Chapter 1

Learning

Learning is defined in this chapter as a relatively permanent change in behaviour as a result of experience. It is the lifelong, dynamic process by which individuals acquire new knowledge or skills and alter their thoughts, feelings, attitudes, and actions. Learning enables individuals to adapt to demands and changing circumstances and is crucial in health care—whether for clients and/or patients acquiring new information and learning necessary skills to manage a diagnosis and/or chronic health condition, or for healthcare staff understanding particular individual learnt health behaviours e.g., phobias and how this may influence interactions with clients and patients. This chapter will cover *learning theories* and apply principles that will describe, explain, or predict how people learn.

After studying this chapter you should be able to:

- Describe Pavlov's classical conditioning model
- Describe the principles of classical conditioning
- Apply principles of classical conditioning to everyday life
- Describe the principles of operant conditioning
- Distinguish between operant conditioning from classical conditioning
- Describe Thorndike's law of effect
- Distinguish reinforcement from punishment and its outcome on behaviour
- Describe the four schedules of reinforcement, and the response pattern associated with each
- Describe cognitive models of learning
- Describe biological influences on learning.

The three most famous figures in the psychology of learning were each colourful characters in their own way. The discoverer of classical conditioning, Ivan Pavlov, was a notoriously compulsive fellow. He ate lunch every day at precisely 12 noon, went to bed at exactly the same time every night and departed St Petersburg, Russia, for holiday the same day every year. Pavlov was also such a rapid walker that his wife frequently had to run frantically to keep up with him. The life of the founder of behaviourism, John B. Watson, was rocked with scandal. Despite becoming one of the world's most famous psychologists, he was unceremoniously booted out of Johns Hopkins University for having an affair with his graduate student, Rosalie Rayner. Watson also had rather unusual ideas about parenting; for example, he believed that all parents should shake hands with their children before bedtime. B. F. Skinner, the founder of radical behaviourism, was something of a prankster in his undergraduate years at Hamilton College in New York. He and a friend once spread a false rumour that comedian Charlie Chaplin was coming to campus. This rumour nearly provoked a riot when Chaplin did not materialise as expected.

By **learning**, we mean a change in an organism's behaviour or thought as a result of experience. When we learn our brain changes along with our behaviours. Remarkably, your brain is physically different now than it was just a few minutes ago, because it underwent chemical changes that allowed you to learn novel facts.

Learning
Change in an organism's behaviour or thought as a result of experience.

Habituation
Process of responding less strongly over time to repeated stimuli.

Learning lies at the heart of just about every domain of psychology. Virtually all behaviours are a complex stew of genetic predispositions and learning. Without learning, we would be unable to do much; we could not walk, talk or read an introductory psychology textbook chapter about learning.

Psychologists have long debated how many distinct types of learning there are. We are not going to try to settle this controversy here. Instead, we will review several types of learning that psychologists have studied in depth, starting with the most basic.

Even foetuses as young as 32 weeks have displayed habituation. While they first experience shock when a gentle vibrator is applied to the mother's stomach, the shock later becomes a mere annoyance.
(**Source:** Olga Makarova/Dreamstime.com.)

Before we do, place your brain on pause, put down your pen or highlighter, close your eyes and attend to several things that you almost never notice: the soft buzzing of the lights in the room, the feel of your clothing against your skin, the sensation of your tongue on your teeth or lips. Unless someone draws our attention to these stimuli, we do not even realise they are there, because we have learned to ignore them. **Habituation** is the process by which we respond less strongly over time to repeated stimuli. It helps to explain why loud snorers can sleep peacefully through the night while keeping their irritated roommates wide awake. Chronic snorers have become so accustomed to the sound of their own snoring that they no longer notice it.

Habituation is the simplest and probably earliest form of learning to emerge in humans. Foetuses as young as 32 weeks display habituation when we apply a gentle vibrator to the mother's stomach. At first, the foetus jerks around in response to the stimulus, but after repeated vibrations it stops moving (Morokuma et al., 2004). What was first a shock to the foetus's system later became a mere annoyance that it could safely ignore.

In research that earned him the Nobel Prize in 2000, neurophysiologist Eric Kandel uncovered the biological mechanism of habituation of *Aplysia*, a 12-centimetre-long sea slug. Prick an *Aplysia* on a certain part of its body and it retracts its gill in a defensive manoeuvre. Touch *Aplysia* in the same spot repeatedly and it begins to ignore the stimulus. This habituation, Kandel found, is accompanied by a progressive decrease in the release of the neurotransmitter serotonin at *Aplysia*'s synapses (Siegelbaum, Camardo & Kandel, 1982). This discovery helped psychologists unravel the neural bases of learning (see Figure 1.1).

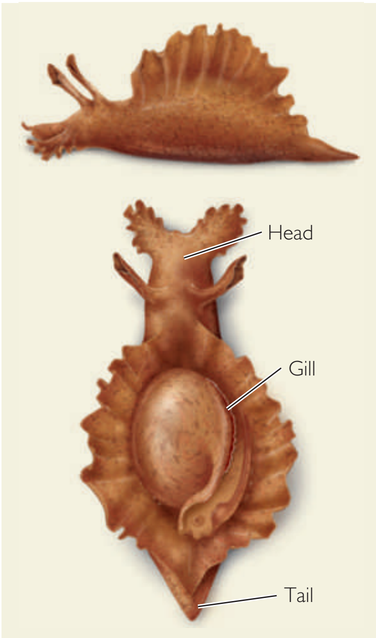

Figure 1.1 Habituation in a simple animal.

Aplysia californicus is a sea slug about 12 centimetres long that retracts its gill when pricked, but then habituates (stops retracting its gill) if pricked repeatedly.

Psychologists have studied habituation by measuring—of all things—sweat. Because perspiration on our fingertips is a good indicator of anxiety (Fowles, 1980), scientists measure it by using an electrical conductivity measure called the skin conductance response. Most research shows that our hands stop sweating sooner for weak stimuli than they do for strong stimuli, meaning that weak stimuli stop producing anxiety fairly quickly compared with strong stimuli. In the case of very strong stimuli, like painful electric shocks, we often see no habituation at all—people continue to sweat anxiously at the same high levels—even across many trials (Lykken et al., 1988).

This research suggests that habituation makes good sense from an evolutionary standpoint. We wouldn't want to attend to every tiny sensation that comes across our mental radar screens, because most pose no threat. Yet we wouldn't want to habituate to stimuli that could be dangerous. Fortunately, not all repeated stimuli lead to habituation, only those that we deem safe or worth ignoring do.

