

CHAPTER 7

LEARNING AND CONDITIONING



If you have ever experienced a panic attack, you know that it is a terrifying experience: your heart is racing, you feel out of breath and perhaps even faint, and you are convinced that something terrible will happen. A panic attack can be thought of as an overreaction to a real or perceived threat in the environment. The symptoms are the result of the excitation of the sympathetic division of the autonomic nervous system (recall the ‘fight-or-flight’ response discussed in Chapter 2). Panic attacks are not at all uncommon, especially during times of stress: up to 40 percent of young adults have occasional panic attacks (see Chapter 15).

Far fewer individuals develop a panic disorder – in these cases the attacks are frequent and the intense worry about them interferes with everyday life. Research has shown that the most effective form of treatment for panic disorders is **cognitive behavior therapy** (see Chapter 16). This is a treatment method that involves procedures to change maladaptive cognitions and beliefs. Cognitive behavior therapy has its roots in **behavior therapy**, a general term referring to treatment methods based on the principles of learning and conditioning. The effectiveness of these forms of therapy suggests that some of the behaviors involved in panic disorders seem to be *learned* responses, which may be unlearned in the therapy.

Learning and conditioning are the topics of this chapter. We will engage in a systematic analysis of learning that will give you insight into how experience alters behavior. **Learning** is defined as a relatively permanent change in behavior that occurs as a result of experience. Behavior changes that are due to maturation or to temporary conditions (such as fatigue or drug-induced states) are not included.

Not all cases of learning are the same, though. There are two basic kinds of learning: non-associative learning and associative learning. **Non-associative learning** involves learning about a single stimulus, and it includes habituation and sensitization. **Habituation** is a type of non-associative learning that is characterized by a decreased behavioral response to an innocuous stimulus. For example, the sound of a horn might startle you when you first hear it. But if the horn toots repeatedly in a short time, the amount that you are startled by each sound progressively decreases. In contrast, **sensitization** is a type of non-associative learning whereby there is an increase in a behavioral response to an intense stimulus. Sensitization typically occurs when noxious or fearful stimuli are presented to an organism. For example, the acoustic startle response to a horn is greatly enhanced if you enter a dark alley right before the loud sound.

CHAPTER OUTLINE

PERSPECTIVES ON LEARNING

CLASSICAL CONDITIONING

Pavlov's experiments

Cognitive factors

Biological constraints

INSTRUMENTAL CONDITIONING

Skinner's experiments

Cognitive factors

Biological constraints

LEARNING AND COGNITION

Observational learning

Prior beliefs

CUTTING EDGE RESEARCH: MAP LEARNING IN LONDON'S TAXI DRIVERS

LEARNING AND THE BRAIN

Habituation and sensitization

Classical conditioning

Cellular basis of learning

LEARNING AND MOTIVATION

Arousal

From incentives to goals

Intrinsic motivation and learning

SEEING BOTH SIDES: WHAT ARE THE BASES OF SOCIAL LEARNING?

Both habituation and sensitization are typically relatively short-lived, lasting for minutes to hours. Although these types of learning are quite simple, they are exceptionally important for determining what an organism attends to in the world. Indeed, the fact that non-associative learning can be demonstrated in all animals, ranging from single-celled paramecia to humans, is a testament to the importance of this form of learning. We will revisit non-associative learning in the section on the brain and learning.

Associative learning is much more complicated than non-associative learning, because it involves learning relationships among events. It includes classical conditioning and instrumental conditioning. Classical and instrumental conditioning both involve forming associations – that is, learning that certain events go together. In this chapter, we will discuss these forms of learning in detail. In classical conditioning, an organism learns that

one event follows another. For example, a baby learns that the sight of a breast will be followed by the taste of milk. In instrumental conditioning, an organism learns that a response it makes will be followed by a particular consequence. For example, a young child learns that striking a sibling will be followed by disapproval from his or her parents.

Besides classical and instrumental conditioning, this chapter will cover a more complex form of learning: observational learning. For other forms of complex learning in humans, the role of memory and cognition are crucial – these are the topics of Chapters 8 and 9. We will also take a look at the neural basis of learning, referring back to concepts introduced in Chapter 2. Lastly, the importance of motivation for learning is briefly discussed – you will see that the topic of motivation is further explored in Chapter 10.

PERSPECTIVES ON LEARNING

Recall from Chapter 1 that three of the most important perspectives on psychology are the behaviorist, cognitive, and biological perspectives. As much as any area in psychology, the study of learning has involved all three of these perspectives.

Most of the early work on learning, particularly on conditioning, was done from a behaviorist perspective. During the early decades of the last century, especially in North America, this approach to the study of behavior took psychology by storm. The most important ‘spokesman’ for behaviorism was the American John Watson. A brief article he published in 1913, entitled ‘Psychology as the Behaviorist Views it’ is referred to as ‘the behavioristic manifesto’. His ideas were formulated in response to the writings by some of the ‘founding fathers’ of psychology, such as William James, E. B. Titchener and Wilhelm Wundt. William James was interested in topics like consciousness and emotion and Titchener was devoting his research to the study of mental structures. The German Wilhelm Wundt, as we saw in Chapter 1, was the first to establish a laboratory dedicated to the study of psychology. His method of inquiry was that of introspection. In Watson’s opinion, the methods of psychology were too subjective. Watson also argued that the subject matter of psychological research should not be consciousness, but rather *behavior*. He was inspired by animal studies carried out by the Russian Ivan Pavlov and believed that his experiments afforded psychologists with a scientific method of inquiry: objective and replicable.

For early behaviorists the focus was on external stimuli and observable responses, in keeping with the behavioristic dictum that behavior is better understood in terms of external causes than mental ones. The behaviorists’ approach to learning included other key assumptions as well. One was that simple associations of the classical or instrumental kind are the basic building blocks of *all* learning processes, regardless of what is being learned or who is doing the learning – a rat learning to run a maze or a child mastering arithmetic (Skinner, 1971, 1938). It follows that something as complex as acquiring a language is presumably a matter of learning many associations (Staats, 1968). These views led behaviorists to focus on how the behaviors of non-human organisms, particularly rats and pigeons, are influenced by rewards and punishments in simple laboratory situations.

The findings and phenomena uncovered in this work continue to form the basis for much of what we know about associative learning. But as we will see, the behavioristic assumptions have had to be modified in light of subsequent work. Understanding conditioning, not to mention complex learning, requires that we consider what the organism *knows* about the relations between stimuli and response (even if the organism is a rat or a pigeon). This brings in the cognitive perspective. Moreover, it now appears that no single set of laws underlies learning in all situations and by all organisms. In particular, different mechanisms of learning seem to be involved in different species, which brings in the biological perspective.

The discoveries described in this chapter set the stage for the ‘cognitive revolution’ in psychology, an intellectual movement in the 1950s championed by Jerome Bruner and others, who rejected the constraints of behaviorism (Bruner, 1997): they believed that mental representations are not only important topics in psychology, but that they can be studied using the scientific method as well. As described in Chapter 1, this movement was strengthened by the development of computers in the second half of the last century. This allowed researchers (for example, Nobel prize winner Herbert Simon) to simulate cognitive processes, ushering in a view of human beings as processors of information – rather than organisms that are simply conditioned to respond to external events.

It remains invaluable to study the work done by behaviorists. As you will discover in this chapter, their experimental paradigms and discoveries have laid the foundation for much of the research into human behavior that has been carried out since.

INTERIM SUMMARY

- Learning is a relatively permanent change in behavior that is the result of experience.
- There are four basic kinds of learning: (1) habituation and sensitization, (2) classical conditioning, (3) instrumental conditioning, and (4) complex learning.



CLASSICAL CONDITIONING


Ivan Pavlov, a Russian physiologist who had already received the Nobel Prize for his research on digestion, made an important discovery in the early years of the twentieth century. For his research, he was measuring

dogs’ salivation in response to food – any dog will salivate when food is placed in its mouth. But Pavlov noticed that the dogs in his laboratory began to salivate at the mere *sight* of a food dish. It occurred to him that the dogs had perhaps learned to associate the sight of the dish with the taste of the food, and he decided to see whether a dog could be taught to associate food with other stimuli, such as a light or a tone. The elegant experiments that Pavlov designed to study this question have contributed much to our understanding of one of the most basic processes of learning: classical conditioning (often referred to as ‘Pavlovian conditioning’). **Classical conditioning** is a learning process in which a previously neutral stimulus becomes associated with another stimulus through repeated pairing with that stimulus. The food dish was originally a *neutral* stimulus: it did not lead to a salivation response. However, the food itself *does* cause salivation when it is placed in the mouth of the dog. After food and food dish are presented together (‘paired’) repeatedly, the mere sight of the food dish is enough to cause a salivation response. The dog has learned that two events (the sight of a food dish, and the taste of food in the mouth) are *associated*.

In this section, you will be introduced to the vocabulary of classical conditioning through a presentation of Pavlov’s initial findings. Over the years, many psychologists have devised interesting variations of Pavlov’s experiments – we will also discuss some of these important and more recent discoveries.

Pavlov’s experiments

In Pavlov’s basic experiment, a tube is attached to the dog’s salivary gland so that the flow of salivation can be measured. Then the dog is placed in front of a pan into which meat powder can be delivered automatically. The dog is hungry and when meat powder is delivered, salivation is registered. This salivation is an **unconditioned**



Ivan Pavlov with his research assistants and one experimental subject (the dog).

response (UR): an unlearned response elicited by the taste of the food. By the same token, the food itself is termed the **unconditioned stimulus (US):** a stimulus that automatically elicits a response without prior conditioning. The researcher can also turn on a light in a window in front of the dog. This event is called a **neutral stimulus (NS)** because it does not cause salivation – though it may of course lead to other responses by the dog (such as tail wagging, jumping, and barking). Next, the researcher will repeatedly pair the presentation of the food with the light: first the light is turned on, then some meat powder is delivered and the light is turned off. This is called the **conditioning phase** of the experiment. After a number of such paired presentations, the dog will salivate in response to the light even if no meat powder is delivered. This teaches us that the dog has learned that the two events (food and light) are associated – the light has become a **conditioned stimulus (CS)**, causing a **conditioned response (CR)**. Figure 7.1 diagrams the different phases of Pavlov's conditioning experiment. In variations on this experiment, Pavlov used a tone (or other stimuli) instead of a light, and found similar results in each case.

In a classical conditioning experiment, the researcher capitalizes on the existence of a certain unconditioned response, typically a reflex – in our basic example the salivation. Such responses are part of the natural behavioral repertoire of the animal or human under study (for example: the eye blink in response to a puff of air on the eye, or a knee jerk reflex in response to a tap on the knee).

In Pavlov's experiments, the form of the conditioned response often mimicked the form of the unconditioned response – in our basic example it was salivation in both cases. In most cases, however, it is a bit more complicated than that. Note that, in our example, you might consider the salivation in response to the light (the CR) to be **anticipatory**: the dog salivates in response to the light, because it has learned that the light *precedes* the food. This anticipatory nature of the conditioned response explains why in some cases it takes on quite a different form from the unconditioned response. In this way, classical conditioning can help to explain the complex response humans have to the repeated intake of specific drugs.

Drug tolerance

Drug tolerance refers to the decreased effect of a drug when it is taken repeatedly. In other words, increased doses are required to produce the same effects that were initially produced with smaller doses. Research has shown that classical conditioning contributes to drug tolerance. These insights are important, not in the last place because drug tolerance is important in drug addiction.

Habitual coffee drinkers will develop a degree of tolerance to caffeine: with repeated intake, the effect of the caffeine (which is to raise blood pressure) is attenuated. Even though the coffee *originally* resulted in an increased blood pressure, it no longer does so after the coffee-drinking

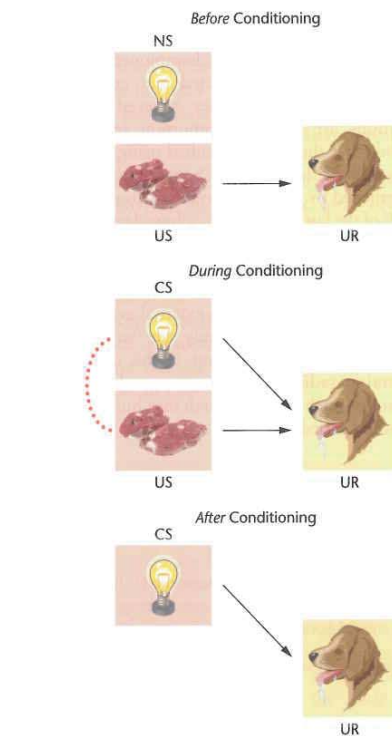


Figure 7.1 A diagram of Classical Conditioning. *Before conditioning, the unconditioned stimulus (US) causes the unconditioned response (UR) – this does not have to be learned. The neutral stimulus (NS) does not lead to a response. During conditioning, the unconditioned stimulus (US) and the conditioned stimulus (CS) are paired, and their association is learned. After conditioning, the conditioned stimulus (CS) causes the conditioned response (CR). In this example, both UR and CR are salivation.*

habit has formed. But when these same habitual coffee drinkers are given caffeine intravenously (injected directly into a vein), the original effect of the caffeine returns (Corti et al., 2002). It appears that drug tolerance is greater when the drug is taken under the usual circumstances. This effect is called the 'situational specificity of drug tolerance', and it can be explained by classical conditioning.

The intake of a drug will trigger a **compensatory response** of the body – recall our discussion of **homeostasis** in Chapter 2. When caffeine (the unconditioned stimulus, US) is consumed and blood pressure is raised (the unconditioned response, UR), the body responds to restore homeostasis by bringing the blood pressure back

down to its normal level. It turns out that when someone habitually drinks a cup of coffee, this compensatory response (the conditioned response, CR) will be elicited by cues related to the habitual caffeine intake (the conditioned stimulus, CS) – the smell of the coffee, for example. Classical conditioning explains how the body learned to respond to the situational cues (the CS) that are associated with regular caffeine intake, simply because of their repeated pairing with the caffeine intake (the US). In this way, classical conditioning explains how tolerance develops: the body's compensatory response (the CR) clearly contributes to tolerance for the drug. Another example is that of **alcohol tolerance**. Imagine someone who habitually drinks a few beers. It has been found that this person will show greater tolerance to the alcohol in a beer (the usual drink), than when the same amount of alcohol is consumed in another drink (Remington et al., 1997).

