light that was slightly yellower than the traditional police-cruiser red were to appear? Would the driver still pull over? Probably so. What if the flashing light were made more and more yellow, and less and less like the red light to which almost all drivers would respond? We would observe less of the response of pulling over to the side of the road.

The process of responding to stimuli similar to, but not identical with, the original training stimulus is called generalization. The conditioned response declines in probability as the similarity between original and test stimuli decreases. This is called a generalization gradient.

Generalization is a very important concept for all of psychology. It is easy to see that if an organism had to learn a new response to each stimulus it might ever encounter, it would be dead before it had even learned enough to survive. As a result animals and humans show a marked ability to generalize. We are able to react appropriately to any stimulus similar to ones with which we have had previous experience. How many different ways are there to open a door? At the first meeting of your class, your instructor did not need to have a "Door-Opening Drill" so that everyone would know how to leave the classroom. We all have a generalized door-opening response based on past experience with doors similar -- but not identical -- to your classroom door.

Discrimination

Quite obviously, there must be limits to generalization. Otherwise, once an organism had learned one response, that response would occur to all stimuli! The ability to react differently to similar but distinct stimuli is called discrimination. It is an ability equally as important as generalization. Both are critical to the survival of any organism.



How can we demonstrate discrimination? It's easy. Many abilities make a person skilled at "body language" (see the Language and Communication Chapter). One is the ability to discriminate among cues. A very simple cue is facial expression. If you had a favor to ask of your psychology instructor, you would be more likely to ask it if the instructor was smiling. You'd be much less likely to ask it if the instructor was frowning. In the past you have often been reinforced by people who are smiling, but you have seldom been reinforced by people who are frowning. You rely on your ability to discriminate a smile from a frown to increase your chances of success with your instructor -- or anyone, for that matter.

The same principle of discrimination is at work when you encounter a traffic light while driving. If the light is green, you proceed; if red, you stop -- in many places, if it's yellow, you speed up! Again, it is simply a matter of discrimination.



Punishment

Decisions about the proper role of punishment in modern society are constantly being made, but those decisions are never easy. When should punishment be used? How effective is punishment in causing a change in behavior? What factors make punishment most likely to succeed? Applying a punishment should be carefully guided by an understanding of learning principles. The following conditions, related to learning, help determine how effective a punishment is likely to be. These conditions, in turn, are impacted by a broader array of environmental and

experiential factors which also impact the effectiveness of any punishment and suggest other alternatives to punishment.

Punishment should be severe enough to cause some discomfort. However, the intensity should not be so severe as to hinder the organism from responding at all. For example, new recruits in military training use blanks in their first field exercises, rather than live "enemy" bullets. The intense, punishing conditions of battle are there. However, death -- the ultimate punishment -- doesn't follow if the forgetful recruit gets caught in a line of fire.

The response to be altered must be punished with consistency. It is much easier to learn what is being done wrong if the punishment for errors always occurs and is always the same. Suppose a child is told she will lose the right to ride her tricycle for 24 hours if that tricycle is not put away properly at the end of each day. That punishment should be applied each time the offense occurs.

As you read elsewhere, the timing of a reinforcement is quite important in learning a response. Timing is also important in punishment. Generally, the greater the time between the response and the punishment it produces, the weaker the effect of that punishment on the response. What's the favorite threat of frustrated baby-sitters when the children get out of hand? "Just wait until your parents get home!" It may be the least effective punishment of all -- even less so because understanding the threat requires a level of abstraction beyond the capabilities of many children.

The strength of the punishment must be related to the strength of the response being punished. The greater the pleasure to be given up, the more severe the punishment must be. For instance, people who enjoy smoking say they feel relaxed when they smoke. They say it helps control their nerves. It keeps their hands busy. In fact, they say it generally makes them feel more socially acceptable or used to, until lately. Such people are not going to be convinced to stop smoking by reading, "SURGEON GENERAL'S WARNING: Smoking Causes Lung Cancer, Heart Disease, Emphysema, And May Complicate Pregnancy." Despite the truth of the warning, the punishment is simply too weak relative to the strength of the response to be much of a deterrent.

A punishment may lose its ability to alter a response. What is at first a punishment strong enough to alter behavior may eventually mean little to the persons or animals being punished. Alternatively, the benefits gained by the organism from the response being punished may outweigh the costs associated with the punishment. The organism experiences

adaptation in which he/she comes to view the punishment as an acceptable cost for the benefits received.

Factors which Influence Punishment

A punishment becomes more effective if it not only punishes the response to be altered or extinguished, but also reinforces, permits, or rewards alternative responses. Traffic signals control our behavior in part by "punishing" us. They make us stop on red, but they hold the promise of greater reinforcers. The reinforcer for waiting at a red light is less interference from cross traffic later.