Some cases of repeated exposure to stimuli lead to *sensitisation*—that is, responding more strongly over time—rather than habituation. Sensitisation is most likely when a stimulus is dangerous, irritating, or both. *Aplysia* show sensitisation as well as habituation. Have you ever tried to study when the person next to you was whispering, and the whispering kept getting more annoying to the point that you could not concentrate? If so, you have experienced sensitisation.

Classical conditioning

The story of habituation could hardly be more straightforward. We experience a stimulus, respond to it and then stop responding after repeated exposure. We have learned something significant, but we have not learned to forge connections between two stimuli. Yet a great deal of learning depends on associating one thing with another. If we never learned to connect one stimulus, such as the appearance of an apple, with another stimulus, such as its taste, our everyday life would be a world of disconnected sensory experiences.

In the nineteenth century, a school of thinkers called the British Associationists believed that we acquire virtually all of our knowledge by *conditioning*—that is, by forming associations among stimuli. Once we form these links, like the connection between our mother's voice and her face, we need only recall one element of the pair to retrieve the other. The British Associationists believed that simple connections provided the mental building blocks for all of our more complex ideas. Their armchair conjectures were to be confirmed by a pioneering Russian physiologist who demonstrated these processes of association in the laboratory.

Pavlov's discovery of classical conditioning

That physiologist's name was Ivan Pavlov. Pavlov's primary research was on digestion in dogs—in fact, his discoveries concerning digestion, not classical conditioning, earned him the Nobel Prize in 1904. Pavlov placed dogs in a harness and inserted a collection tube into their salivary glands to study their digestive responses to meat powder. In doing so, he observed something unexpected: dogs began salivating (more informally, they started to drool), not only at the meat powder itself, but at previously neutral stimuli that had become associated with it, such as the research assistants who brought in the powder. Indeed, the dogs even salivated to the sound of these assistants' footsteps as they approached the laboratory. The dogs seemed to be anticipating the meat powder and responding to stimuli that signalled its arrival.

We call this process of association **classical conditioning** (or **Pavlovian or respondent conditioning**): a form of learning in which animals come to respond to a previously neutral stimulus that had been paired with another stimulus that elicits an automatic response. Pavlov's initial observations were merely anecdotal; so, like any good scientist, he put his informal observations to a more rigorous test.

Classical (Pavlovian or respondent) conditioning Form of learning in which animals come to respond to a previously neutral stimulus that had been paired with another stimulus that elicits an automatic response.

The classical conditioning phenomenon

This is how Pavlov first demonstrated classical conditioning systematically (see also Figure 1.2).

1. He started with an initially neutral stimulus, one that didn't elicit any particular response. In this case, Pavlov used a metronome, a clicking pendulum that keeps time (in other studies, Pavlov used a tuning fork or whistle; contrary to urban legend, Pavlov did not use a bell).
2. He then paired the neutral stimulus again and again with an **unconditioned stimulus (UCS)**, a stimulus that elicits an automatic—that is, a reflexive—response. In the case of Pavlov's dogs, the UCS was the meat powder, and the automatic, reflexive response it elicits is the **unconditioned response (UCR)**. For the dogs, the UCR was salivation. The key point is that the animal does not need to learn to respond to the UCS with the UCR: dogs naturally drool in response to food. The animal generates the UCR without any training at all, because the response is a product of nature, not nurture.
3. As Pavlov repeatedly paired the CS and the UCS, he observed something remarkable. If he now presented the metronome alone, it elicited a response, namely salivation. This new response is the **conditioned response (CR)**: a response previously associated with a non-neutral stimulus that comes to be elicited by a neutral stimulus. Lo and behold, learning has occurred. The metronome had become a **conditioned stimulus (CS)**—a previously neutral stimulus that comes to elicit a conditioned response as a result of its association with an unconditioned stimulus. The dog, which previously did nothing when it heard the metronome except perhaps turn its head towards it, now salivates when

Factoid

Classical conditioning may occur not only in animals but in plants. One researcher found that a *Mimosa* plant that folds its leaves (UCR) when touched (UCS) can be conditioned to fold its leaves (CR) in response to a change in lighting condition (CS) that has been repeatedly paired with a touch (Haney, 1969). Nevertheless, this finding is scientifically controversial.

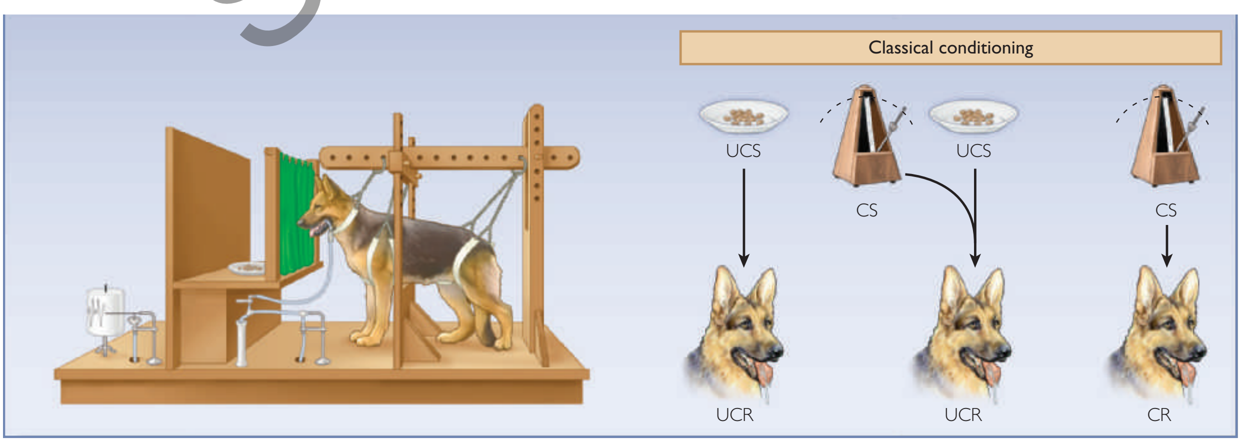

Figure 1.2 Pavlov's classical conditioning model.

The UCS (meat powder) is paired with the CS (metronome clicking) and produces the UCR (salivation). Then the CS is presented alone, and the CR (salivation) occurs.

Factoid

Backward conditioning—in which the UCS is presented *before* the CS—is extremely difficult to achieve. Because the CS fails to predict the UCS and the UCR often begins before the CS has even occurred, organisms have difficulty using the CS to anticipate the UCS.

it hears the metronome. The CR, in contrast to the UCR, is a product of nurture, not nature.