So, when a habitual user takes a drug under unusual circumstances (for example an injection of caffeine or alcohol in an unusual beverage), tolerance to the drug is reduced because the conditioned compensatory response is not triggered. This analysis explains the perplexing finding that most deaths due to an 'overdose' of a recreational drug (such as heroin or cocaine) are in fact *not* the result of an actual *overdose* (Siegel, 2001). It has been reported that, in most of these cases, the habitual user of the drug took *no more* than their normal dose of the drug – but rather, took it under unusual circumstances (for example, by injecting in a different part of the body, or in a different room than normally). The unusual circumstances deprived the user of the life-saving compensatory response,

thereby reducing tolerance to the drug and making it lethal.

Acquisition

We will return to Pavlov's original experiments to introduce a few more important aspects of learning through classical conditioning. Each paired presentation of the CS (light) followed by the US (food) is called a reinforced trial. Repeated pairings of the CS and the US strengthen the association between the two, as illustrated by the increase in the magnitude of the CR (the salivation response) in the left panel of Figure 7.2. This is the **acquisition** stage of the experiment, and the figure represents the **learning curve**. The largest change in the magnitude of the CR happens in the earliest conditioning trials, and there is little change in the CR later on.

Extinction

If the US is subsequently omitted, the CR will gradually diminish, as illustrated by the middle panel of Figure 7.2. As you see, after about ten trials or so there is no salivation in response to the light, if it is not followed by food. **Extinction** represents learning that the CS no longer predicts the US.

Spontaneous recovery

When the experimenter allows the dog to rest for a certain period, and then presents again *only* the light, the (extinguished) salivation response reappears – see right panel of Figure 7.2. This is called **spontaneous recovery**: no reinforced trials are needed, and the CS again leads to a CR. As you can see, the recovered CR is weaker than it

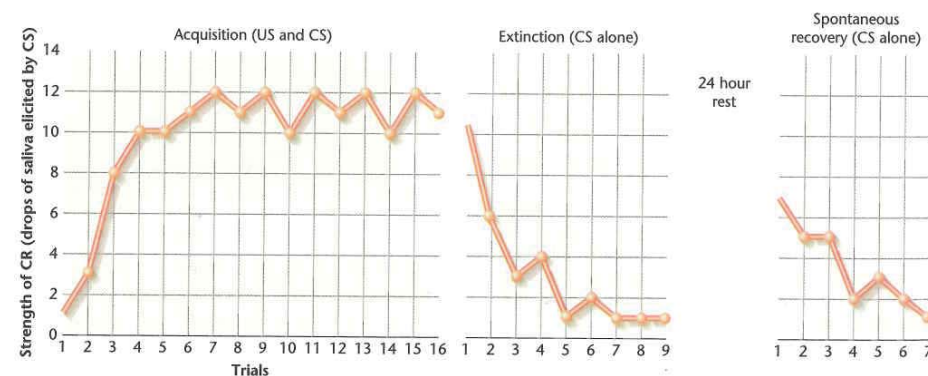


Figure 7.2 Acquisition and Extinction of a Conditioned Response. *The curve in the panel on the left depicts the acquisition phase of an experiment. Drops of saliva in response to the CS (before the onset of the US) are plotted on the vertical axis; the number of trials is plotted on the horizontal axis. After 16 acquisition trials, the experimenter switched to extinction; the results are presented in the panel in the middle. The panel on the right shows spontaneous recovery of the response after a 24 hour rest period. (Adapted from Conditioned Reflexes, by E. P. Pavlov. Copyright © 1927 by Oxford University Press. Reprinted by permission of Oxford University Press.)*

was after acquisition. With repeated presentation of the CS alone, the CR will again diminish. Spontaneous recovery reflects that the association between the CS and the US that was originally learned, does not simply disappear during extinction. Rather, extinction seems to involve the formation of a *new* association (between CS and *no* US). The spontaneous recovery of the CR means that the dog 'remembers' that the light used to predict food – even though the response itself was completely extinguished.

Extinction can also be undone by reinforcing the original association through repeated pairing of the CS and the US, as it was originally done during acquisition. The re-learning curve would be steeper than the learning curve presented in the left panel of Figure 7.2 (re-learning an association is faster than originally learning it). This suggests again that the association between the CS and the US was not forgotten, even though the CR was extinguished. Consider again the example of our habitual coffee drinker: the smell of coffee (the CS) causes the compensatory response to decrease blood pressure (the CR). This compensatory response will eventually be extinguished if the coffee drinker switches to decaffeinated coffee, which constitutes the presentation of the CS *in the absence* of the US (the caffeine). But when this person switches back to drinking regular coffee, the body will respond by quickly re-learning the old association.

Stimulus generalization

Pavlov noticed that the dogs that had been trained to have a conditioned response to a certain tone, would show the same response to a tone that was slightly higher or lower in pitch. This is called **response generalization**: the more similar the new stimuli are to the original CS, the more likely they are to evoke the same response. Suppose that a person is conditioned to have a mild emotional reaction to the sound of a tuning fork producing a tone of middle C. This emotional reaction can be measured by the *galvanic skin response*, or GSR, which is a change in the electrical activity of the skin that occurs during emotional stress. That person will show a change in GSR in response to higher or lower tones without further conditioning (see Figure 7.3).

Stimulus generalization accounts in part for a human or animal's ability to react to novel stimuli that are similar to familiar ones – an ability that is clearly adaptive. Organisms might not be exposed to exactly the same stimulus very often, but similar stimuli are likely to predict similar events.

Stimulus discrimination

A process that is complementary to generalization is discrimination. Stimulus generalization is a reaction to similarities, and **stimulus discrimination** is a reaction to differences. Conditioned discrimination is brought about through differential conditioning, as shown in Figure 7.4.

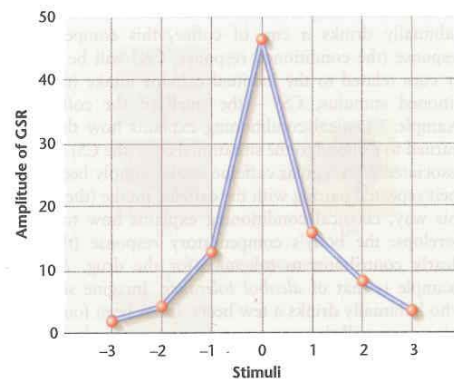


Figure 7.3 The Gradient of Generalization. Stimulus 0 denotes the tone to which the galvanic skin response (GSR) was originally conditioned. Stimuli +1, +2, and +3 represent test tones of increasingly higher pitch; stimuli -1, -2, and -3 represent tones of lower pitch. Note that the amount of generalization decreases as the difference between the test tone and the training tone increases. ("The Sensory Generalization of Conditioned Responses with Varying Frequencies of Tone," from *Journal of General Psychology*, Vol. 17, p. 125–148, 1937. Reprinted by permission of the Helen Dwight Reid Educational Foundation.)

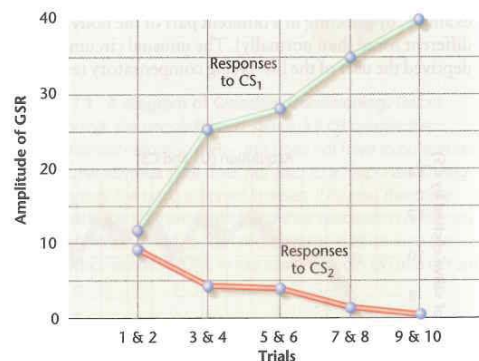


Figure 7.4 Conditioned Discrimination. The discriminative stimuli were two tones of clearly different pitch (CS₁ = 700 Hertz and CS₂ = 3,500 Hertz). The unconditioned stimulus, an electric shock applied to the left forefinger, occurred only on trials when CS₁ was presented. The strength of the conditioned response, in this case the GSR, gradually increased following CS₁ and extinguished following CS₂. (Adapted from "Differential Classical Conditioning: Verbalization of Stimulus Contingencies," by M. J. Fuhrer & P. E. Baer, reprinted by permission from *Science*, Vol. 150, December 10, 1965, pp. 1479–1481. Copyright © 1965 by American Association for the Advancement of Science.)

Instead of just one tone during conditioning, now there are two. The low-pitched tone, CS₁, is always followed by a mild forefinger shock, and the high-pitched tone, CS₂, is not. Initially, participants show a GSR to both tones. During the course of conditioning, however, the amplitude of the conditioned response to CS₁ gradually increases while the amplitude of the response to CS₂ decreases. Through this process of *differential reinforcement*, participants are conditioned to discriminate between the two tones. It is important to note that the presentation of CS₂ leads to a suppression of the response (lowered GSR). This is because its presentation contains information for the subject, namely that *no shock* will follow. Most of the examples of conditioning we discussed thus far were examples of *excitatory conditioning*, in which case the CS leads to an increase in the probability or magnitude of a certain response. But differential reinforcement teaches us that another possible consequence of classical conditioning is a decrease in the probability or magnitude of a behavioral response – this is *inhibitory conditioning*.

Generalization and discrimination occur frequently in everyday life. A young child who has learned to associate the sight of her pet dog with playfulness may initially approach all dogs. Eventually, through discrimination, the child may expect playfulness only from dogs that look like hers. The sight of a threatening dog has come to inhibit the child's response to approaching dogs.

Second-order conditioning

Once a dog has been conditioned to salivate in response to a light, it is possible to condition the dog to salivate in response to another stimulus (for example, a tone), simply by repeatedly pairing the light and the tone. This is called **second-order conditioning**. In other words, once the light has taken on the role of a conditioned stimulus, it acquires the power of an unconditioned stimulus. If the dog is now put in a situation in which it is exposed to a tone (CS₂) followed by the light (CS₁), the tone alone will eventually elicit the conditioned response – even though it was never paired with food. During this conditioning there must also be trials that reinforce the association between the light and the food; otherwise, the originally conditioned association will be extinguished.

The existence of second-order conditioning greatly increases the scope of classical conditioning. Especially in humans, most conditioned responses are established through second-order conditioning. The original US is usually a biologically significant stimulus, such as food, pain or nausea. All that is needed for conditioning to occur is the pairing of that stimulus with another. Consider the plight of cancer patients who are undergoing chemotherapy to stop the growth of their tumors. Chemotherapy involves injecting toxic substances (the US) into the patient, who as a result often becomes nauseated

(the UR). Young cancer patients are often given ice cream before the chemotherapy session. The ice cream is intended to lighten the child's distress about the treatment, but unfortunately it becomes associated with it. The ice cream can take on the role of a CS and cause nausea by itself (Bernstein, 1978, 1999). If the child is then repeatedly presented with other stimuli, such as certain toys, followed by ice-cream, the patient may start to experience unpleasant feelings in response to the toys alone. This would be a consequence of second-order conditioning, since the toys were never directly paired with treatment or nausea.

Conditioning and fear

Classical conditioning also plays a role in emotional responses like fear. Suppose that a rat in an enclosed compartment is periodically subjected to electric shock. Just before the shock occurs, a tone sounds. After repeated pairings of the tone (the CS) and the shock (the US), the tone alone will produce reactions in the rat that indicate fear, including freezing and crouching. In addition, its blood pressure increases. The rat has been conditioned to be fearful when exposed to what was previously a neutral stimulus. Humans, too, can be conditioned to be fearful (Jacobs & Nadel, 1985; Watson & Rayner, 1920). Indeed, classical conditioning of fear seems to be at the root of several anxiety disorders, such as post-traumatic stress disorder and panic disorder (Bouton, Mineka, & Barlow, 2001).

We have seen repeatedly that a conditioned stimulus leads to a conditioned response, precisely because it *predicts* the occurrence of a certain unconditioned stimulus. Predictability is also important for emotional reactions. If a particular CS reliably predicts that pain is coming, the *absence* of that CS predicts that pain is *not* coming so that the organism can relax. The CS has become a 'danger' signal, and its absence a 'safety' signal. When such signals are erratic, the emotional toll on the organism can be devastating. When rats have a reliable predictor that shock is coming, they respond with fear only when the danger signal is present; if they have no reliable predictor, they appear to be continually anxious and may even develop ulcers (Seligman, 1975). There are clear parallels to human emotionality. If a dentist gives a child a danger signal by saying that a procedure will hurt, the child will be fearful until the procedure is over. In contrast, if the dentist always tells a child that it won't hurt, when in fact it sometimes does, the child has no danger or safety signals and may become terribly anxious whenever in the dentist's office. As adults, many of us have experienced the anxiety of being in a situation where something disagreeable is likely to happen but no warnings exist for us to predict it. Unpleasant events are, by definition, unpleasant, but unpredictable unpleasant events are downright intolerable (see also Chapter 14).

Cognitive factors

Pavlov and others believed that it was enough for conditioning to occur if the CS and the US were **temporally contiguous** – that is, the CS and the US occur close together in time. Pavlov was careful not to make any claims about the organism's cognitive understanding of relationships between stimuli; such internal events were considered not to be observable. From our previous discussion, however, it would seem that conditioning occurs if the CS **predicts** the US. In such cases, we say that the US is **contingent** on the CS (the US is more likely to occur when the CS is presented, than when it is not presented). Some researchers indeed argued that the critical factor behind classical conditioning is what the animal knows (Bolles, 1972; Tolman, 1932). In this cognitive view, classical conditioning gives an organism new knowledge about the relationship between two stimuli: given the CS, the organism has learned to expect the US (Rescorla, 1968).

In a series of important and elegantly designed experiments, Rescorla (1968) contrasted contiguity and contingency. He was able to show that the CS must be a reliable predictor of the US. Mere temporal contiguity is *not* enough for conditioning to occur. The procedure for one of these experiments is depicted in Figure 7.5. There are two groups of rats, group A and B. The number of temporally contiguous pairings of tone and shock was the same in both groups. So, if *temporal contiguity* determines conditioning, both groups of rats should show equal amounts of conditioning. What was different however, was the contingency of the shock on the tone: for group A all shocks were preceded by tones, whereas for group B shocks were equally likely in the presence and absence of

the tone. Therefore, the tone was highly predictive of the shock for group A, but it had no predictive power for group B. So, if *contingency* determines conditioning, we would expect only group A to exhibit conditioning. And this is exactly what Rescorla found: only the rats in group A developed a conditioned fear response. In other groups in the experiment (not shown in Figure 7.5), the strength of the conditioning was directly related to the predictive value of the CS in signaling the occurrence of the US. Subsequent experiments supported the conclusion that the predictive relationship between the CS and the US is more important than either temporal contiguity or the frequency with which the CS and US are paired (Rescorla, 1972).

Biological constraints

Early behaviorists assumed that the laws of learning were the same for all species. Moreover, they assumed that *any* CS could be associated with *any* US through classical conditioning.