How behavior is affected by punishment depends on how and why a response was gained -- its response history. The results of the same punishment applied to responses with different histories will vary widely.

Age is another factor that may influence the effects of a punishment. My youngest child was discouraged from behaviors such as stirring the fish bowl water by a firm "No, no." For about six months -- between 15-21 months -- any time he was about to engage in a behavior which he knew would elicit a "No, no", he starting saying "No, no" as he was positioning himself to do the errant behavior. At that age, "No, no" was not an effective punishment. His older brother always enjoyed watching his mother make cookies. Even at age two he tried to stick a finger into the batter. A disapproving and stern "No, no" would change this behavior since the fun of watching could be continued just the same. The "punishment fits the crime."

Alternatively, suppose a child is seldom given much attention except for misbehavior. In that case, the punishment serves the child's purpose -- getting attention -- and bad behavior will continue or even increase!

Punishment: Cues and Alternatives

Often, the response being punished is not the only one possible in a given situation. As long as other responses can be made, punishment need not be very severe in order to guide the behavior.

Has something like this ever happened to you? You were taken as a child to your grandparents' home. You started to talk at the beginning of the evening meal when your grandmother was attempting to bless the food. It may have happened only because your parents had foregone the food-blessing tradition. With only a hint from your parents you could easily select

silence as a better way to respond. Little punishment was necessary.

There is still a lot of argument among psychologists as to the precise effects of punishment. Clearly, refraining for a while from an action that has been punished doesn't mean that one won't do it in the future. Punishments work better if they suggest a better course of action. To be most effective, punishments should involve learning specific alternative responses. Time in prison is not enough. Time spent learning another set of responses in prison is more effective.

Which should be used? Reinforcement? Punishment?

The issue of punishment and when to use it is one we all face time after time. Is punishment ever justified as a means of control? Is corporal punishment desirable in the school classroom? How about in prisons? These are oft-debated issues in our society, about which people hold many different views. The science of psychology offers a powerful set of well-understood learning principles that can be applied to modify behavior. But the scientific view is not the only one to be considered. There are moral and religious, legal, and ethical issues that also must be examined in connection with the modification of behavior.

Should we reinforce the good or punish the bad? There's no easy answer. A punished response is simply suppressed; once the threat of punishment is removed, the unwanted response often returns. A variety of sources indicate that the most effective means of altering behavior -- if you must choose between reinforcement or punishment -- is to reinforce an alternative response rather than to punish the undesired response. By reinforcing the better, alternative response, you strengthen the possibility it will recur.

REVIEW QUESTIONS

DEFINING LEARNING

1. How do psychologists define learning?
2. Can you observe learning? How is it measured?

ELEMENTS OF LEARNING

1. Why is learning best viewed as an intervening variable?
2. What independent variables influence learning?
3. What dependent variables illustrate learning has occurred?

REINFORCEMENT

1. What is the Law of Effect?

CLASSICAL CONDITIONING

1. Who pioneered classical conditioning?
2. What is classical conditioning?

OPERANT CONDITIONING

1. Describe trial and error learning.
2. Define positive reinforcer and negative reinforcer. Give an example of each.

OPERANT VS. CLASSICAL

1. What is the difference between trial and error learning and operant conditioning?
2. Are operant and classical conditioning actually the same thing?

OBSERVATIONAL LEARNING

1. What is the main advantage of observational learning?
2. How does observational learning compare with operant and classical conditioning?

SCHEDULES OF REINFORCEMENT

1. Describe the five different schedules of reinforcement and the effect of each on behavior.
2. How do primary and secondary reinforcers differ?
3. Define the following terms: extinction, partial reinforcement effect, generalization, and discrimination.
4. What factors do operant and classical conditioning have in common?

PUNISHMENT

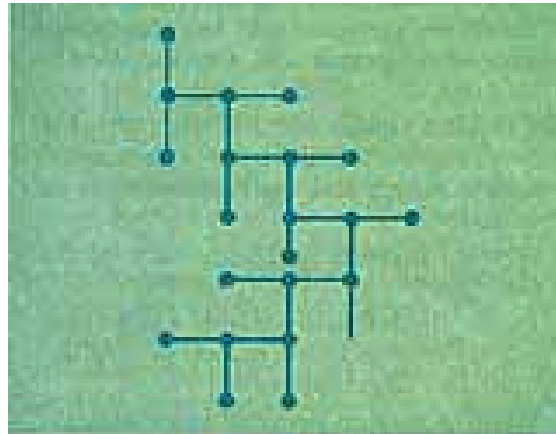
1. Describe the factors involved in the effectiveness of punishment.
2. Which do you think is better -- to punish the bad or to reinforce the good? Explain.