In most cases, the CR is similar to the UCR but it is rarely identical to it. For example, Pavlov found that dogs salivated less in response to the metronome (the CS) than to the meat powder (the UCS).

Few findings in psychology are as replicable as classical conditioning. We can apply the classical conditioning paradigm to just about any animal which has an intact nervous system, and demonstrate it repeatedly without fail. If only all psychological findings were so dependable!

Principles of classical conditioning

We will next explore the major principles underlying classical conditioning. Pavlov noted, and many others have since confirmed, that classical conditioning occurs in three phases—acquisition, extinction and spontaneous recovery. In addition, as we will see, once classical conditioning to a stimulus occurs, it often extends to a host of related stimuli, making its everyday life influence surprisingly powerful.

Acquisition

In **acquisition**, we gradually learn—or acquire—the CR. If you look at Figure 1.3(a), you will see that, as the CS and the UCS are paired over and over again, the CR increases progressively in strength. The steepness of this curve varies somewhat depending on how close together in time the CS and UCS are presented. In general, the closer in time the pairing of the CS and the UCS, the faster learning occurs, with about a half-second delay typically being the optimal pairing for learning. Longer delays usually decrease the speed and strength of the organism's response.

Extinction

In a process called **extinction**, the CR decreases in magnitude and eventually disappears when the CS is repeatedly presented alone—that is, without the UCS (see Figure 1.3[b]). After numerous presentations of the metronome without the meat power, Pavlov's dogs eventually stopped salivating. Most psychologists once believed that extinction was similar

Unconditioned stimulus (UCS) Stimulus that elicits an automatic response.

Unconditioned response (UCR) Automatic response to a non-neutral stimulus that does not need to be learned.

Conditioned response (CR) Response previously associated with a non-neutral stimulus that is elicited by a neutral stimulus through conditioning.

Conditioned stimulus (CS) Initially neutral stimulus.

Acquisition Learning phase during which a conditioned response is established.

Extinction Gradual reduction and eventual elimination of the conditioned response after the conditioned stimulus is presented repeatedly without the unconditioned stimulus.

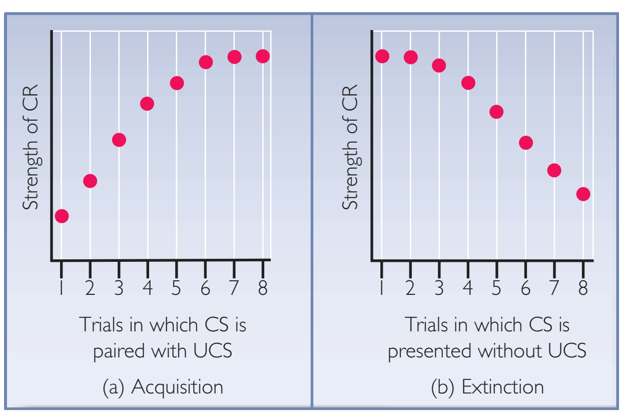

Figure 1.3 Acquisition and extinction.

Acquisition is the repeated pairing of the UCS and the CS, increasing the CR's strength (a). In extinction, the CS is presented again and again without the UCS, resulting in the gradual disappearance of the CR (b).

to forgetting: the CR fades away over repeated trials, just as many memories gradually decay. Yet the truth is more complicated and interesting than that. Extinction is an active, rather than passive, process. During extinction a new response, which in the case of Pavlov's dogs was the *absence* of salivation, gradually 'writes over' or inhibits the CR, namely salivation. The extinguished CR does not vanish completely; it is merely overshadowed by the new behaviour. This contrasts with most forms of traditional forgetting, in which the memory itself disappears. Interestingly, Pavlov had proposed this hypothesis in his writings, although few people believed him at the time. How do we know he was right? Read on.

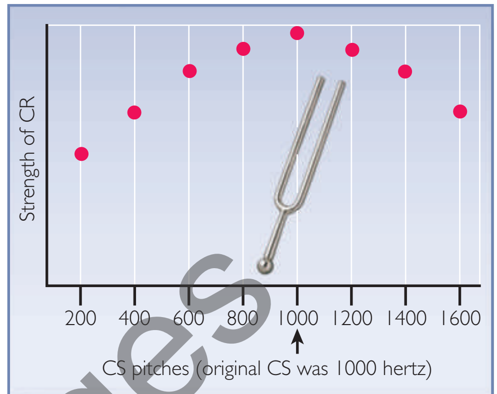

Figure 1.4 Generalisation gradient.

The more similar to the original CS the new CS is (for example, Pavlov using a tone pitched close to the original tone's pitch), the stronger the CR will be.

Spontaneous recovery

In a phenomenon called **spontaneous recovery**, a seemingly extinct CR reappears (often in a somewhat weaker form) if the CS is presented again. It is as though the CR were lurking in the background, waiting to appear following another presentation of the CS. In a classic study, Pavlov (1927) presented the CS (the tone from a metronome) alone again and again, and extinguished the CR (salivation) because there was no UCS (mouth-watering meat powder) following it. Two hours later, he presented the CS again and the CR returned. The animal had not really forgotten the CR, just suppressed it.

A related phenomenon is the **renewal effect**, which occurs when we extinguish a response in a setting different from the one in which the animal acquired it. When we restore the animal to the original setting, the extinguished response reappears (Bouton, 1994). The renewal effect may help to explain why people with *phobias*—intense, irrational fears—who have overcome their phobias often experience a reappearance of their symptoms when they return to the environment in which they acquired their fears (Denniston, Chang & Miller, 2003). Even though it may sometimes lead to a return of phobias, the renewal effect is often adaptive. If you have been bitten by a snake in one part of a forest, it makes sense to experience fear when you find yourself there again, even years later. That same snake or his slithery descendants may still be lying in wait in the same spot.

Spontaneous recovery
Sudden re-emergence of an extinct conditioned response after a delay in exposure to the conditioned stimulus.

Renewal effect
Sudden re-emergence of a conditioned response following extinction when an animal is returned to the environment in which the conditioned response was acquired.