This doctrine places these early behaviorists firmly on the *nurture* side of the nature–nurture debate: what an organism learns, depends entirely on its experiences with the environment. Others, however, had emphasized the biological function of the learning process: it allows the organism to adapt and survive. Early ethologists (for example, European Nobel Prize winners Konrad Lorenz, Nikolaas Tinbergen, and Karl von Frisch) made discoveries that revealed powerful *biological predispositions* in human and animal behavior (Tinbergen, 1951).

Ethologists, like behaviorists, are concerned with the behavior of animals, but place greater emphasis on

evolution and genetics – and they study the behavior of animals in their natural environment. This perspective on learning draws attention to the fact that exactly *what* an organism needs to learn depends on its evolutionary history – to some extent animals are ‘pre-programmed’ to learn particular things in particular ways.

Consider the example of a **learned taste aversion**. Many of us have had the experience of becoming ill after eating a certain food, and would not want to eat that particular food ever again. Garb and Stunkard (1974) found that over one-third of people have had at least one such experience. Typically, a novel food was eaten and the person got ill (nausea and vomiting) within a few hours. Learned taste aversions at first seem typical instances of classical conditioning: the taste of the food has become associated with the illness. However, upon closer inspection, the conditioning does not entirely comply with the rules of classical conditioning. First of all, most taste aversions occur after just one bad experience with the food – no repeated pairings are necessary. Secondly, the CS–US interval is usually very long: the illness (the US) occurs a few hours after the ingestion of the food (the CS). From an evolutionary perspective, it is very easy to see what is adaptive about the ability of an organism to be able to learn to avoid particular foods in a single trial: the organism will avoid food that is potentially harmful. The existence of learned taste aversions shows that organisms are very selective in what they are able to learn: certain associations are learned very readily, while others may never be learned.

Garcia and Koelling (1966) carried out a series of controlled experiments that reveal the importance of biological predispositions in learning. One of their experiments is diagrammed in Table 7.1. In the first stage of the experiment, an experimental group of rats is allowed to lick at a tube that contains a flavored solution. Each time the rat licks the tube, a click and a light are presented. The rat experiences three stimuli simultaneously – the taste of the solution, as well as the light and the click. In the second stage of the experiment, rats in the experimental group are mildly poisoned with lithium

chloride. Which stimuli – the sweet taste or the light-plus-click – will become associated with feeling sick? To answer this question, in the third and final stage, rats in the experimental group are again presented with the tube. Sometimes the solution in the tube has the same flavor as before but there is no light or click, and at other times the solution has no flavor but the light and click are presented. The animals avoid the solution when they experience the taste, but not when the light-plus-click is presented. Therefore, the rats have associated only taste with feeling sick. These results cannot be attributed to taste being a more potent CS than light-plus-click, as shown by the control condition of the experiment, which is diagrammed at the bottom of Table 7.1. In the second stage, instead of being mildly poisoned, the rat is shocked. In the final stage, the animal avoids the solution only when the light-plus-click is presented, not when it experiences the taste alone (Garcia & Koelling, 1966).

So, taste is a better signal for sickness than for shock, and light-plus-click is a better signal for shock than for sickness. Why does this selectivity of association exist? It does not fit with the early behaviorist idea that equally potent stimuli can be substituted for one another. Because taste and light-plus-click can both be effective conditioned stimuli, and being sick and being shocked are both effective unconditioned stimuli, it should have been possible for either CS to become associated with either US. On the other hand, selectivity of association fits perfectly with the ethological perspective and its emphasis on an animal's evolutionary adaptation to its environment. In their natural habitat, rats rely on taste to select their food. Consequently, there may be a genetically determined relationship between taste and intestinal reactions that fosters an association between taste and sickness but not between light and sickness. Moreover, in a rat's natural environment, pain resulting from external factors like cold or injury is invariably due to external stimuli. As a result, there may be a built-in relationship between external stimuli and ‘external pain’, which fosters an association between light and shock but not one between taste and shock.




Figure 7.5 Rescorla's Experiment. For each group, the events for 16 trials are presented. On some trials the CS occurs and is followed by the US (CS + US); on other trials the CS or US occurs alone; and on still other trials, neither the CS nor the US occurs. The boxes to the far right give a count of these trial outcomes (CS+US, CS only, US only, or neither stimulus) for the two groups. The number of CS + US trials is identical for both groups, as is the number of trials on which only the CS occurs. But the two groups differ in the number of trials on which the US occurred alone (never in Group A and as frequently as any other type of trial in Group B). A conditioned response to CS developed readily for Group A but did not develop at all for Group B. (R. A. Rescorla (1967) "Pavlovian Conditioning & Its Proper Control Procedures," from *Psychological Review*, Vol. 74:71–80. Copyright © 1967 by the American Psychological Association.)

Table 7.1

An experiment on constraints and taste aversion The design of an experiment showing that taste is a better signal for sickness than shock, whereas light-plus-sound is a better signal for shock than sickness. (J. Garcia and R. A. Koelling (1966) 'The Relation of Cue to Consequence in Avoidance Learning,' *Psychonomic Science*, 4: 123–124. Reprinted by permission of the Psychonomic Society.)

Condition	Conditioned stimuli (CS)	Unconditioned stimulus (US)	Result
Poison	Sweet taste; light + click	Lithium chloride	Taste → suppression of drinking Light + click → no suppression of drinking
Shock	Sweet taste; light + click	Footshock	Taste → no suppression Light + click → suppression of drinking

If rats learn to associate taste with sickness because it fits with their natural means of selecting food, another species with a different means of selecting food might have trouble learning to associate taste with sickness. This is exactly what happens. Birds naturally select their food on the basis of looks rather than taste, and they readily learn to associate a light with sickness but not to associate a taste with sickness (Wilcoxon, Dragoin, & Kral, 1971). Here, then, is a perfect example of different species learning the same thing – what causes sickness – by different means. In short, if we want to know what may be conditioned to what, we cannot consider the CS and US in isolation. Rather, we must focus on the two in combination and consider how well that combination reflects built-in relationships. This conclusion differs considerably from the assumption that the laws of learning are the same for all species and situations. In fact, several recent theorists have explored classical conditioning by using a behavior systems approach that considers the evolutionary history of the behaviors under study (Fanselow, 1994).

INTERIM SUMMARY

- In classical conditioning, a conditioned stimulus (CS) that consistently precedes an unconditioned stimulus (US) comes to serve as a signal for the US and will elicit a conditioned response (CR) that often resembles the unconditioned response (UR).
- For classical conditioning to occur, the CS must be a reliable predictor of the US; that is, there must be a higher probability that the US will occur when the CS has been presented than when it has not.
- The ability of stimuli to become associated in a classical conditioning experiment is constrained by biology and evolution.




B. F. Skinner was a pioneer in the study of instrumental conditioning.

INSTRUMENTAL CONDITIONING

In classical conditioning, the conditioned response is a response that was part of the animal's natural repertoire – like salivation. But how do dogs learn *new* 'tricks', like rolling over and playing dead? If you have ever trained a dog to perform such tricks, you know that it involves rewarding the dog whenever it does what you want it to do. Initially, you will reward the dog for approximating the desired behavior, but eventually you will only reward it if it performs the entire trick. In **instrumental conditioning**, certain behaviors are learned because they *operate* on the environment. Your dog learns that performing the trick results in food: the behavior is *instrumental* in producing a certain change in the environment. If we think of the dog as having food as a goal, instrumental conditioning (which is also called operant conditioning) amounts to learning that a particular behavior (called the '*response*' – in this case rolling over) leads to a particular goal (Rescorla, 1987). Classical conditioning involves learning the relationship between events; instrumental conditioning (also called '*operant conditioning*') involves learning the relationship between responses and their outcomes.

In this section, we will review the findings of B. F. Skinner, an American psychologist who contributed much to our understanding of instrumental conditioning.



NINA LEEB, © LIFE MAGAZINE/TIME PIX

By the 1950s, Skinner was the leading proponent of behaviorism in the United States. As before, we will also discuss more recent discoveries and insights.

The study of instrumental conditioning did not begin with Skinner's work. E. L. Thorndike carried out a series of important experiments at the turn of the twentieth century (Thorndike, 1898). He was inspired by the writings of Charles Darwin, which contained many anecdotes about animals revealing seemingly intelligent and insightful behavior. But Thorndike felt that, to study animal intelligence, controlled experiments should be carried out. From his experiments, Thorndike concluded that animals, unlike humans, do not learn by developing some **insight** (an understanding of the situation, leading to the solution of a problem) – rather, they learn through trial-and-error. In a typical experiment, a hungry cat is placed in a cage whose door is held fast by a simple latch, and a piece of fish is placed just outside the cage. Initially, the cat tries to reach the food by extending its paws through the bars. When this fails, the cat moves about the cage, engaging in a variety of behaviors. At some point it inadvertently hits the latch, frees itself, and eats the fish. Researchers then place the cat back in its cage and put a new piece of fish outside. The cat goes through roughly the same set of behaviors until once more it happens to hit the latch. The procedure is repeated again and again. Over a number of trials, the cat eliminates many of its irrelevant behaviors, and eventually it opens the latch and frees itself as soon as it is placed in the cage. The cat has learned to open the latch to obtain food.

It may sound as if the cat is acting intelligently, but Thorndike argued that there is little 'intelligence' operating here. There is no moment in time when the cat seems to have an insight about the solution to its problem. Instead, the cat's performance improves gradually over a series of trials. The cat appears to be engaging in **trial-and-error learning**, and when a reward immediately follows one of those behaviors, the learning of the action is strengthened. Thorndike referred to this strengthening as the **law of effect**. He argued that in instrumental learning, the law of effect selects from a set of random responses only those that are followed by positive consequences.

Skinner's experiments

Skinner's method of studying instrumental conditioning was simpler than Thorndike's: he studied only one response at a time. In a Skinnerian experiment, a hungry animal – usually a rat or a pigeon – is placed in a box like the one shown in Figure 7.6, which is called an **operant chamber** (also referred to as a *Skinner box*). The inside of the box is bare except for a protruding bar with a food dish beneath it. A small light above the bar can be turned on at the experimenter's discretion. Left alone in the box,



Figure 7.6 Apparatus for Instrumental Conditioning: the Operant Chamber. This photograph shows an operant chamber (often called a 'Skinner box') with a magazine for delivering food pellets. The computer is used to control the experiment and record the rat's responses.

the rat moves about, exploring. Occasionally it inspects the bar and presses it. The rate at which the rat first presses the bar is the **baseline level**.

Acquisition and extinction

After establishing the baseline level, the experimenter activates a food magazine located outside the box. Now, every time the rat presses the bar, a small food pellet is released into the dish. The rat eats the food pellet and soon presses the bar again. The food reinforces bar pressing, and the rate of pressing increases dramatically. If the food magazine is disconnected and pressing the bar no longer delivers food, the rate of bar pressing diminishes. An instrumental response that is not reinforced undergoes extinction, just as a classically conditioned response does.

Instrumental conditioning increases the likelihood of a response by following the behavior with a reinforcer (often something like food or water). Because the bar is always present in the Skinner box, the rat can respond to it as frequently or as infrequently as it chooses. The organism's rate of response is therefore a useful measure of the instrumental learning; the more frequently the response occurs during a given time interval, the greater the learning.

Reinforcement versus punishment

In instrumental conditioning, an environmental event that follows behavior produces either an increase or a decrease in the probability of that behavior. **Reinforcement** refers to the process whereby the delivery of a stimulus *increases* the probability of a behavior. Reinforcement can be done by giving an appetitive stimulus (**positive reinforcement**) or by the removal of an aversive stimulus (**negative reinforcement**). In other words: there may be either a *positive* or a *negative* contingency between the behavior and reinforcement. A positive contingency means that something is given: for example, bar pressing is followed by food. A negative contingency means that something is taken away: for example, bar pressing terminates or prevents shock. **Punishment** is the converse of reinforcement: it *decreases* the probability of a behavior, and consists of the delivery of an aversive stimulus (**positive punishment**, or simply 'punishment') or the removal of an appetitive stimulus (**negative punishment** or 'omission training'). Again, note that there may be either a positive contingency between the behavior and punishment (bar pressing is followed by shock) or a negative contingency (bar pressing terminates or prevents food delivery). (See the Concept Review Table.)

Although rats and pigeons have been the favored experimental subjects, instrumental conditioning applies to many species, including our own. Indeed, instrumental conditioning has a good deal to tell us about child rearing. A particularly illuminating example is the following case. A young boy had temper tantrums if he did not get enough attention from his parents, especially at bedtime. Because

the parents eventually responded to the tantrums, their attention probably reinforced the boy's behavior. To eliminate the tantrums, the parents were advised to go through the normal bedtime ritual and then ignore the child's protests, painful though that might be. If the reinforcer (attention) was withheld, the behavior should be extinguished – which is just what happened. The time the child spent crying at bedtime decreased from 45 minutes to not at all over a period of only seven days (Williams, 1959). This is an example of omission training because withholding something the boy wanted (parental attention) decreased the behavioral response (bedtime crying).

Shaping

Suppose that you want to use instrumental conditioning to teach your dog a trick – for instance, to get the mail from the slot in your front door. You cannot wait until the dog does this naturally and then reinforce it, because you may wait forever. When the desired behavior is truly novel, you have to condition it by taking advantage of natural variations in the animal's actions. To train a dog to get the mail, you can give the animal a food reinforcer each time it approaches the door, requiring it to move closer and closer to the mail for each reinforcer until finally the dog grabs the mail. This technique, called **shaping**, is reinforcing only variations in response that deviate in the direction desired by the experimenter. Animals can be taught elaborate tricks and routines by means of shaping. Two psychologists and their staff trained thousands of animals of many species for television shows, commercials, and county fairs

CONCEPT REVIEW TABLE

Types of reinforcement and punishment

Type	Definition	Effect	Example
Positive reinforcement	Delivery of a pleasant or appetitive stimulus following a behavioral response	Increases the frequency of the behavioral response	If studying is followed by a high grade on an exam, then the incidence of studying before exams will increase
Negative reinforcement	Removal of an unpleasant or aversive stimulus following a behavioral response	Increases the frequency of the behavioral response	If leaving a study area removes you from a noisy classmate, then the time you spend away from the study area will increase
Positive punishment ('Punishment')	Delivery of an unpleasant or aversive stimulus following a behavioral response	Decreases the frequency of the behavioral response	If your professor embarrasses you for asking a question in class, then the likelihood you will ask questions in class will decrease
Negative punishment ('Omission training')	Removal of a pleasant or appetitive stimulus following a behavioral response	Decreases the frequency of the behavioral response	If your girl- or boyfriend withholds affection whenever you watch TV, the time you spend in front of the TV will decrease



Figure 7.7 Search and Rescue by Pigeons. The Coast Guard has used pigeons to search for people lost at sea. Shaping methods are used to train the pigeons to spot the color orange, the international color for life jackets. Three pigeons are strapped into a Plexiglas chamber attached to the underside of a helicopter. The chamber is divided into thirds so that each bird faces in a different direction. When a pigeon spots an orange object, or any other object, it pecks a key that buzzes the pilot. The pilot then heads in the direction indicated by the bird that responded. Pigeons are better suited than people for the task of spotting distant objects at sea. They can stare over the water for a long time without suffering eye fatigue, they have excellent color vision, and they can focus on a 60- to 80-degree area, whereas a person can focus only on a 2- to 3-degree area. (After Simmons, 1981)

(Breland & Breland, 1966). One popular show featured 'Priscilla, the Fastidious Pig'. Priscilla turned on the TV set, ate breakfast at a table, picked up dirty clothes and put them in a hamper, vacuumed the floor, picked out her favorite food, and took part in a quiz program by answering questions from the audience by flashing lights that indicated yes or no. She was not an unusually bright pig; in fact, because pigs grow so fast, a new 'Priscilla' was trained every three to five months. The ingenuity was not the pig's but the experimenters', who used instrumental conditioning and shaped the pig's behavior to produce the desired result. Shaping has been used to train pigeons to locate people lost at sea (see Figure 7.7), and porpoises have been trained to retrieve underwater equipment.