ACTIVITIES

1. This idea is easier to describe than it is to do, but give it a try. When you talk with a friend keep track of how many times he/she says the word "I," or uses a negative word (such as "not"), or uses a plural form. Once you have established a base rate, start reinforcing your friend every time he or she uses the word you want used (let's say "I" for our purposes). What's reinforcing? A smile. A nod of agreement. A mumbled "Mhmm." You are using a positive,

secondary, or social reinforcer to influence your friend's frequency of use of a specific word. Does the average use of the word increase? How much? Can you extinguish the behavior by withholding your reinforcement? Try it!

2. You can make a maze using thumb-tacks and string or rubber bands. By placing the thumb tacks at every choice point and every end of a blind alley, construct a maze similar to the pattern you see here:



See if you can find some friends who will help you. Blindfold them before they can see your maze. Put their finger on the starting point and explain to them that they are to find a path to the goal making as few errors as possible. Make a graph of your results, identifying each attempt (or trial) on the X axis (abscissa, see Chapter 2) and the number of errors (or the length of time) on the Y axis (ordinate). How many trials does it take before your friend can go through the whole maze with no errors? How does your graph show you that learning is occurring? When is learning best? Why?

3. In addition to classical and operant conditioning, we've also discussed observational learning. Complex activities are sometimes learned best by using a "Watch me, this is how it's done" approach. Interview a member of your family or a friend who is involved in a trade such as carpentry, plumbing, electrical wiring, or auto mechanics. How much of their training in the trade was achieved by reading? How much was learned simply by watching to see how a senior crafts-person did the job? Would they agree that observational learning was an efficient means of learning their trade?


4. What controls your behavior? Are you punished at home? At school? Make a list of the punishers and the positive and

negative reinforcers that are used by friends, family, and teacher to control your behavior. How do you think these punishers and reinforcers influence you? Which would you rather "earn," a punisher or a reinforcer?

5. Career Search. Advertisers employ psychologists to study consumer behavior. You read that we learn to link products with slogans. What else do advertisers hope we'll associate with their products? Pay close attention to the commercials on television for the next few days. What positive themes do you see advertisers pairing with their products? How do the commercials encourage you to associate the positive aspects of some other event or person with the product being advertised? Pool your observations with your classmates.

6. Here's a way to prove to yourself that observational learning is very efficient in some learning situations. Ask ten people you know to take the following brain-teaser: Connect all nine dots in the following diagram with four straight lines without lifting your pencil from the paper.

When they give up (as we bet they will!), show them how it's done. Then a week later, ask the same ten people to try it again. How many of the ten could solve the problem after seeing the solution just once? What is the solution? That's easy. The answer is as follows:



INTERESTED IN MORE About LEARNING?

FERSTER, C. B. & CULBERTSON, S A. (1982). *Behavior Principles*, 3rd ed. Prentice-Hall. An advanced introductory text on principles of operant conditioning. A systematic analysis of research studies of human behavior in classroom and clinic and of everyday behavior.

LEFRANCOIS, G. R. (1972). *Psychological Theories and Human Learning: Kongor's Report*. Brooks/Cole. An irreverent look at a Variety of traditional theories of learning. Shy on modern

learning theory. The man-from-outer-space viewpoint allows the author to poke fun at humans. Nice reading.

MARTIN, G. & PEAR, J. (1983). *Behavior Modification: What It Is and How to Do It*, 2nd ed. Prentice-Hall. The title tells it all. Includes a discussion of basic terms and concepts as well as "hands on" examples of behavior modification principles applied to everyday life.

PAVLOV, I. P. (1960). *Conditioned Reflexes*. Dover. An English translation of the original work in the study of classical conditioning by the Nobel-prize-winning scientist. A treat if you already have a basic understanding of the process.

PETERSON, L. B. (1975). *Learning*. Scott, Foresman. A brief review of classical and operant conditioning. Includes an extension of basic principles into more complex activities.

SKINNER, B. F. (1972). *Cumulative Record: A Selection of Papers*, 3rd ed. Prentice-Hall. A wide-ranging collection covering everything from critiques of other theories to the use of operant conditioning in teaching and learning.

SKINNER, B. F. (1948). *Walden II*. Macmillan. Can a society be run on the principles of reinforcement, with no use of punishment? Would children raised in such an environment be free from jealousy and envy, or would they lose their individuality? Skinner's program for a utopian society.