Stimulus generalisation

Pavlov found that following classical conditioning his dogs salivated not merely to the original metronome sound, but to sounds similar to it. This phenomenon is **stimulus generalisation**: the process by which CSs that are similar, but not identical, to the original CS elicit a CR. Stimulus generalisation occurs along a *generalisation gradient*: the more similar to the original CS the new CS is, the stronger the CR will be (see Figure 1.4). Pavlov found that his dogs showed their largest amount of salivation to the original sound, with progressively less salivation to sounds that were less and less similar to it in pitch. Stimulus generalisation is typically adaptive, because it allows us to transfer what we have learned to new things. For example, once we have learned to drive our own car, we can borrow a friend's car without needing a full tutorial on how to drive it.

Stimulus generalisation
Process by which conditioned stimuli similar, but not identical, to the original conditioned stimulus elicit a conditioned response.

Stimulus discrimination

Stimulus discrimination is the flip side of the coin to stimulus generalisation; it occurs when we exhibit a less pronounced CR to CSs that differ from the original CS. Stimulus discrimination helps us understand why we can enjoy scary movies. Although we may hyperventilate a bit while watching sharks circle the capsized sailors in the movie *The Reef*,

Stimulus discrimination
Displaying a less pronounced conditioned response to conditioned stimuli that differ from the original conditioned stimulus.

Higher-order conditioning Developing a conditioned response to a conditioned stimulus by virtue of its association with another conditioned stimulus.

we would respond even more strongly if we went on a shark dive while on holiday. We have learned to discriminate between a motion picture stimulus and the real-world version of it, and to modify our response as a result.

Higher-order conditioning

Taking conditioning a step further, organisms learn to develop conditioned associations to CSs that are associated with the original CS. If after conditioning a dog to salivate to a tone, we pair a picture of a circle with that tone, a dog eventually salivates to the circle as well as to the tone. That is **higher-order conditioning**: the process by which organisms develop classically conditioned responses to CSs associated with the original CS (Gewirtz & Davis, 2000, pp. 257–266). There are several levels of higher-order conditioning, defined by the number of steps between the CS and the UCS. As we might expect, second-order conditioning—in which a new CS is paired with the original CS—tends to be weaker than garden-variety classical conditioning, and third-order conditioning—in which a third CS is paired with the second-order CS—is even weaker. Fourth-order conditioning and beyond are typically difficult or impossible to achieve.

A person hiking through the forest may experience fear when he approaches an area if he has previously spotted a dangerous animal there.
(**Source:** Yuri Arcurs/Dreamstime.)

Higher-order conditioning allows us to extend classical conditioning to a host of new stimuli. It helps explain why we feel hungry after someone merely says 'kebab' after a late-night party. We have already come to associate the sight, sound and smell of a kebab with satisfying our hunger, and we eventually came to associate the word 'kebab' with these CSs.

Applications of classical conditioning to daily life

Classical conditioning applies to myriad domains of everyday life. We consider five here: advertising, the acquisition of fears and phobias, the acquisition of fetishes, and disgust reactions.

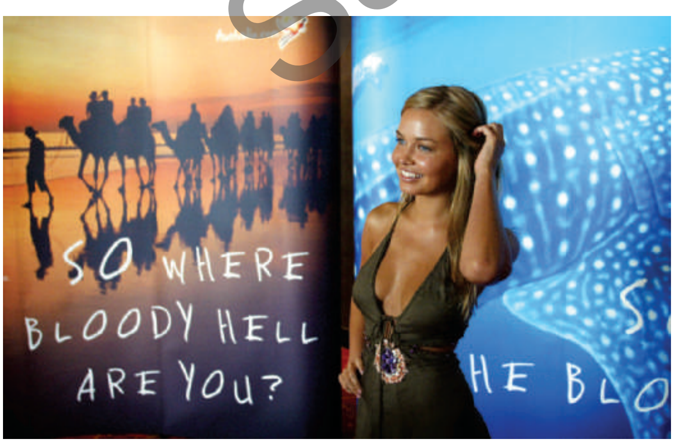

Advertisers use higher-order classical conditioning to get customers to associate their products with an inherently enjoyable stimulus.
(**Source:** Toby Zerna/Newspix/News Ltd.)

Classical conditioning and advertising

Few people grasp the principles of classical conditioning, especially higher-order conditioning, better than advertisers. By repeatedly pairing the sights and sounds of products with photographs of spunky football players and bikini-clad models, marketing whizzes try to establish classically conditioned connections between their brands and positive emotions. They do so for a good reason: it works.

Researcher Gerald Gorn (1982) paired slides of either blue or beige pens (the CSs) with music that participants had rated as either enjoyable or not enjoyable (the UCSs). Then he gave participants the opportunity to select a pen upon departing the lab. Whereas 79 per cent of participants who heard music they liked picked the pen that had been paired with the music, only 30 per cent of those who heard music they disliked picked the pen that had been paired with the music.

Nevertheless, not all researchers who have paired products with pleasurable stimuli have succeeded in demonstrating classical conditioning effects (Gresham & Shimp, 1985, pp. 10–49; Smith, 2001). But many of the negative findings are open to a rival explanation: latent inhibition. **Latent inhibition** refers to the fact that when we have experienced a CS alone many times, it is difficult to classically condition it to another stimulus (Vaitl & Lipp, 1997, pp. 85–93). Because some investigators who failed to obtain classical conditioning effects for products relied on brands with which participants were already familiar, their negative findings may be attributable to latent inhibition. Indeed, when researchers have used novel brands, they have generally been able to show classical conditioning effects (Stuart, Shimp & Engle, 1987, pp. 334–349).

Latent inhibition
Difficulty in establishing classical conditioning to a conditioned stimulus we have repeatedly experienced alone; that is, without the unconditioned stimulus.

Classic study in which a nine-month-old boy was conditioned to fear white furry objects. Here, Little Albert, with John B. Watson and Rosalie Rayner, is crying in response to a Santa Claus mask.
(**Source:** Benjamin Harris, PhD.)

The acquisition of fears and phobias: the strange tale of Little Albert

Can classical conditioning help explain how we come to fear or avoid stimuli? John B. Watson, the founder of behaviourism, answered this question in 1920 when he and his graduate student, Rosalie Rayner, performed what must be regarded as one of the most ethically questionable studies in the history of psychology.

Watson and Rayner (1920) set out in part to show that the Freudian view of phobias, which proposed that phobias stem from deep-seated conflicts buried in the unconscious, was wrong. To do so, they recruited a nine-month-old infant who will be forever known in the psychological literature as Little Albert. Little Albert was fond of furry little creatures, like white rats. But Watson and Rayner were about to change that.