Importantly, the Brelands' work also indicated that not *all* behaviors could be shaped. For example, they had great difficulty training raccoons to drop coins into a piggy bank to receive a food reward. Rather than drop the coins in the bank to obtain a food reinforcer, the raccoons would rub them together incessantly, drop them in the bank, pull them out again, and continue rubbing them together. This behavior, of course, resembles the behavior that raccoons normally display to natural food items. The behavioral predisposition of the raccoon to vigorously manipulate an object associated with food made it difficult to shape a novel response. The phenomenon of animals resorting to biologically natural behaviors is called *instinctive drift*. It reveals that instrumental conditioning, like classical conditioning, operates under biological constraints.

Conditioned reinforcers

Most of the reinforcers we have discussed are called *primary* because they satisfy basic drives. If instrumental

conditioning occurred only with primary reinforcers, it would not occur very often because primary reinforcers are not that common. However, virtually any stimulus can become a *secondary* or **conditioned reinforcer**, which is a stimulus that has been consistently paired with a primary reinforcer. Conditioned reinforcers greatly increase the generality of instrumental conditioning. A minor variation in the typical instrumental conditioning experiment illustrates how conditioned reinforcement works. When a rat in a Skinner box presses a lever, a tone sounds momentarily and is followed shortly by delivery of food (the food is a primary reinforcer; the tone will become a conditioned reinforcer). After the animal has been conditioned in this way, the experimenter begins the extinction process, so that when the rat presses the lever, neither the tone nor the food occurs. In time, the animal ceases to press the lever. Then the tone is reconnected but not the food magazine. When the animal discovers that pressing the lever turns on the tone, its rate of pressing increases markedly, overcoming the extinction even though no food is delivered. The tone has acquired a reinforcing quality of its own through classical conditioning. Because the tone was reliably paired with food, it came to signal food. Secondary reinforcers apply to human behavior as well: our lives abound with conditioned reinforcers. Two of the most prevalent are money and praise. Presumably, money is a powerful reinforcer because it has been paired so frequently with so many primary reinforcers – we can buy food, drink, and comfort, to mention just a few of the obvious things. And mere praise can sustain many activities without even the promise of a primary reinforcer.



Praise is an effective reinforcer for many people.

Generalization and discrimination

Again, what was true for classical conditioning holds for instrumental conditioning as well: Organisms generalize what they have learned, and generalization can be curbed by discrimination training. If a young child is reinforced by her parents for petting the family dog, she will soon generalize this petting response to other dogs. Because this can be dangerous (the neighbors might have a vicious watchdog), the child's parents may provide some discrimination training so that she is reinforced when she pets the family dog but not the neighbor's.

Discrimination training will be effective to the extent that there is a discriminative stimulus (or a set of them) that clearly distinguishes cases in which the response should be made from those in which it should be suppressed. Our young child will have an easier time learning which dog to pet if her parents can point to an aspect of dogs that signals friendliness (a wagging tail, for example). In general, a discriminative stimulus will be useful to the extent that its presence predicts that a response will be followed by reinforcement and its absence predicts that the response will not be followed by reinforcement (or

vice versa). Just as in classical conditioning, the predictive power of a stimulus seems to be critical to conditioning.

Schedules of reinforcement

In real life, not every instance of a behavior is reinforced. For example, hard work is sometimes followed by praise, but often it goes unacknowledged. If instrumental conditioning occurred only with continuous reinforcement, it might play a limited role in our lives. Once a behavior is established, however, it can be maintained when it is reinforced only a fraction of the time. This phenomenon, *partial reinforcement*, can be illustrated in the laboratory by a pigeon that learns to peck at a key for food. Once this instrumental response is established, the pigeon continues to peck at a high rate, even if it receives only occasional reinforcement. In some cases, pigeons that were rewarded with food an average of once every five minutes (12 times an hour) pecked at the key as often as 6,000 times per hour – 500 pecks per pellet of food received!

Moreover, extinction following the maintenance of a response on partial reinforcement is much slower than extinction following the maintenance of a response on continuous reinforcement. Extinction of pecking in pigeons reinforced every five minutes takes days, whereas pigeons reinforced continuously extinguish in a matter of minutes. This phenomenon is known as the *partial-reinforcement effect*. It makes intuitive sense because there is less difference between extinction and maintenance when reinforcement during maintenance is only partial.

When reinforcement occurs only some of the time, we need to know exactly how it is scheduled – after every third response? After every five seconds? It turns out that the schedule of reinforcement determines the pattern of responding. There are four basic schedules of reinforcement (see the Concept Review Table).

Some schedules are called **ratio schedules**, because reinforcement depends on the number of responses the organism makes. It's like being a factory worker who gets paid per piece of work finished. The ratio can be either

CONCEPT REVIEW TABLE

Schedules of Reinforcement

Ratio schedules

Fixed Fixed ratio (FR): Reinforcement is provided after a fixed number of responses

Variable Variable ratio (VR): Reinforcement is provided after a certain number of responses, with the number varying unpredictably

Interval schedules

Fixed Fixed interval (FI): Reinforcement is provided after a certain amount of time has elapsed since the last reinforcement

Variable Variable interval (VI): Reinforcement is provided after a certain amount of time has elapsed since the last reinforcement, with the duration of the interval varying unpredictably

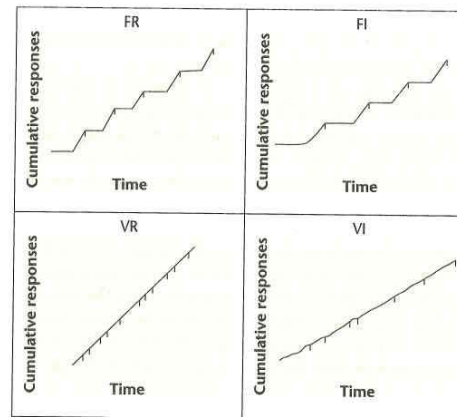


Figure 7.8 Typical Patterns of Responding on the Four Basic Schedules of Reinforcement. Each curve plots an animal's cumulative number of responses as a function of time; the slope of the curve thus indicates the animal's rate of responding. The short tick marks on each line indicate the moment reinforcement occurred. In the curve for the FR schedule, note the horizontal segments, which correspond to pauses (they show no increase in the cumulative number of responses). In the curve for the FI schedule, note again that the horizontal segments correspond to pauses. (Adapted from Barry Schwartz, *Psychology of Learning and Behavior*, 3/e, with the permission of W. W. Norton & Co., Inc.)

fixed or variable. On a **fixed ratio schedule** (called an FR schedule), the number of responses that have to be made is fixed at a particular value. If the number is 5 (FR 5), 5 responses are required for reinforcement; if it is 50 (FR 50), 50 responses are required; and so on. In general, the higher the ratio, the higher the rate at which the organism responds, particularly when the organism is initially trained on a relatively low ratio (say, FR 5) and then is continuously shifted to progressively higher ratios, culminating, say, in FR 100. It is as if our factory worker initially got \$5 for every 5 hems sewn, but then times got tough and he needed to do 100 hems to get \$5. But perhaps the most distinctive aspect about behavior under an FR schedule is the pause in responding right after the reinforcement occurs (see Figure 7.8). It is hard for the factory worker to start on a new set of hems right after he has just finished enough to obtain a reward.

On a **variable ratio schedule** (a VR schedule), the organism is still reinforced only after making a certain number of responses, but that number varies unpredictably. In a VR 5 schedule, the number of responses needed for reinforcement may sometimes be 1, at other times 10, with an average of 5. Unlike the behavior that occurs under FR schedules, there are no pauses when the organism is operating under a VR schedule (see Figure 7.8),



Gamblers who play the slot machines are reinforced with payoffs on a variable ratio schedule. Such a schedule can generate very high rates of responding.

presumably because the organism has no way of detecting how far it is from a reinforcement. A good example of a VR schedule in everyday life is the operation of a slot machine. The number of responses (plays) needed for reinforcement (payoff) keeps varying, and the operator has no way of predicting when reinforcement will occur. Of a schedules of reinforcement, VR schedules can generate the highest rates of responding (as casino owners appear to have figured out).

Other schedules of reinforcement are called **interval schedules**, because under these schedules reinforcement is available only after a certain time interval has elapsed (and the animal makes a response). Again, the schedule can be either fixed or variable. On a **fixed interval schedule** (an FI schedule), the organism is reinforced for its first response after a certain amount of time has passed since its last reinforcement. On an FI 2 (minutes) schedule, for example, reinforcement is available only when 2 minutes have elapsed since the last reinforced response;

responses made during that 2-minute interval have no effect. One distinctive aspect of responding on an FI schedule is a pause that occurs immediately after reinforcement (see Figure 7.8). This post-reinforcement pause can be even longer than the one that occurs under FR schedules. Another distinctive aspect of responding on an FI schedule is an increase in the rate of responding as the end of the interval approaches, producing a pattern often described as a scallop (see again Figure 7.8). A good example of an FI schedule in everyday life is mail delivery, which comes just once a day (FI 24 hours) or in some places twice a day (FI 12 hours). Right after your mail is delivered, you would not check it again, but as the end of the mail-delivery interval approaches, you will start checking again.

On a *variable interval schedule* (a VI schedule), reinforcement still depends on a certain interval having elapsed, but the interval's duration varies unpredictably. In a VI 10 (minute) schedule, for example, sometimes the critical interval is 2 minutes, sometimes 20 minutes, and so on, with an average of 10 minutes. Unlike the variations in responding under an FI schedule, organisms tend to respond at a uniform high rate when the schedule is a VI schedule (see Figure 7.8). For an example of a VI schedule in everyday life, consider redialing a phone number after hearing a busy signal. To receive reinforcement (getting your call through), you have to wait some time interval after your last response (dialing), but the length of that interval is unpredictable.

Aversive conditioning

Negative or aversive events, such as a shock or a painful noise, are often used in instrumental conditioning. In punishment training, a response is followed by an aversive stimulus or event, which results in the response being weakened or suppressed on subsequent occasions. It can effectively eliminate an undesirable response if it is consistent and delivered immediately after the undesired response – especially if an alternative response is rewarded. Rats that have learned to take the shorter of two paths in a maze to reach food will quickly switch to the longer one if they are shocked when taking the shorter path. The temporary suppression produced by punishment provides an opportunity for the rat to learn to take the longer path. In this case, punishment is an effective means of redirecting behavior because it is informative, which seems to be the key to the effective use of punishment.

Applying punishment training to correct human behavior has not always been successful. It is often used in an attempt to increase safe behavior, for example in driving, by using the possibility of an accident as a threat or future punishment: 'If you speed you may die in a road accident.' The problem is that all drivers who are still alive have the experience of *not* dying when speeding. So, speeding cannot really be controlled by conditioning,

unless perhaps we change the threat into 'If you speed, you will be fined.' But, again, most drivers who speed do not get caught and are not fined.

So, even though punishment can suppress an unwanted response, it has several disadvantages. First, its effects are often not as informative as the results of reward. Reward essentially says, 'Repeat what you have done.' Punishment says, 'Stop it!', but it fails to give an alternative. As a result, the organism may substitute an even less desirable response for the punished one. Second, the by-products of punishment can be unfortunate. Through classical conditioning, punishment often leads to dislike or fear of the punishing person (traffic police, parent, or teacher) and the situation (traffic, home, or school) in which it occurred. Finally, extreme or painful punishment may elicit bad behavior that is more serious than the original undesirable behavior.

Escape and avoidance behavior

We have seen that punishment training can sometimes work to inhibit unwanted behaviors. But aversive events can also be used in the learning of new responses. Organisms can learn to make a response to terminate an ongoing aversive event (for example, we may leave a room if there is a painfully loud noise there): this is called **escape learning**. Often, escape learning is followed by **avoidance learning**; the organism learns to make a certain response to prevent an aversive event from even starting (for example, avoiding a certain room if it was associated with a loud noise in the past). To study escape and avoidance learning in animals, psychologists have used a device called a *shuttle box* (see Figure 7.9). The shuttle box consists of two compartments divided by a barrier. On each trial, the animal is placed in one of the compartments. At some point a warning light is flashed, and




Figure 7.9 Shuttle Box. The shuttle box is used to study escape and avoidance learning in animals.

five seconds later the floor of that compartment is electrified. To get away from the shock, the animal must jump over the barrier into the other compartment. Initially, the rat jumps over the barrier only when the shock starts – this is escape learning. With practice, it learns to jump upon seeing the warning light, thereby avoiding the shock entirely – this is avoidance learning. An analysis of the two stages of escape and avoidance learning will shed light on the fact that phobias (fears of specific objects or situations) can be extremely resistant to extinction.

The first stage involves classical conditioning. Through repeated pairings of the warning light (the CS) and the shock (the US), the animal learns that the light predicts the shock, and exhibits a conditioned response of fear (the CR) in response to the light alone.