Watson and Rayner first allowed Little Albert to play with a rat. But only seconds afterward, Watson snuck up behind Little Albert and struck a gong with a steel hammer, creating an ear-splitting noise, which startled him out of his wits and made him cry. After seven such pairings of the rat and UCS (the loud sound from the gong), Little Albert displayed a CR (crying) to the rat alone, demonstrating that the rat had now become a CS. This conditioned response was still present when Watson and Rayner exposed Little Albert to the rat five days later. Moreover, Watson and Rayner observed that Little Albert had become a victim of stimulus generalisation, coming to fear not merely rats, but also a rabbit, a dog, a furry coat and, to a lesser extent, a Santa Claus mask and John B. Watson's hair. Fortunately, Little Albert also demonstrated at least some stimulus discrimination, as he did not display much fear towards cotton balls or the hair of Dr Watson's research assistants.

Incidentally, no one knows what became of poor Little Albert. His mother withdrew him from the study about a month after it began, never to be heard from again. Needless to say, because inducing a phobia-like condition in an infant raises a host of serious ethical questions, Watson and Rayner's Little Albert study would never get past a modern-day university ethical review board.

Stimulus generalisation, like that experienced by Little Albert, allows our learning to be remarkably flexible—which is often, although not always, a good thing. It allows us to develop fears of many stimuli. Certain phobias, like those of snakes, spiders, heights,

Table 1.1 Phobias galore. This sampling of phobias—some common, some exceedingly rare—illustrates just how enormously varied people's fears can be. Many of these phobias can be acquired at least partly by classical conditioning.

Alliumphobia: Fear of garlic	*Melissophobia:* Fear of bees
Arachibutyrophobia: Fear of peanut butter sticking to the roof of your mouth	*Ophidiophobia:* Fear of snakes
Brontophobia: Fear of thunderstorms	*Peladophobia:* Fear of bald people
Bufonophobia: Fear of toads	*Pentheraphobia:* Fear of one's mother-in-law
Catoptrophobia: Fear of mirrors	*Pogonophobia:* Fear of beards
Elurophobia: Fear of cats	*Rhytiphobia:* Fear of getting wrinkles
Epistaxiophobia: Fear of nose bleeds	*Samhainophobia:* Fear of Halloween
Latrophobia: Fear of doctors	*Taphephobia:* Fear of being buried alive
Lachanophobia: Fear of vegetables	*Xyrophobia:* Fear of razors

Michael Domjan and his colleagues used classical conditioning to instil a fetish in male quails.
(**Source:** Dr Michael Domjan.)

water and blood, are considerably more widespread than others (www.anxietyaustralia.com.au). And some are downright strange, as Table 1.1 illustrates.

The good news is that if classical conditioning can contribute to our acquiring phobias, it can also contribute to our conquering them. Mary Cover Jones, a student of Watson, treated a three-year-old named Little Peter, who had a phobia of rabbits. Jones (1924) treated Peter's fear successfully by gradually introducing him to a white rabbit while giving him a piece of his favourite lolly. As she moved the rabbit increasingly close to him, the sight of the rabbit eventually came to elicit a new CR: pleasure rather than fear. Modern-day psychotherapists, although rarely feeding their clients lollies, use similar practices to eliminate phobias. They may pair feared stimuli with relaxation or other pleasurable stimuli (Wolpe, 1990).

Fetishes

Fetishism
Sexual attraction to non-living things.

On the flip side of the coin from phobias, **fetishism**—sexual attraction to non-living things—may also arise in part from classical conditioning (Akins, 2004; Hoffmann, 2011). Like phobias, fetishes come in a bewildering variety of forms: shoes, stockings, dolls, stuffed animals, automobile engines (yes, that's right), and just about anything else (Lowenstein, 2002).

Although the origins of human fetishes are controversial, Michael Domjan and his colleagues were successful in classically conditioning fetishes in male Japanese quails. In one study, they presented male quails with a cylindrical object made of terrycloth, followed by a female quail with which they happily mated. After 30 such pairings, about half of the male quails attempted to mate with the cylindrical object when it appeared alone (Köksal et al., 2004). Although the generalisability of these findings to humans is unclear, there is good evidence that at least some people develop fetishes by the repeated pairing of neutral objects with sexual activity (Rachman & Hodgson, 1968; Weinberg, Williams & Calhan, 1995).

Disgust reactions

Imagine that a researcher asked you to eat a piece of fudge. No problem, right? Well, now imagine the fudge were shaped like dog faeces. If you are like most participants in the studies of Paul Rozin and his colleagues, you would hesitate (D'Amato, 1998; Rozin, Millman & Nemeroff, 1986).

Rozin and his colleagues have found that we acquire disgust reactions with surprising ease. In most cases, these reactions are probably the product of classical conditioning, because CSs associated with disgusting UCSs come to elicit disgust themselves. In many cases, disgust reactions are tied to stimuli that are biologically important to us, such as animals or objects that are dirty or potentially poisonous (Rozin & Fallon, 1987).

In another study, Rozin and his collaborators asked participants to drink from two glasses of water, both of which contained sugar (sucrose). In one case, the sucrose came from a bottle labelled 'Sucrose'; in another, it came from a bottle labelled 'Sodium Cyanide, Poison'. The investigators told participants that both bottles were completely safe. They even asked participants to select which label went with which glass, proving the labels were meaningless. Even so, participants were hesitant to drink from the glass that contained the sucrose labelled as poisonous (Rozin, Markwith & Ross, 1990, pp. 383–384). Participants' responses in this study were irrational, but perhaps understandable: they were probably relying on the heuristic 'better safe than sorry'. Classical conditioning helps keep us safe, even if it goes too far on occasion.

Some researchers have suggested that classical conditioning may similarly help to explain some human fetishes.
(**Source:** Photostogo/Photolibrary.)

Assess your knowledge | FACT or FICTION?

1. Habituation to meaningless stimuli is generally adaptive. **(True/False)**
2. In classical conditioning, the conditioned stimulus (CS) initially yields a reflexive, automatic response. **(True/False)**
3. Conditioning is most effective when the CS precedes the UCS by a short period of time. **(True/False)**
4. Extinction is produced by the gradual 'decay' of the CR over time. **(True/False)**

Answers: (1) T; (2) F; (3) T; (4) F

Operant conditioning

What do the following four examples have in common?