The avoidance learning seems to present a puzzle: we know that a conditioned response will extinguish if the conditioned stimulus is presented in the absence of the unconditioned stimulus. And that seems to be the case here: once the animal has learned to avoid being shocked (by escaping on time), the CS is no longer followed by the US (the shock). So, why doesn't the conditioned response extinguish? What reinforces the animal for jumping over the barrier? You might say that it is the absence of the shock, but that is a non-event. The solution to this puzzle – and the second stage of our analysis – involves instrumental conditioning. The animal has learned that jumping over the barrier removes an aversive event, namely the conditioned fear itself (see Figure 7.10). Therefore, what first appears to be a non-event is actually fear, and the avoidance behavior is reinforced because it reduces this fear (Mowrer, 1947; Rescorla & Solomon, 1967).

Now, consider someone who has developed a particular fear – let's say, test anxiety – because of past experiences, such as failure on tests. The conditioned response (fear) can be reduced by avoiding having to take the test, for example by sleeping through the alarm, or by asking

for a later test date. The successful reduction of the aversive stimulus (the conditioned fear response) reinforces the avoidance behavior, and will strengthen it in the future. And though it may lead to temporary relief, the consequences of such avoidance behavior are clearly detrimental in the long run. But what to do, when test anxiety is a real problem? Our analysis makes it clear that this response will not extinguish if there is no more exposure to tests. Students who suffer from test anxiety will have to be convinced that their fear response is a learned reaction to past events, which can and will be unlearned with repeated experiences of successful test-taking. See Chapter 15 for further discussion of anxiety disorders and phobias.

Cognitive factors

Cognitive factors play an important role in instrumental conditioning, just as they do in classical conditioning. As we will see, it is useful to view the organism in an instrumental conditioning situation as acquiring new *knowledge* about relationships between responses and reinforcers. As with classical conditioning, we want to know what factor is critical for instrumental conditioning to occur. Again, one of the options is temporal contiguity: an instrumental response is conditioned whenever it is immediately followed by reinforcement (Skinner, 1948). A more cognitive option, closely related to predictability, is that of *control*: an instrumental response is conditioned only when the organism interprets the reinforcement as being controlled by its response.

Important experiments by Maier and Seligman (1976) provide support for the control view. Their basic experiment has two stages. In the first stage, some dogs learn that whether they receive a shock or not depends on (is controlled by) their own behavior, while other dogs learn that they have no control over the shock. Think of the dogs as being tested in pairs. Both members of a pair are in a harness that restricts their movements, and occasionally the pair receives an electric shock. One member of the pair, the 'control' dog, can turn off the shock by pushing a nearby panel with its nose; the other member of the pair, the 'yoked' dog, cannot exercise any control over the shock. Whenever the control dog is shocked, so is the yoked dog, and whenever the control dog turns off the shock, the yoked dog's shock is also terminated. The control and yoked rats therefore receive the same amount of electrical shocks.

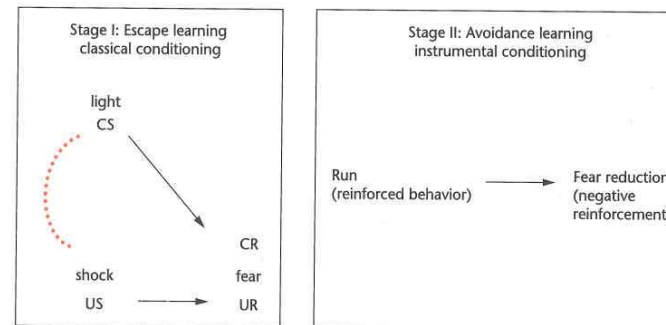


Figure 7.10 Two-stage Analysis of Escape and Avoidance Learning.

To find out what the dogs learned in the first stage of the experiment, a second stage is needed. In this stage, the experimenter places both dogs in shuttle box. On each trial a tone is first sounded, indicating that the compartment the animal currently occupies is about to be subjected to an electric shock. To avoid the shock, the dog must learn to jump the barrier into the other compartment when it hears the warning tone. Control dogs learn this response rapidly – as we saw before in avoidance learning in rats. But the yoked dogs are another story. Initially, the yoked dogs make no movement across the barrier, and as trials progress, their behavior becomes increasingly passive, finally lapsing into utter helplessness. Why? Because during the first stage the yoked dogs learned that shocks were not under their control, and this non-control made avoidance learning in the second stage impossible. In other words: during the first stage of the experiment the animals had learned that they were *helpless*, and this ‘discovery’ prevents them from learning to avoid shock later on, even when they could. The phenomenon of **learned helplessness** has important implications. It supports the notion that instrumental conditioning occurs only when the organism perceives reinforcement as being under its control (Seligman, 1975). (See Chapter 15 for a detailed discussion of learned helplessness, control, and stress.)

We can also talk about these findings in terms of contingencies. We can say that instrumental conditioning occurs only when the organism perceives a contingency between its responses and reinforcement. In the first stage of the preceding study, the relevant contingency is between pushing a panel and the absence of shock. Perceiving this contingency amounts to determining that the likelihood of avoiding shock is greater when the panel is pushed than when it is not. Dogs that do not perceive this contingency in the first stage of the study appear not to look for any contingency in the second stage. This contingency approach makes it clear that the results of research on instrumental conditioning fit with the findings about the importance of predictability in classical conditioning: knowing that a CS predicts a US can be interpreted as showing that the organism has detected a contingency between the two stimuli. In both classical and instrumental conditioning, what the organism seems to learn is a contingency between two events: In classical conditioning, a behavior is contingent on a particular stimulus; in instrumental conditioning, a behavior is contingent on a particular response.

Our own ability to learn contingencies develops very early. In a study of three-month-old infants, each infant was lying in a crib with its head on a pillow (Watson, 1967). Beneath each pillow was a switch that closed whenever the infant turned its head. For infants in the control group, whenever they turned their heads and

closed the switch, a mobile on the opposite side of the crib was activated. For these infants, there was a contingency between head turning and the mobile moving – the mobile was more likely to move with a head turn than without. These infants quickly learned to turn their heads, and they reacted to the moving mobile with signs of enjoyment (they smiled and cooed). The situation is quite different for infants in the non-control group. For these infants, the mobile was made to move roughly as often as it did for infants in the control group, but whether it moved or not was not under their control: there was no contingency between head turns and the mobile movements. These infants did *not* learn to turn their heads more frequently, and after a while they showed no signs of enjoying the moving mobile at all. The mobile appears to have gained its reinforcing character when its movement could be controlled and lost it when its movement could not be controlled. Interestingly, people sometimes suffer from what has been termed an ‘illusion of control’: they believe that they have control over the outcome of a chance event. Langer (1975) describes gamblers who believe that their winnings in a game are the result of their skill, whereas they think of their losses as chance events. For addicted gamblers, this cognitive illusion is likely to contribute to their addiction.

Biological constraints

As with classical conditioning, biology imposes constraints on what may be learned through instrumental conditioning. The instinctive drift discussed under the ‘shaping’ section above is one example of that. Consider pigeons in two experimental situations: reward learning, in which the animal acquires a response that is reinforced by food, and escape learning, in which the animal acquires a response that is reinforced by the termination of shock. In the case of reward, pigeons learn much faster if the required response is pecking a key than if it is flapping their wings. In the case of escape, the opposite is true: pigeons learn faster if the required response is wing flapping than if it is pecking (Bolles, 1970). These seem inconsistent with the assumption that the same laws of learning apply to all situations, but they make sense from an ethological perspective. The reward case with the pigeons involved eating, and pecking (but not wing flapping) is part of the birds’ natural eating activities. A genetically determined connection between pecking and eating is reasonable. Similarly, the escape case involved a danger situation, and the pigeon’s natural reactions to danger include flapping its wings (but not pecking). Birds are known to have a small repertoire of defensive reactions, and they will quickly learn to escape only if the relevant response is one of these natural reactions.

INTERIM SUMMARY

- In instrumental conditioning, animals learn that their behavior has consequences. For example, a rat may learn to press a lever to obtain food reinforcement. The rate of response is a useful measure of response strength. The rate and pattern of responding during instrumental conditioning is determined by schedules of reinforcement.
- Reinforcers increase the probability of a response, whereas punishers decrease the probability of behavioral responses. Reinforcers and punishers can be arranged in either positive or negative contingencies with a particular behavior.



LEARNING AND COGNITION

A famous quote by Watson is this one: ‘Give me a dozen healthy infants well-formed, and my own specified world to bring them up in and I’ll guarantee to take any one at random and train him to become any type of specialist I might select – doctor, lawyer, artist, merchant-chief and yes, even beggar-man and thief, regardless of his talents, penchants, tendencies, abilities, vocations, and race of his ancestors’ (1930, p. 104). The doctrine of early behaviorists can be summarized as follows: to predict human behavior – to control it, even – we need to know only the situation that the human reacts to. And to study the mechanics of learning, it suffices to study simple animals. Since the assumption is that learning results *only* from experience (with stimulus–response relationships, and with the consequence of responses), there is no reason to study or assume ‘higher mental processes’.

We have seen that the empirical approach to the study of behavior had an enormous impact on the history of psychology, especially in the United States. But we have also seen that many of the experiments that were carried out by behaviorists later in the century revealed the

importance of cognition. Recall the experiments by Rescorla, showing that not all stimulus–response relationships are learned equally easily (contingency matters), as well as the experiments by Seligman, showing that reinforcers can lose their ‘power’ if the organism perceives no control over them.

But the basic behaviorist doctrine actually never went unchallenged. Already in the 1930s, Edward C. Tolman, an American psychologist, described findings showing **latent learning** in simple animals: he was able to show that animals were learning, while their behavior did not change in a corresponding way (Tolman & Honzik, 1930). In a typical study, rats would learn to run a complicated maze. One group of rats was rewarded with food for finding their way through the maze: these rats improved gradually in solving the maze, over the course of a number of days. A second group was not rewarded initially, and consequently showed little improvement in solving the maze. However, when a reward was introduced for this second group of rats, their performance almost instantly caught up with the performance of the first group. This showed that the second group of rats had ‘latent knowledge’ of the maze, which was only expressed behaviorally once the food was introduced. Tolman concluded that a rat running through a complex maze was not learning a sequence of right- and left-turning responses, but rather was developing a **cognitive map** – a mental representation of the lay-out of the maze (Tolman, 1932). And more importantly: that this learning occurs even when the animal is not reinforced.

Observational learning

Humans, too, learn many things without immediately being reinforced for the behavior. Consider how you learned to give a presentation in class: when you prepared for it, you probably considered how others go about giving a lecture, and you might have even picked up a book for some advice on how to structure your presentation. Clearly, you did not learn how to give a successful presentation through simple conditioning, which would involve randomly trying out many possible behaviors and repeating only those that were rewarded with a good grade. Rather, you learned through **imitation** and **observational learning**: you copied the behavior of others, whose behavior you observed to be successful.

The researcher whose name is connected with the study of observational learning is Albert Bandura. Early on, Bandura emphasized that observational learning occurs through the principles of operant conditioning (Bandura & Walters, 1963): models inform us about the **consequences** of our behaviors. **Models** often are actual persons whose behaviors we observe, but they can also be more abstract (for example, the written instructions found in a book). Reinforcement in many cases is ‘**vicarious**’: the imitator expects to be reinforced just like the model was.



Figure 7.11 Bandura's 'Bobo doll study'. Bandura showed that children learned to behave aggressively towards a Bobo doll toy, after watching a model behave similarly.

One of Bandura's early studies concerned the observational learning of aggressive behavior in young children (Bandura et al., 1961). In this study, one group of children was shown adult models behaving aggressively towards a Bobo doll (see Figure 7.11). Another group of children was exposed to adult models behaving non-aggressively. Afterwards, the children were led into a room in which they could play with many different toys. The first group of children was shown to display more aggressive behavior towards the Bobo doll than the second group of children. Bandura later showed that the effects are very similar if the children are exposed to aggressive behavior by models presented in film-sequences on a TV screen (Bandura et al., 1963). For this reason, Bandura's work is often cited in discussions concerning the effects of media violence on aggressive tendencies in children. For a more detailed discussion of Bandura's work on aggression as a learned response, see Chapter 11. In his later work, Bandura emphasized the cognitive abilities

that are necessary for observational learning to occur (Bandura, 1977, 2001). The learner must be able to (1) pay attention to the model's behavior and observe its consequences, (2) remember what was observed, (3) be able to reproduce the behavior, and (4) be motivated to do so. In other words: observational learning involves the ability to imagine and anticipate – thoughts and intentions are essential.

Most of Bandura's work focuses on the importance of cognition in social learning in humans. In his view, humans are agents of their own experiences, not 'undergoers' (Bandura, 2001). His theory on social learning is further discussed in Chapter 13. For now, it suffices to say that Bandura's 'agentic perspective' draws our attention to the fact that cognitions motivate actions, and that a sense of *self-efficacy* (an individual's belief in their own effectiveness) is essential for complex and social learning. If you believe that you are simply incapable of giving a good presentation in class, you are unlikely to motivate

yourself to plan and anticipate the effects of the decisions you make regarding that talk.

Prior beliefs

Humans and animals alike are very sensitive to learning relationships between stimuli, as we have seen. When relationships between stimuli or events are less than perfectly predictable, humans can even estimate the degree of objective relationships between stimuli (Shanks & Dickinson, 1987; Wasserman, 1990). This has been shown with experimental tasks that were novel to the subjects, and that did not concern stimuli about which the subjects had any prior beliefs. But when similar experiments are carried out using stimuli about which the subjects do hold prior beliefs, the situation changes in an interesting way: such studies show that prior beliefs can constrain what the subjects learn. This again indicates that learning involves processes in addition to those that form associations between inputs.

In these studies, a different pair of stimuli – for example, a picture and a word – is presented on each trial, and the participant's task is to learn the relationship between the members of the pairs. Subjects might detect, for example, that certain pictures are more likely to appear alongside certain words. Some striking evidence for the role of prior beliefs comes from cases in which there is no objective association between the pairs of stimuli, but participants nevertheless detect such a relationship. In one experiment, each trial presented the subjects with a picture of a person drawn by a mental patient, alongside a description of the symptoms of that patient. These symptoms included statements such as 'suspiciousness of other people' and 'concerned with being taken care of'. The participant's task was to determine whether any aspects of the drawings were associated with any of the symptoms. The experimenters had paired the symptoms randomly with the drawings so that there was no objective association between them. Yet, participants consistently reported such associations, and the relationships they reported were ones that they probably believed before participating in the experiment – for example: that large eyes are associated with suspiciousness or that a large mouth is associated with a desire to be taken care of by others. These non-existent but plausible relationships detected by the subjects are referred to as *spurious associations* (Chapman & Chapman, 1969).

Even when there is an objective association to be learned, prior beliefs affect what subjects actually learn. This was shown in studies similar to the one described above (Jennings, Amabile, & Ross, 1982). On each of a set of trials, participants were presented with two measures of an individual's honesty taken from two completely different situations. For example, one measure

might have been how often a young boy copied another student's homework in school, and the second an indication of how often that same boy was dishonest at home. Most people believe (erroneously) that two measures of the same trait (such as honesty) will always be highly correlated. This is the critical prior belief. In fact, the objective relationship between the two measures of honesty varied across different conditions of the experiment, sometimes being quite low. The participants' task was to estimate the strength of this relationship by choosing a number between 0 (which indicated no relation) and 100 (a perfect relation). The results showed that participants consistently overestimated the strength of the relationship. Their prior belief that an honest person is honest in all situations led them to see more than was there. Other research has shown that our prior beliefs can be overcome, if the data (the objective association) are made salient enough – only then do subjects learn what is actually there (Alloy & Tabachnik, 1984).

The results of these studies are reminiscent of what we called top-down processing in perception (see Chapter 5), in which perceivers combine their expectations of what they are likely to see with the actual input to yield a final percept. In top-down processing in learning, the learner combines prior belief about an associative relationship with the objective input about that relationship to yield a final estimate of the strength of that relationship.

The importance of prior beliefs in human learning strengthens the case for a cognitive approach to learning. The research also has a connection to the ethological approach to learning. Just as rats and pigeons may be constrained to learn only associations that evolution has prepared them for, so we humans seem to be constrained to learn associations that our prior beliefs have prepared us for. Without prior constraints of some sort, perhaps there would simply be too many potential associations to consider, and associative learning would be chaotic, if not impossible.

INTERIM SUMMARY

- According to the cognitive perspective, the crux of learning is an organism's ability to represent aspects of the world mentally and then operate on these mental representations rather than on the world itself.
- Learning through imitation and observation happens as a result of vicarious reinforcement: by observing a model's behavior, the imitator expects to be reinforced just like the model was.
- When learning relationships between stimuli that are not perfectly predictive, people often invoke prior beliefs.

CUTTING EDGE RESEARCH



Map learning in London's taxi drivers: Structural and functional consequences

Taxi drivers in London are famous for their extensive training. All London taxi drivers have to pass an exam at the Public Carriage Office. To pass it, they spend multiple years acquiring 'The Knowledge': the detailed lay-out of the city with 25,000 streets and thousands of places of interest. Maguire and her co-workers used magnetic resonance imaging (MRI) to show that these London taxi drivers have greater gray matter volume in the posterior (back) part of their hippocampi and smaller gray matter volume in the anterior (front) part of their hippocampi, compared to an age-matched control group (Maguire et al., 2000; Maguire et al., 2003). These results are interesting, because they suggest that the hippocampus in healthy adult humans has the ability to change structurally as new spatial knowledge is acquired.

Other recent findings show similar 'environmentally driven plasticity': the ability of the human neural system to change structurally in response to specific demands. For example: Draginsky et al. (2004) showed structural changes in the brains of subjects who trained their juggling skills. Musicians also show an increase in gray matter volume in motor and auditory areas, associated with time spent practicing and practice intensity (Gaser and Schlaug, 2003).

In a more recent study, Maguire and her co-workers compared London taxi drivers with a control group who also spend all day driving in busy London: London bus drivers (Maguire et al., 2006). The two groups of subjects were similar on many dimensions (driving experience, stress levels, age, handedness, education, IQ) but differed in one important way: whereas taxi drivers navigate the city freely (relying on their superior memory of the city's lay-out), bus drivers use only a constrained set of routes. Earlier MRI findings were replicated: taxi drivers have greater gray matter volume in posterior hippocampi and less volume in anterior hippocampi than bus drivers (Maguire et al., 2006). Because of the carefully chosen control group, this finding lends further support to the hypothesis that the gray matter differences are a result of the specific demands placed on spatial memory. Interestingly, the study also revealed that there might be a price that London's taxi drivers pay for acquiring 'The Knowledge'. The two groups were tested for functional differences, and it was found that the ability to acquire new visuo-spatial information was worse in taxi drivers than in bus drivers. In fact, the taxi drivers did worse than would be expected for healthy men their age. This might be a cognitive trade-off, and a consequence of the reduced anterior hippocampal gray matter volume found in the taxi drivers.

processes. Hebb saw humans as biological organisms and the product of evolution. He believed that mental processes should be regarded as processes that involve the nervous system and the brain – and that learning is a process that involves *changes in neural activity*. Moreover, he believed that it was possible to speculate about these processes in a meaningful way – a clear departure from the influential ideas of behaviorism at that time. Hebb formulated ideas about learning and the brain, that were inferences based on observations (Hebb, 1966).

Hebb's main contribution to the study of learning concerns his ideas about possible neurological changes underlying learning. Hebb hypothesized that if input from neuron A repeatedly increases the firing rate of neuron B, then the connection between neurons A and B will grow stronger (Hebb, 1958). In other words: repetition of the same response leads to permanent changes at the synapses between neurons. This idea is known as the **Hebbian learning rule**. At Hebb's time, this notion was a theoretical speculation. Current knowledge of the biochemistry underlying neurological changes has confirmed Hebb's ideas, as we will see.

In this section we will discuss **neural plasticity**: the ability of the neural system to change in response to experience. To appreciate these ideas, you need to recall from Chapter 2 the basic structure of a neural connection

and how it transmits an impulse. An impulse is transmitted from one neuron to another by the axon of the sending neuron. Because the axons are separated by the synaptic gap, the sender's axon secretes a neurotransmitter, which diffuses across the synaptic gap and stimulates the receiving neuron. The key ideas regarding learning are (1) that a change in the synapse is the neural basis of learning and (2) that the effect of this change is to make the synapse more (or less) efficient.

Habituation and sensitization

To understand the neural basis of complex psychological phenomena, it is best to examine simple forms of learning and memory. Perhaps the most elementary form of learning is non-associative learning. Habituation and sensitization are examples of this type of learning. During *habituation*, a behavioral response, such as orienting to an unfamiliar sound, decreases over successive presentations of that stimulus. During *sensitization*, a behavioral response increases during presentations of intense stimuli, such as very loud noises. In both cases, learned changes in behavior can persist for hours to days.

To study these learning processes at the neural level, a team of researchers led by Nobel prize winner Eric Kandel has chosen to work with an organism with a very simple nervous system: the marine slug, *Aplysia californica* (Kandel, Schwartz, & Jessell, 1991). *Aplysia* has proven to be an excellent experimental model to study non-associative learning, because it has a simple and accessible nervous system. Learning in *Aplysia* has been studied by measuring the gill withdrawal reflex, which can be elicited by gentle mechanical stimulation of the gill or surrounding tissue. The gill withdrawal reflex is a defensive response that protects the fragile gill from injury.

When the gill is lightly stimulated with a water jet, the gill is withdrawn. However, repeated stimulation of the gill produces weaker and weaker withdrawal responses. Researchers have shown that this habituation learning is accompanied by a decrease in the amount of neurotransmitter secreted by gill sensory neurons onto a motor neuron that controls gill withdrawal (Figure 7.12).

The gill withdrawal reflex also exhibits sensitization. If an intense stimulus, such as an electric shock to the tail or head is administered, then the light touch to the gill will elicit a much larger withdrawal response. Like habituation, sensitization learning involves a change in synaptic transmission between sensory and motor neurons that control the gill. In this case, the intense stimulus causes an increase in the amount of neurotransmitter secreted by the sensory neuron. This increase depends on the activation of interneurons that release serotonin onto the gill sensory neurons. These findings provide relatively direct evidence that elementary learning is mediated by synaptic changes at the neuronal level.

Classical conditioning

What about associative learning? Do synaptic changes like the ones just described mediate classical conditioning? Indeed, researchers have proposed a neural model of classical conditioning in *Aplysia* that is remarkably similar to that for sensitization (Hawkins & Kandel, 1984). Incredible progress has also been made in understanding the neural mechanisms of classical conditioning in mammals, including humans. Two experimental models have been used with great success: eyeblink conditioning and fear conditioning.

Eyeblink conditioning

When a stimulus, such as an air puff (the US), is directed at the eye, it elicits a reflexive blink. This unconditional eyeblink response can be conditioned if a CS, such as a tone, precedes the puff. After training, the CS will come to elicit eyeblink CRs even when the air puff is not presented.

Detailed mapping studies in rabbits by Richard Thompson and colleagues have revealed the neural circuitry in this form of classical conditioning (Thompson & Krupa, 1994). The essential site of synaptic plasticity appears to reside in the cerebellum. Animals with cerebellar lesions cannot learn or remember the conditioned eyeblink (although they show normal eyeblink URs). Interestingly, eyeblink conditioning is associated with changes in synaptic transmission in the cerebellum. This change is called **long-term depression** (LTD) and is associated with a long-lasting decrease in synaptic transmission at synapses in the cerebellar cortex. This change occurs in the pathway that transmits information about the CS to cerebellar cortical neurons. The decrease in CS transmission in the cerebellar cortex results in a behavioral CR because the cerebellar cortex normally inhibits the CR-producing part of the eyeblink conditioning circuit.

Fear conditioning

As we saw in earlier in this chapter, emotional responses such as fear are easily conditioned. Laboratory work with rats has yielded important insights into the brain mechanisms of this sort of learning. In this model, rats are conditioned to fear a place or a cue that has been paired with an aversive stimulus, such as foot shock. Fear is often assessed by measuring the freezing response – the immobility that rodents show when they are afraid. As in the eyeblink conditioning paradigm, a specific brain area is essential for learning and remembering fearful experiences. In this case, it is the amygdala, a limbic system structure deep within the brain that is important for emotions, including fear (Klüver & Bucy, 1937). The amygdala receives sensory information from thalamic and cortical brain areas, associates these stimuli, and translates these associations into fear responses mediated by the hypothalamus, midbrain, and medulla

LEARNING AND THE BRAIN

The transition from behaviorism to a more cognitive approach to the study of learning was also stimulated by ideas concerning the brain. The Canadian researcher Donald Hebb contributed much to early theories about learning and the brain; his ideas have been very influential in the field of behavioral neuroscience.

We have seen that early behaviorists focused on the study of observable events, rather than on mental

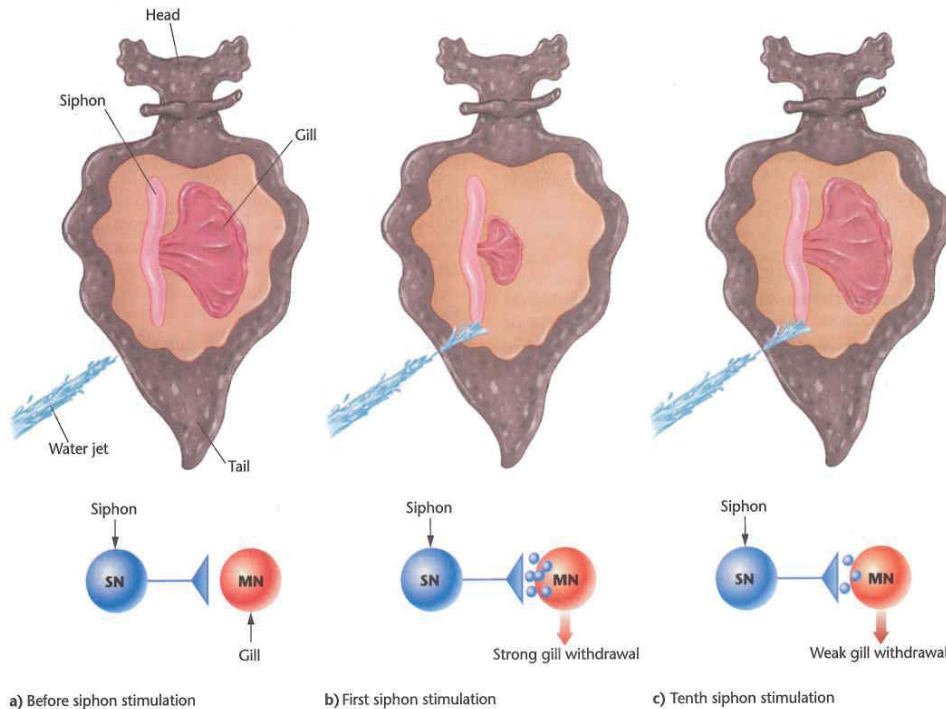


Figure 7.12 Habituation in *Aplysia californica*. (a) Before mechanical stimulation of the siphon, the gill is extended. (b) When water is squirted on the siphon for the first time during habituation training, the gill withdraws vigorously. A simple circuit involving siphon sensory neurons (SN) that form excitatory synaptic contacts onto motor neurons (MN) mediates gill withdrawal. (c) After the 10th siphon stimulus, the magnitude of gill withdrawal is small. The gill withdrawal response has habituated. Habituation is mediated by a decrease in presynaptic neurotransmitter release at the SN-MN synapse.

(Figure 7.13). Animals with amygdala damage cannot learn or remember fear memories (Davis, 1997; Fendt & Fanselow, 1999; Maren, 2001; Maren & Fanselow, 1996). Moreover, neurons in the amygdala exhibit many changes during new fear learning. For example, amygdala neurons increase their activity in response to CSs that have been associated with aversive UCSs. It appears that learning in the amygdala is mediated by **long-term potentiation (LTP)**, which is a persistent increase in synaptic transmission in pathways that send CS information to the amygdala (Rogan & LeDoux, 1996).

Hence, in both eyeblink conditioning and fear conditioning, changes in synaptic transmission in defined brain areas are responsible for the behavioral changes that accompany associative learning.




Figure 7.13 Neural Circuit for Classical Fear Conditioning. The amygdala receives sensory information from many sensory areas, including the thalamus, neocortex, and hippocampus. The amygdala associates this information during fear conditioning and then generates fear CRs by projecting to brain areas, such as the midbrain, hypothalamus, and medulla, that mediate a number of different fear responses.

Another study shows that what holds for other mammals applies to humans as well (Bechara et al., 1995). This study involved a human patient, referred to as S.M., who had a rare disorder (Urbach-Wiethe disease) that results in degeneration of the amygdala. S.M. was exposed to a fear-conditioning situation in which a neutral visual stimulus (the CS) was predictably followed by the sound of a loud horn (the US). Despite repeated trials, S.M. showed no evidence of fear conditioning. Yet S.M. had no trouble recalling the events associated with the conditioned and unconditioned stimuli. Another patient, who had a normal amygdala but had suffered damage to a brain structure involved in the learning of factual material, showed normal fear conditioning but was unable to recall the events associated with the conditioning. The two patients had the opposite problems, indicating that the amygdala is involved in the learning of fear, not learning in general.