- Using bird feed as a reward, a behavioural psychologist teaches a pigeon to distinguish paintings by Monet from paintings by Picasso. By the end of the training, the pigeon is a veritable art aficionado.
- Using fish as a treat, a trainer teaches a dolphin to jump out of the water, spin three times, splash in the water and propel itself through a hoop.
- In his initial attempt at playing tennis, a frustrated 12-year-old hits his opponent's serve into the net the first 15 times. After two hours of practice, he returns his opponent's serve successfully more than half the time.
- A hospitalised patient with dissociative identity disorder (formerly known as multiple personality disorder) displays features of an 'alter' personality whenever staff members pay attention to him. When they ignore him, his alter personality seemingly vanishes.

The answer: all are examples of **operant conditioning**. The first, incidentally, comes from an actual study (Watanabe, Sakamoto & Wakita, 1995). Operant conditioning is learning

Operant conditioning
Learning controlled by the consequences of the organism's behaviour.

controlled by the consequences of the organism's behaviour. In each of these examples, superficially different as they are, the organism's behaviour is shaped by what comes after it, namely reward. Psychologists also refer to operant conditioning as 'instrumental conditioning', because the organism's response serves an instrumental function. That is, the organism 'gets something' out of the response—like food, sex, attention or avoiding something unpleasant.

Behaviourists refer to the behaviours emitted by the animal to receive a reward as *operants*, because the animal 'operates' on its environment to get what it wants. Dropping $1.70 into a drink machine is an operant, as is asking an appealing classmate to go out for coffee after the tutorial. In the first case, our reward is a refreshing drink, and in the second, an enjoyable afternoon.

Distinguishing operant conditioning from classical conditioning

Operant conditioning differs from classical conditioning in three important ways, which we have highlighted in Table 1.2.

Law of effect
Principle asserting that if a stimulus followed by a behaviour results in a reward, the stimulus is more likely to elicit the behaviour in the future.

Table 1.2 Key differences between operant and classical conditioning

	Classical conditioning	Operant conditioning
Target behaviour is …	Elicited automatically	Emitted voluntarily
Reward is …	Provided unconditionally	Contingent on behaviour
Behaviour depends primarily on …	Autonomic nervous system	Skeletal muscles

How could operant conditioning principles explain this boy's ability to improve his tennis game with practice?
(**Source:** Photostogo/Photolibrary.)

1. In classical conditioning, the organism's response is *elicited*—that is, 'pulled out' of the organism by the UCS, and later the CS. Remember that in classical conditioning the UCR is a reflexive and automatic response that does not require training. In operant conditioning, the organism's response is *emitted*—that is, generated by the organism in a seemingly voluntary fashion.
2. In classical conditioning, the animal's reward is independent of what it does. Pavlov gave his dogs the meat powder regardless of whether, or how much, they salivated. In operant conditioning, the animal's reward is contingent—that is, dependent—on what it does. If the animal does not emit a response in an operant conditioning paradigm, it comes out empty-handed (or in the case of a dog, empty-pawed).
3. In classical conditioning, the organism's responses depend primarily on the autonomic nervous system. In operant conditioning, the organism's responses depend primarily on the skeletal muscles. That is, in contrast to classical conditioning, in which learning involves changes in heart rate, breathing, perspiration and other bodily systems, in operant conditioning learning involves changes in voluntary motor behaviour.

The law of effect

The famous **law of effect**, put forth by psychologist E. L. Thorndike, forms the basis of much of operant conditioning: if a response, in the presence of a stimulus, is followed by a satisfying state of affairs, the bond between stimulus and response will be strengthened.

This statement means that if we are rewarded for a response to a stimulus, we are more like to repeat that response to the stimulus in the future. Psychologists sometimes refer to early forms of behaviourism as S-R psychology ('S' stands for *stimulus*, 'R' for *response*). According to S-R theorists, most of our complex behaviours reflect the accumulation of associations between stimuli and responses: the sight of a close friend and saying hello, or the smell of a delicious hamburger and reaching for it on our plate. S-R theorists maintain that almost everything we do voluntarily—driving a car, eating a sandwich or planting a kiss on someone's lips—results from the gradual build-up of S-R bonds due to the law of effect.

Thorndike (1898) discovered the law of effect in a classic study of cats and puzzle boxes. He placed a hungry cat in a box and put a tantalising piece of fish just outside. To escape from the box, the cat needed to hit upon (literally) the right solution, which was pressing on a lever or pulling on a string inside the box (see Figure 1.5). When Thorndike first placed the cat in the puzzle box, it typically flailed around aimlessly in a frantic effort to escape. Then, by sheer accident, the cat eventually found the correct solution, scurried out of the box and gobbled up its delectable treat. Thorndike wanted to find out what would happen to the cat's behaviour over time. Once it figured out the solution to the puzzle, would it then get it right every time?

He found that the cat's time to escape from the puzzle box decreased *gradually* over 60 trials. There was no point at which the cat abruptly realised what it needed to do to escape. According to Thorndike, his cats were learning by trial and error through the steady build-up of associations. Indeed, Thorndike and many other S-R theorists went so far as to conclude that all learning, including all human learning, occurs by trial and error. For them, S-R bonds are gradually 'stamped into' the organism by reward.

These findings, Thorndike concluded, provide a crushing blow to the hypothesis that cats learn by **insight**—that is, by grasping the nature of the problem. Had his cats possessed insight into the nature of the problem, the results presumably would have looked like what we see in Figure 1.6. This figure illustrates what psychologists term the 'Aha! reaction': 'Aha—I got it!' Once the animal solves the problem, it gets it correct just about every time after that. Yet Thorndike never found an 'Aha!' moment: the time to a correct solution decreased only gradually.

Insight
Grasping the nature of a problem.

B. F. Skinner and reinforcement

Thorndike's pioneering discoveries laid the groundwork for research on operant conditioning. B. F. Skinner then kicked it up a notch using electronic technology.

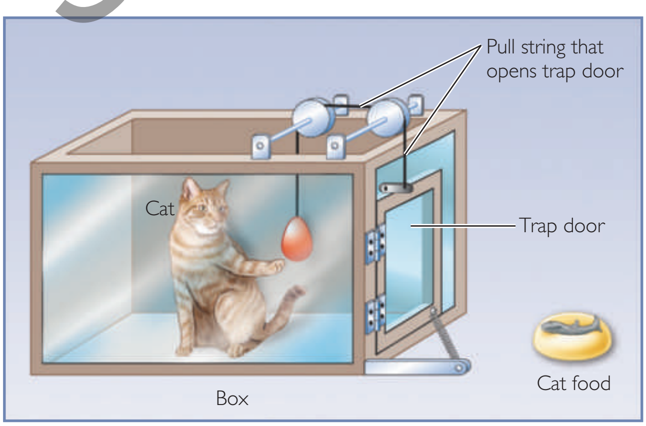

Figure 1.5 Thorndike's puzzle box.