Cellular basis of learning

As we have seen, learning results in changes in synaptic transmission in both slugs and mammals. We have not been very specific about what causes these changes in synaptic transmission. There are several possibilities. One is that learning results in an increase or decrease in the amount of neurotransmitter secreted by the sending neuron, perhaps because of an increase or decrease in the number of axon terminals that secrete the neurotransmitter (as we saw with sensitization and habituation in the *Aplysia*). Alternatively, there may be no change in the amount of neurotransmitter sent, but there may be a change in the number of postsynaptic receptors. Other possibilities are that the synapse could change in size or that entirely new synapses could be established. All of these changes are examples of **synaptic plasticity**: changes in the morphology and/or physiology of synapses involved in learning and memory. Indeed, learning may also be accompanied by the growth of new neurons (Gould, Beylin, Tanapat, Reeves, & Shors, 1999; van Praag, Kemperman, & Gage, 1999).

A critical advance in understanding the cellular basis of memory was the finding that synapses in several brain areas can exhibit long-lasting increases in synaptic transmission under some conditions (Berger, 1984; Bliss & Lomo, 1973). For example, rapid electrical stimulation of synapses in the hippocampus causes an enhancement in the magnitude of synaptic responses that lasts for days or even weeks (Figure 7.14). This long-term potentiation requires a special type of neurotransmitter receptor, the NMDA receptor (Malinow, Otmakhov, Blum, & Lisman, 1994; Zalutsky & Nicoll, 1990). The NMDA receptor is unlike other receptors, in that two conditions must be satisfied for the receptor to open. First, presynaptic glutamate must bind to the NMDA receptor. Second, the


postsynaptic membrane in which the receptor resides must be strongly depolarized. Once opened, the NMDA receptor allows a very large number of calcium ions to flow into the neuron. That influx of ions appears to cause a long-term change in the membrane of the neuron, making it more responsive to the initial signal when it recurs at a later time (see Figure 7.14).

Interestingly, activation of NMDA receptors could arise during classical conditioning, in which weak (CS) and strong (US) inputs converge onto single neurons. In this case, LTP would be induced at synapses transmitting CS information because conditioning would result in both presynaptic activity (during the CS) and postsynaptic depolarization (during the US) in the neurons upon which CS and UCS information converge (Maren & Fanselow, 1996).

Such a mechanism, in which two divergent signals strengthen a synapse, provides a possible explanation of how separate events become associated in memory. For example, learning someone's name requires that you make an association between the person's appearance and his or her name. LTP strengthens synapses so that the sight of the person will prompt you to recall the person's name. In classical fear conditioning, an association is established between a relatively neutral CS and an aversive US. The NMDA mechanism thus offers an intriguing theory to explain how events are associated in memory (Maren, 1999).

INTERIM SUMMARY

- Habituation is mediated by a decrease in synaptic transmission, and sensitization by an increase in transmission.
- Synapses in the mammalian brain are involved in storing information during learning. Increases in synaptic transmission, such as long-term potentiation, are part of these learning processes.



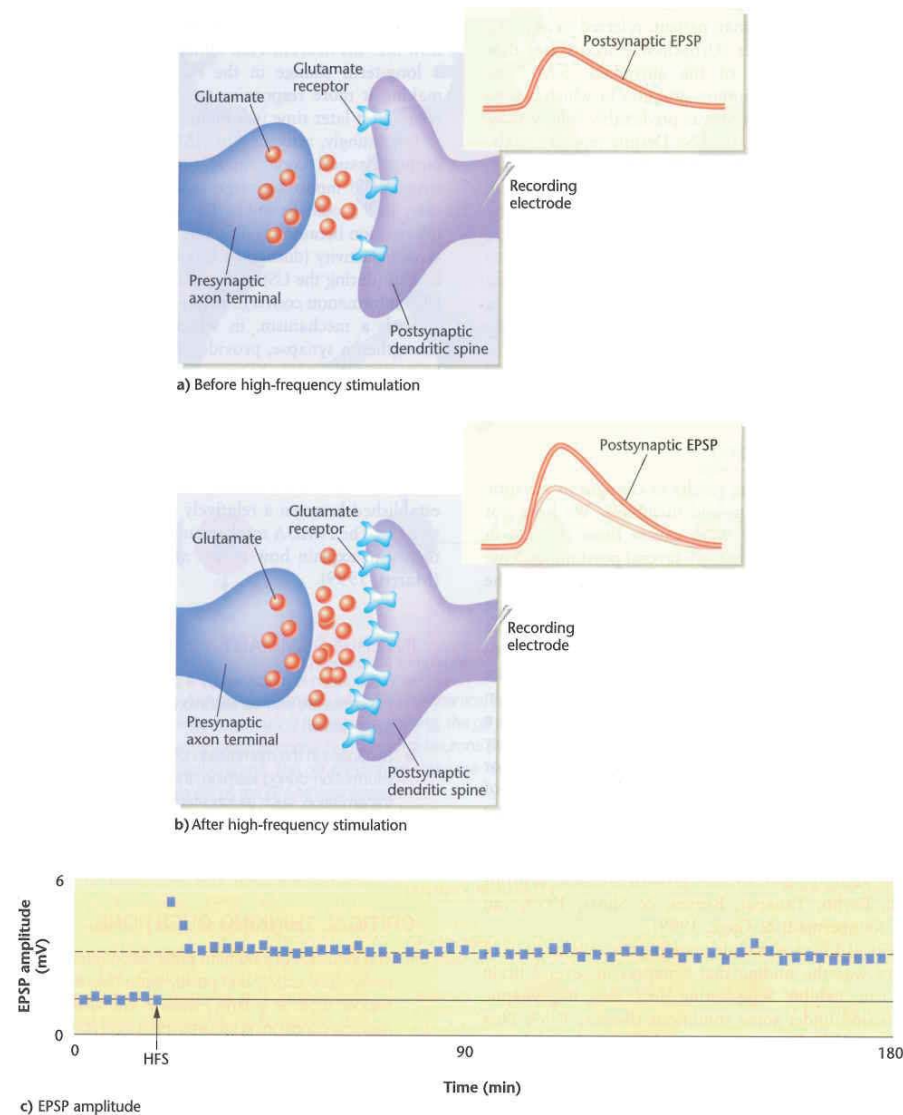


Figure 7.14 Long-term Potentiation in the Hippocampus. (a) Before high-frequency stimulation (HFS), pre-synaptic glutamate release activates post-synaptic glutamate receptors to produce an excitatory post-synaptic potential (EPSP). (b) After high-frequency stimulation of the pre-synaptic neuron, the post-synaptic EPSP is greatly increased in amplitude. This increase is due to an enhancement of pre-synaptic neurotransmitter release and an increase in the number of post-synaptic glutamate receptors. (c) Graph illustrating the amplitude of the EPSP before and after HFS. Long-term potentiation is indicated by the persistent increase in EPSP amplitude.

LEARNING AND MOTIVATION

Coming to the end of this chapter on learning, you may be surprised to have read preciously little about the kind of learning you are engaging in at this very moment: studying. We have focused instead on very basic learning processes. However, psychology does have much to say about the kind of processes involved in the *how* and the *why* of complex learning. Most of this will be covered in the next couple of chapters in this book: the 'how' of complex human learning is described in Chapters 8 and 9, which address memory and cognition, respectively. Questions regarding the 'why' of certain behaviors will be addressed in Chapter 10, which concerns motivation. In this section, we will briefly review some of the most relevant theories that tie concepts from the field of motivation to the study of complex human learning.

Arousal

We have already discussed some of Hebb's work on the neural underpinnings of learning. Hebb also formulated an arousal theory of motivation. This aspect of his work was also instrumental in 'closing the gap' between behavioral versus physiological approaches to learning. **Arousal** has both a physiological and a psychological dimension. Physiologically, the term refers to the level of alertness of an organism. Psychologically, the term refers to the tension that can accompany different levels of arousal, ranging from calmness to anxiety. In Hebb's view, arousal is an important motivational concept (Hebb, 1955). He proposed that any organism is motivated to maintain that level of arousal which is appropriate for the behavior it is engaged in. Hebb's insights were based on the **Yerkes-Dodson law** (Yerkes & Dodson, 1908), which relates performance to arousal. This law states that most tasks are best performed at intermediate levels of physiological arousal. Since very complex tasks have enough arousal associated with them, they drive the individual to seek out calmness. Very simple tasks, on the other hand, can become boring at low levels of arousal. According to Hebb, the bored individual will seek out other activities or novel stimuli to increase arousal. Others have since argued that the **exploratory behavior** of humans (our desire to discover and learn novel things) is the result of a desire for stimulation, which can be explained by arousal theory (Berlyne, 1966).

From incentives to goals

The history of the study of motivation mirrors what we saw in the history of the study of learning. Early theorists focused on **incentives**: a behavior is motivated by its expected reward – for example: a hungry animal is driven



Learning is more enjoyable and more effective when you are intrinsically motivated.

to eat because that will reduce the hunger it experiences (Hull, 1943). Hebb (1966), Tolman (1951), as well as others at the time, pointed out that many human behaviors cannot be motivated by the expectation of an immediate reward. Consider again the example of studying: you are probably motivated to study this book partly because you would like to do well in the course and attain your degree. Your desire to graduate is a long-term goal that motivates your current behavior – an example of complex **goal-oriented behavior**. It is clear that cognition plays a role in our ability to anticipate the long-term consequences of current behavior.

Some of the most complex human behavior can be said to arise from our psychological needs, and have to do with intellectual and emotional aspects of our functioning – our needs for social belonging and self-esteem, for example. The study of human emotion (the topic of Chapter 11) is closely linked to the study of motivation.

Intrinsic motivation and learning

In a cognitive approach to the study of motivation, the emphasis is on the individual's understanding and interpretation of their own actions: Why do we think we do things? In other words: what do we **attribute** our own motivations to? Ask yourself why you are studying this chapter, right now. Is it because you are interested in the material, and comprehending it gives you a sense of competence and pride? If so, you are **intrinsically motivated** by these feelings. Or perhaps you are studying because you think it is necessary in order to do well on your exam and get a good grade in your course. If that is the case, you are **extrinsically motivated** by the external rewards that you anticipate.

Research has shown that intrinsically motivated individuals are more persistent at a task, that their memory of complex concepts is better, and that they handle complex material in cognitively more creative ways (Deci, Ryan, & Koestner, 1999). This suggests that studying is not only



SEEING BOTH SIDES

WHAT ARE THE BASES OF SOCIAL LEARNING?

Social learning cannot be explained through 'simple' associative learning

Juan-Carlos Gómez, University of St. Andrews

Social learning is a complex affair relying upon a plurality of cognitive and motivational mechanisms in which associative learning plays only a limited role. I discuss three pieces of evidence indicating that social learning cannot be the result of simple associative learning.

A key social learning skill is *gaze following*, the reaction of looking in the same direction as others to identify their objects of attention. This is an old evolutionary skill shared with other primates (for example chimpanzees follow the gaze of other chimpanzees), but it is not a reflex reaction. Gaze following is learned during the first year of life, but not through simple associations. This was dramatically demonstrated by an experiment Corkum and Moore (1998) conducted with 8-9 month-old infants who had not yet learned to follow gaze on their own. They tried to teach them the gaze following response with selective reinforcement. Thus a group of children consistently found a reinforcing event if they looked in the same direction as an adult; however, a second group found the reinforcing event only if they looked in the direction opposite to where the adult looked. If gaze following is learned through simple association, this group of children should have learned to look in the direction opposite to the adult. However, they were completely unable to learn this reverse, unnatural contingency, whereas children in the normal gaze following group learned easily to follow the gaze of the adult. Even more surprisingly, children in the reverse contingency group spontaneously learned to follow gaze in the natural direction, despite the fact that they were never rewarded for doing so. Gaze direction is not just an arbitrary stimulus: there seems to be something intrinsically directional in gaze that tightly constrains what can be learned and how it is learned. The rules of simple associative learning do not apply here.

Social learning involves a complex interaction among various social cognitive adaptations that modulate what is learned and how it is learned. For example, *imitation*, another key skill for social learning, is not an automatic mechanism controlled by simple contingencies. Thus, children can imitate behaviors that they actually do not see completed. In a study by Meltzoff (1995), an adult tried but failed to pull apart the two parts of an object

because they were stuck. However, young infants correctly imitated the *intended* action when handed an unstuck version of the same object. Children were filling up the behavioral gap in the model's demonstration with their own representation of the intended outcome. Similarly Gergely, Bekkering, and Kiraly (2002) report that young children imitate 'rationally'. When confronted with a bizarre action performed by a model, switching on a light-box by leaning forward and pressing its top with the head instead of the hand, children imitate this unusual action only if it is presented without a justifying context. However, if the adult had her hands busy holding a blanket around her shoulders because she felt cold, children did not imitate the bizarre action but used their hand to turn on the light. Children make a rational evaluation of the situation in terms of goals, available means, and context.

Finally let's consider the case of autism. Children with autism have good associative learning skills. Indeed associative learning is very useful in teaching them adaptive behaviors (e.g., some speech) and extinguishing undesirable habits (e.g., self-injury behaviors). However, associative learning has striking limitations when it comes to acquiring advanced social skills. For example, when learning new words typical children assume that a particular word corresponds to the object the person who utters the word is looking at. However, children with autism (who lack gaze following skills, or acquire them much later than typical children) learn an association between the word and what they themselves are looking at. In this way, they may acquire peculiar, idiosyncratic meanings for some words. In an experiment Baron-Cohen, Baldwin and Crowson (1997) found that children with autism learned the meaning of an invented word if it was broadcast from a loudspeaker when they touched a particular toy in the room. In contrast, typical children failed to learn the meaning of the word with this method, they need the social context of a real speaker to learn words. In social learning, typical children modulate the use of associative learning with social cognitive skills. Children with autism seem to engage in pure associative learning, and this frequently leads them to insufficient or maladaptive learning. They are good at detecting simple and straightforward physical contingencies, but they have difficulty dealing with the imperfect, context-dependent, contingencies of social interaction. For this, specific social cognitive adaptations that go beyond simple associative learning are needed. The case of autism clearly illustrates the limitations of associative learning in explaining the complexity of social learning and cognition.

SEEING BOTH SIDES

WHAT ARE THE BASES OF SOCIAL LEARNING?



Learning, not instinct, determines behavior: social or otherwise

Phil Reed, Swansea University

In the early twentieth century, a great debate raged between those who believed that behavior is best explained by learning (e.g., behavioral psychologists, such as Watson), and those who believed that behavior is best accounted for by inherited instincts (e.g., 'instinct psychologists', such as McDougall). This debate remains central to understanding the great theories in psychology. At the height of this debate, Holt (1931, p. 4) famously commented on 'instinct psychology': *'Man is impelled to action, it is said, by his instincts...if he twiddles his thumbs, it is the thumb-twiddling instinct; if he does not twiddle his thumbs, it is the thumb-not-twiddling instinct. Thus, everything is explained by magic - word magic.'* This statement remains relevant now to explain flaws in contemporary views of social learning which rely on notions such as instinct or innate drives.

By reducing the argument for instinct to an absurdity, Holt highlighted three problems. Firstly, the circular nature of the explanation offered; it merely re-describes the observed behavior as if it were a theory about that behavior: Why does she twiddle her thumbs? Because she has a 'thumb-twiddling' instinct! How do you know she has a 'thumb-twiddling' instinct? Because she twiddles her thumbs! This argument has been central to many critiques of cognitive psychology. Secondly, the naïve view of the phenomenon to be explained; assuming that a set of complex behaviors can be characterised as a single entity, which can be explained by reference to a small set of constructs. If 'thumb twiddling' were replaced by 'social learning', the assumption that there is one entity called 'social learning', that can be explained by reference to a very small number of instincts, seems overly simplistic. Finally, instinct theories do not offer explanations of where and how such instincts arise.

Tomasello (1999) suggests that social learning underlies human cultural evolution, allowing a cumulative growth in knowledge not apparent in other species. Other species are claimed not to engage in the kinds of social learning that enable this incremental cultural learning to occur, rather each generation has to acquire knowledge afresh (Kummer & Goodall, 1985). He suggests that some innate mechanism, highly-developed in humans, helps drive critical processes such as: joint attention, language learning, and cultural learning (Tomasello, 2003). This mechanism has been termed an 'interactional instinct' (Lee et al., 2009), and this labelling reveals the true nature of this form of theorising: this is 1920s 'instinct psychology' reborn, as if a century of progress in empirical findings in learning theory had not occurred!

Social learning is regarded as having two major forms (Whiten & Ham, 1992). 'Non-imitative social learning' occurs when the presence of another facilitates the acquisition of knowledge, but not necessarily the specifics of an observed behavior. Whereas, in 'true imitation' an observer learns to exactly copy the actions of a model. Learning theory supplies explanations of both forms across the species: non-imitative social learning is explained by classical conditioning (Mineka & Cook, 1988); and true imitation is explained by discriminated operant learning (learning when certain actions will have particular consequences; Miller & Dollard, 1941). Both forms can be shown to occur in nonhumans, and to relate to cultural transmission. The availability of such explanations, and the supporting evidence, suggests there is little need to argue for special social learning instincts in humans.

There are many examples of non-imitative social learning in nonhumans, which illustrate the application of classical conditioning (see Olsson & Phelps, 2007). A seminal example relates to the way in which rats learn food preferences, and how this learning spreads throughout a colony. Galef (1996) presented rats with another rat, together with a novel food (typically avoided by rats), and found that an observer subsequently ate the food more readily than a rat presented with the food in the absence of another rat. Similarly, Mineka and Cook (1988) demonstrated that laboratory-reared monkeys learned to fear snakes when exposed to wild monkeys showing fear of snakes. These examples can be explained by the observer learning the relationship between a stimulus and an outcome through classical conditioning. Importantly, this form of learning produces changes in 'cultural practice', which is neither based on true imitation, nor restricted to humans.

It has been argued that true imitation is uniquely human due to the highly cognitively demanding 'cross-model matching' required to match an observer's visual representation of a behavior to the kinaesthetic senses of their own movements (Tomasello, 1996). However, Heyes and Dawson (1990) have shown that, when a rat was placed in a cage, opposite another rat pushing a bar, either right or left, to earn food, then, when later exposed to the bar, the observer rat would press in the same direction as the demonstrator rat. As the observer was moved 180° before being exposed to the bar, this meant that it must have learned to press in the same direction as the demonstrator, and not in the direction that it witnessed the bar moving across its own visual field; the observer rat had learned about the specific actions of another rat. However, Mitchell et al. (1999) found that observer rats may detect odour on the side of a bar that demonstrators had pushed, suggesting that the observers were not encoding the visual representation of the

other rats. This does not mean that true imitation cannot occur, but that it may not occur in a visual medium for largely non-visual species. Similarly, Reed et al. (1996) noted that imitation only occurred in rats who had been socially-reared, not in those reared in isolation, suggesting that imitation needs to be learned in a social environment (see Baer et al., 1967).

more fun, but also more effective when you are intrinsically motivated. According to some researchers, the attribution of motives to intrinsic causes results in a feeling that one is in control of one's own actions, that one is self-determined (Deci & Ryan, 1985). When external rewards become important, they take away from our sense of self-determination. Persistence is reduced, and – especially for difficult tasks – the individual will be more easily discouraged. These ideas are closely related to ideas expressed by Bandura; we saw earlier that he emphasized the importance of self-efficacy.

There is experimental evidence showing that external rewards can harm intrinsic motivation. One example is research with children that was carried out by Lepper and Green (1975). One group of children was solving puzzles, expecting no reward. The other group of children were told that they would be allowed to play with certain toys, if they worked on the puzzles first. At a later time, both groups were allowed to play with the puzzles spontaneously (neither group expecting a reward). More of the children who had initially not expected a reward, chose to work with the puzzles spontaneously. This type of research has been repeated many times, confirming the detrimental effects of external rewards for persistence and performance on a task that was initially intrinsically motivating (Deci, Ryan, & Koestner, 1999). When rewards are introduced, it seems that 'play becomes work': the individual attributes their own engagement with the task to the anticipated external reward, rather than to the inherent satisfaction associated with it. This effect is called the **overjustification effect**: the external reward becomes the justification for performing the task – a cognitive interpretation of the situation that is detrimental to intrinsic motivation.

Let's assume for the moment that you are intrinsically motivated to study your psychology text book – as of course we hope you are. The research shows that your intrinsic motivation might suffer once you realize that effective studying also holds the promise of an external reward: the good grade. And that would be a pity! Motivation researchers point to the importance of self-determination and self-efficacy, as we have seen. This

In summary, learning theory argues that an 'imitative instinct' is an empty explanatory concept, and there is ample evidence that social learning can occur in many species, certainly in its non-imitative (classically conditioned) form, as well as in its imitative (instrumentally conditioned) form, and that both types of social learning can produce cultural transmission.

means that – besides studying – you should try to protect your intrinsic motivation. Spend some time actively asking yourself what interests you about the material. How does it relate to questions that you ask yourself, and to other topics that interest you? And also realize that grades are not only external rewards – grades also provide information about your level of achievement. A good grade tells you that you have mastered something, and a poor grade – especially when there is also some meaningful feedback – informs you about what might have been lacking in your preparation. Reinterpreting the meaning of a grade in this way (from external reward to a source of information) is an active way to increase your own sense of control.

INTERIM SUMMARY

- In humans, complex learning can be thought of as goal-oriented behavior arising from our psychological needs for self-determination and achievement.
- Intrinsically motivated individuals are more persistent at a task than extrinsically motivated individuals.
- External rewards can be detrimental to intrinsic motivation.



CHAPTER SUMMARY

- 1 Learning may be defined as a relatively permanent change in behavior that is the result of practice. There are four basic kinds of learning: (a) habituation, in which an organism learns to ignore a familiar and inconsequential stimulus; (b) classical conditioning, in which an organism learns that one stimulus follows another; (c) instrumental conditioning, in which a particular response leads to a particular consequence; and (d) complex learning, in which learning involves more than the formation of associations.
- 2 Early research on learning was done from a behaviorist perspective. It often assumed that behavior is better understood in terms of external causes than internal ones, that simple associations are the building blocks of all learning, and that the laws of learning are the same for different species and different situations. These assumptions have been modified in light of subsequent work. The contemporary analysis of learning includes cognitive factors and biological constraints, as well as behaviorist principles.
- 3 In Pavlov's experiments, if a conditioned stimulus (CS) consistently precedes an unconditioned stimulus (US), the CS comes to serve as a signal for the US and will elicit a conditioned response (CR) that often resembles the unconditioned response (UR). Stimuli that are similar to the CS also elicit the CR to some extent, although discrimination training can curb such generalization. These phenomena occur in organisms as diverse as flatworms and humans.
- 4 Cognitive factors also play a role in conditioning. For classical conditioning to occur, the CS must be a reliable predictor of the US; that is, there must be a higher probability that the US will occur when the CS has been presented than when it has not.
- 5 According to ethologists, what an animal learns is constrained by its genetically determined 'behavioral blueprint'. Evidence for such constraints on classical conditioning comes from studies of taste aversion. Although rats readily learn to associate the feeling of being sick with the taste of a solution, they cannot learn to associate sickness with a light. Conversely, birds can learn to associate light and sickness but not taste and sickness.
- 6 Instrumental conditioning deals with situations in which the response operates on the environment rather than being elicited by an unconditioned stimulus. The earliest systematic studies were performed by Thorndike, who showed that animals engage in trial-and-error behavior and that any behavior that is followed by reinforcement is strengthened; this is known as the law of effect.
- 7 In Skinner's experiments, typically a rat or pigeon learns to make a simple response, such as pressing a lever, to obtain reinforcement. The rate of response is a useful measure of response strength. Shaping is a training procedure that is used when the desired response is novel; it involves reinforcing only variations in response that deviate in the direction desired by the experimenter.
- 8 A number of phenomena can increase the generality of instrumental conditioning. One is conditioned reinforcement, in which a stimulus associated with a reinforcer acquires its own reinforcing properties. Other relevant phenomena are generalization and discrimination; organisms generalize responses to similar situations, although this generalization can be brought under the control of a discriminative stimulus. Finally, there are schedules of reinforcement. Once a behavior is established, it can be maintained when it is reinforced only part of the time. Exactly when the reinforcement comes is determined by its schedule; the basic types of reinforcement schedules are fixed ratio, variable ratio, fixed interval, and variable interval schedules.
- 9 There are three kinds of aversive conditioning. In punishment, a response is followed by an aversive event, which results in the response being suppressed. In escape, an organism learns to make a response in order to terminate an ongoing aversive event. In avoidance, an organism learns to make a response to prevent the aversive event from even starting.
- 10 Cognitive factors play a role in instrumental conditioning. For instrumental conditioning to occur, the organism must believe that reinforcement is at least partly under its control; that is, it must perceive a contingency between its responses and the reinforcement. Biological constraints are also a factor in instrumental conditioning. There are constraints on what reinforcers can be associated with what responses. With pigeons, when the reinforcement is food, learning is faster if the response is pecking a key rather than flapping the wings, but when the reinforcement is termination of shock, learning is faster when the response is wing flapping rather than pecking a key.

- 11 According to the cognitive perspective, the crux of learning is an organism's ability to represent aspects of the world mentally and then operate on these mental representations rather than on the world itself. In complex learning, the mental representations depict more than associations, and the mental operations may constitute a strategy. Studies of complex learning in animals indicate that rats can develop a cognitive map of their environment, as well as acquire abstract concepts such as cause.
- 12 Learning through imitation and observation happens as a result of vicarious reinforcement: by observing a model's behavior, the imitator expects to be reinforced just like the model was. Humans learn many complex and social behaviors through observational learning.
- 13 When learning relationships between stimuli that are not perfectly predictive, people often invoke prior beliefs. This can lead to the detection of relationships that are not objectively present (spurious associations). When the relationship is objectively present, having a prior belief about it can lead to overestimating its predictive strength; when an objective relationship conflicts with a prior belief, the learner may favor the prior belief.
- These effects demonstrate top-down processing in learning.
- 14 The neural mechanisms of non-associative forms of learning have been studied in invertebrate slugs. Habituation is mediated by a decrease in synaptic transmission, and sensitization by an increase in transmission. Regression and growth, respectively, of synapses are also involved in these types of learning.
- 15 Synapses in the mammalian brain take part in storing information during learning. The cerebellum is particularly important for motor conditioning, and the amygdala is essential for emotional conditioning. Increases in synaptic transmission, termed long-term potentiation, are involved in these learning processes.
- 16 Intrinsically motivated individuals are more persistent at a task than individuals motivated by an external reward. Experiments show that adding external rewards can lead to overjustification of the behavior. As a consequence, the individual attributes his or her engagement with the task to the external rewards. This is damaging to intrinsic motivation, as well as to performance. Complex tasks are best accomplished if the individual perceives a sense of control and self-determination.

CORE CONCEPTS

cognitive behavior therapy	response generalization	learned helplessness
behavior therapy	stimulus discrimination	latent learning
learning	second-order conditioning	cognitive map
non-associative learning	temporal contiguity	observational learning
habituation	contingency	self-efficacy
sensitization	learned taste aversion	Hebbian learning rule
associative learning	instrumental conditioning	neural plasticity
classical conditioning	insight	long-term depression (LTD)
unconditioned response	trial-and-error learning	long-term potentiation (LTP)
unconditioned stimulus	law of effect	synaptic plasticity
neutral stimulus	positive and negative reinforcement	arousal
conditioned stimulus	positive and negative punishment	Yerkes-Dodson law
conditioned response	shaping	exploratory behavior
drug tolerance	conditioned reinforcer	incentive
acquisition	fixed and variable ratio schedule	intrinsic motivation
learning curve	fixed and variable interval schedule	extrinsic motivation
extinction	escape learning	overjustification effect
spontaneous recovery	avoidance learning	

WEB RESOURCES



<http://www.atkinsonhilgard.com/>

Take a quiz, try the activities and exercises, and explore web links.

<http://www.healthinfluence.com/Primer/classical.htm>

Learn more about how classical conditioning can influence your actions.

<http://psych.athabascau.ca/html/prtut/reinpair.htm>

This site discusses positive reinforcement. After you have read the information, try the practice exercise to test your knowledge.

http://nobelprize.org/nobel_prizes/medicine/laureates/1904/pavlov-bio.html

Read a detailed biography of Pavlov here at the official site of the Noel Prize Organization.

CD-ROM LINKS



Psyk.Trek 3.0

Check out CD Unit 5, Learning

- 5a Overview of classical conditioning
- 5b Basic processes in classical conditioning
- 5c Overview of operant conditioning
- 5d Schedules of reinforcement
- 5e Reinforcement and punishment
- 5f Avoidance and escape learning