Thorndike's classic puzzle box research seemed to suggest that cats solve problems solely through trial and error.

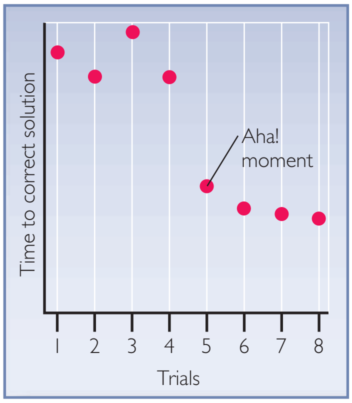

Figure 1.6 'Aha!' reaction. Insight learning: once the animal solves the problem, it gets the answer right almost every time after that.

Skinner found Thorndike's experimental set-up unwieldy, because the researcher had to stick around to place the unhappy cat back into the puzzle box following each trial. This limitation made it difficult to study the build-up of associations in ongoing operant behaviour over hours, days or weeks. So he developed what came to be known as a **Skinner box** (more formally, an operant chamber), which electronically records an animal's responses and prints out a *cumulative record*, or graph, of the animal's activity. A Skinner box typically contains a bar that delivers food when pressed, a food dispenser and often a light that signals when reward is forthcoming (see Figure 1.7). With this set-up, Skinner studied the operant behaviour of rats, pigeons and other animals, and mapped out their responses to reward. By allowing a device to record behaviour without any direct human observation, Skinner ran the risk of missing some important behaviours that the box was not designed to record. Nonetheless, his discoveries forever altered the landscape of psychology.

Terminology of operant conditioning

To understand Skinner's research, you need to learn some psychological jargon. There are three key concepts in Skinnerian psychology: reinforcement, punishment and discriminant stimulus.

Skinner box
Small animal chamber constructed by b. F. Skinner to allow sustained periods of conditioning to be administered and behaviours to be recorded unsupervised.

Reinforcement
Outcome or consequence of a behaviour that strengthens the probability of the behaviour.

Positive reinforcement
Positive outcome or consequence of a behaviour that strengthens the probability of the behaviour.

Negative reinforcement
Removal of a negative outcome or consequence of a behaviour that strengthens the probability of the behaviour.

Reinforcement

Up to this point, we have used the term *reward* to refer to any pleasant consequence that makes a behaviour more likely to occur. But Skinner found this term imprecise. He preferred the term **reinforcement**, meaning any outcome that strengthens the probability of a response (Skinner, 1953; 1971).

Skinner distinguished **positive reinforcement**, when we administer something pleasant, from **negative reinforcement**, when we take away something unpleasant. Positive reinforcement could be giving a child a sweet biscuit when he picks up his toys; negative reinforcement could be ending a child's time-out for bad behaviour once she has stopped whining. In both cases, the most frequent outcome is an increase or strengthening of the response. Note, though, that Skinner would call these actions 'reinforcements' *only* if they make the response more likely to occur in the future.

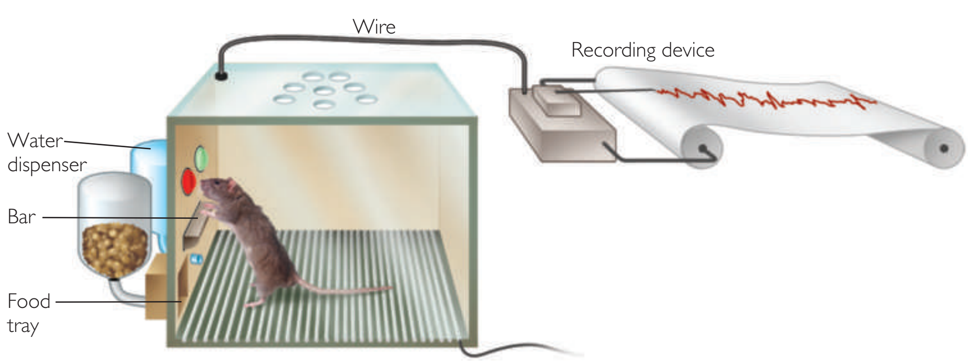

Figure 1.7 Rat in Skinner box and electronic device for recording the rat's behaviour.

B. F. Skinner devised a small chamber (the Skinner box) containing a bar that the rat presses to obtain food, a food dispenser, and often a light that signals when reward is forthcoming. An electronic device graphs the rat's responses in the researcher's absence.

Hundreds of psychology students over the years have demonstrated the power of reinforcement using an unconventional participant: their lecturer. In the game Condition Your Professor (Vyse, 1997), a class of introductory psychology students agrees to provide positive reinforcement—like smiling or nodding their heads—to their lecturer whenever he or she moves in a particular direction, such as to the far left side of the room. One famous introductory psychology teacher spent almost all of his time lecturing from behind a podium. During one class, his students smiled profusely and nodded their heads whenever he ventured out from behind the podium. Sure enough, by the end of class the lecturer was spending most of his time away from the podium. You and your classmates might want to attempt a similar stunt with your introductory psychology lecturer: just don't mention we suggested it.

Punishment

Negative reinforcement should not be confused with **punishment**, which is any *outcome* that weakens the probability of a response. Like reinforcements, punishment can be either positive or negative. If a punishment involves administering a stimulus, then it is positive; if it is taking away a stimulus, then it is negative.

Punishment
Outcome or consequence of a behaviour that weakens the probability of the behaviour.

Positive punishment typically involves administering a stimulus that the organism wants to avoid, such as a physical shock or a smack, or an unpleasant social outcome, such as laughing at someone. Negative punishment involves the removal of a stimulus that the organism wishes to experience, such as a favourite toy or treat.

Punishment also should not be confused with the disciplinary practices often associated with it. Skinner, who insisted on precision in language, argued that certain actions that might superficially appear to be punishments are actually reinforcements. He defined reinforcements and punishments solely in terms of their consequences. Consider this scenario. A mother rushes into her three-year-old son's bedroom and yells 'Stop that!' each time she hears him kicking the wall. Is she punishing the child's demanding behaviour? There is no way to tell without knowing the consequences. If he kicks the wall more often following the scolding, then the mother is actually reinforcing his behaviour—strengthening the probability of a response. If his kicking decreases or stops altogether after he was scolded, then the mother's scolding is a punishment—weakening the probability of a response.

Try labelling each of the following examples as an instance of either negative reinforcement or punishment and explain why. (You can find the answers written upside-down in the margin at the bottom of the next page.)

1. A boy keeps making noise in the back of the classroom despite the teacher's repeated warnings. The teacher finally sends him to the principal's office. When he returns two hours later, he is much quieter.

Table 1.3 Distinguishing reinforcement from punishment

	Procedure	Effect on behaviour	Typical example
Positive reinforcement	Presenting a desirable stimulus	Increases target behaviour	Giving a gold star on homework, resulting in a student studying more
Negative reinforcement	Removing an undesirable stimulus	Increases target behaviour	Static on phone that subsides when you move to a different spot in your room, causing you to stand there more often
Positive punishment	Presenting a stimulus	Decreases target behaviour	Scolding by a pet owner, reducing a dog's habit of chewing on shoes
Negative punishment	Removing a stimulus	Decreases target behaviour	Taking away a favourite toy, stopping a child from throwing future tantrums

Forcing a student to see the principal is typically a form of punishment; nevertheless, it can instead serve as a negative reinforcement if it allows the student to escape from an unpleasant class.
(**Source:** Monkey Business Images/Dreamstime.)

In some countries, such as China and Thailand, children are rarely punished.
(**Source:** Squall/Fotolia.)

2. A woman with diabetes works hard to control her blood sugar through diet and exercise. As a result, her doctor allows her to discontinue administering her unpleasant daily insulin shots, which increases her attempts to eat healthily and exercise.
3. A parole board releases a previously aggressive criminal from prison early for being a 'model citizen' within the institution over the past five years. On his release, he continues to behave in a law-abiding manner.
4. A woman yells at her roommate for leaving dirty clothing scattered all around her apartment. Her roommate apologises and never makes a mess again.

Does punishment work in the long run? Popular wisdom tells us that it usually does: 'Spare the rod, spoil the child.' Yet Skinner (1953) and most of his followers argued against the routine use of punishment to change behaviour. They believed that reinforcement alone could shape most human behaviours for the better.

According to Skinner and others (Azrin & Holz, 1966, pp. 380–447), punishment has several disadvantages:

1. Punishment tells the organism only what *not* to do, not *what* to do. A child who is punished for throwing a tantrum will not learn how to deal with frustration more constructively.
2. Punishment often creates anxiety, which in turn interferes with future learning.
3. Punishment may encourage subversive behaviour, prompting people to become sneakier about the situations in which they can and cannot display forbidden behaviour. A child who is punished for grabbing his brother's toys may learn to grab his brother's toys only when his parents are not looking.
4. Punishment from parents may provide a model for children's aggressive behaviour (Straus, Sugarman & Giles-Sims, 1997). A child whose parents slap him when he misbehaves may 'get the message' that slapping is acceptable.

Numerous researchers have reported that the use of physical punishment by parents is positively correlated with aggressive behaviour in children (Fang & Corso, 2007; Gershoff, 2002). Across many studies, Murray Straus and his colleagues (1998) found that physical punishment is associated with more behavioural problems in children.

Elizabeth Gershoff (2002) conducted a meta-analysis of 88 studies of corporal punishment based on a whopping 39 309 participants. Although she found some evidence that corporal punishment is associated with short-term improvements in children's behaviour, she also found that a history of such punishment in childhood is associated with an increased probability of becoming an abuser in adulthood.

Yet we must remember that these studies are correlational and do not demonstrate causality. Other interpretations are possible. For example, because children share half of their genes with each parent, and because aggression is partly heritable (Krueger, Hicks & McGue, 2001), the correlation between parents' physical aggression and their children's aggression may be due to the fact that parents who are physically aggressive pass on this genetic predisposition to their children (DiLalla & Gottesman, 1991, pp. 125–129). Or the correlation might be driven by environmental factors, such as family exposure to a lot of media violence, which could make the parents more likely to use harsh discipline and independently make their children more aggressive.

A recent study addressed this problem by investigating children of identical twins. Children of 887 twins participating in the Australian Twin Registry provided information about how they were disciplined by their parents; they also were interviewed about their lives, including their drug and alcohol use and their tendencies towards positive and negative behaviours. This experimental design is powerful because it untangles the effects of genes (shared by both twins and passed onto their children) from punishment styles (which sometimes differed from twin to twin). The results indicated higher levels of aggressive and delinquent behaviour and greater alcohol and drug use in children whose parents used harsh physical punishment compared with those whose parents used milder methods. The results ruled out a genetic explanation, because in pairs of twin parents who differed in their disciplinary styles, children of the harsh-punishing twin displayed more of these negative behaviours than children of the mild-punishing twin (Lynch et al., 2006).

That is not to say that we should never use any kind of punishment, only that we should use milder methods and use them sparingly. Most research suggests that punishment works best when it is delivered consistently and follows the undesired behaviour promptly (Brennan & Mednick, 1994). In particular, immediate punishment tends to be effective, whereas delayed punishment is often useless (Church, 1969; McCord, 2006; Moffitt, 1983). Punishment of an undesired behaviour also works best when we simultaneously reinforce a desired behaviour (Azrin & Holz, 1966).

Discriminative stimulus

The final critical term in operant conditioning lingo is **discriminative stimulus (S_d)**, which is any stimulus that signals the presence of reinforcement. When we snap our fingers at a dog in the hopes of having it come over to us, the dog may approach us to get a much-appreciated petting. For the dog, our finger snapping is a discriminative stimulus: it is a signal that if it comes near us, it will receive reinforcement. According to behaviourists, we are responding to discriminative stimuli virtually all the time, even if we are not consciously aware of it. A friend's waving at us from across campus is another common discriminative stimulus: it often signals to us that our friend wants to chat with us, thereby reinforcing us for responding to her wave.

Discriminative stimulus (S_d) Stimulus associated with the presence of reinforcement.

Acquisition, extinction, spontaneous recovery, stimulus discrimination and stimulus generalisation

If you feel you are experiencing a case of déjà vu upon reading these bolded terms, do not be concerned, because you have indeed seen all of them before. Acquisition, extinction, spontaneous recovery, stimulus discrimination and stimulus generalisation apply just as much to operant conditioning as to classical conditioning. You can find the definitions in Table 1.4. Next, we examine how three of these concepts apply to operant conditioning.

Extinction

In operant conditioning, extinction occurs when we stop delivering reinforcement to a previously reinforced behaviour. Gradually, this behaviour declines in frequency and disappears. If parents give a screaming child a toy to quieten her, they may be inadvertently reinforcing her behaviour, because she is learning to scream to get something. If her parents