Chapter 3

The Behaviourist Approach

Chapter contents

LEARNING OBJECTIVES

In this chapter, the objectives are to learn

- the basic assumptions of behaviourism
- the nature of stimuli and responses
- the principles of classical conditioning, including
  - unconditioned stimulus and response
  - conditioned stimulus and response
- the phenomena of classical conditioning, including
  - stimulus generalization and discrimination
  - extinction and spontaneous recovery
  - higher order conditioning
- the applications of classical conditioning, including
  - conditioned emotional responses
  - conditioned drug and immune responses
- the principles of operant conditioning, including
  - reinforcers and reinforcement
  - contingencies of reinforcement
  - schedules of reinforcement
- the phenomena of operant conditioning, including
  - shaping
  - extinction
  - discriminative stimuli
  - non-contingent reinforcement
- the applications and implications of conditioning, including
  - aversive control of behaviour
  - the interrelationships of classical and operant conditioning
  - autonomic conditioning and biofeedback
  - biological constraints on learning

MIND DOESN'T MATTER

One of the basic themes of this book is that behaviour can often be understood in different ways, represented by the five approaches. As we will see in this chapter, the behaviourists emphasize links between the environment and behaviour. In doing so, they tend to ignore both physiological processes and mental events, even in circumstances that might invite such interpretations, as in an example reported by a psychologist named Israel Goldiamond (1973).

As the result of a car accident, Goldiamond spent several months in hospital undergoing treatment for a spinal injury. At one point, he shared a room with a man who had suffered brain damage. This patient was often disoriented, urinating on walls and muttering, 'What the hell am I doing in Panama?' (He wasn't.)

Thus far, this story appears unremarkable – after all, Chapter 2 explores how the brain controls behaviour. Yet, rather than focusing on the brain damage, Goldiamond examined the influence of the environment on the man's behaviour. He noted, for example, that the man did not act oddly in the hospital cafeteria – a fact which the staff had overlooked. Goldiamond accounted for this by noting that the features of cafeterias are fairly universal, while those of a rehabilitation hospital are not familiar to most people. Hence, the man was only disoriented when he was in an unfamiliar setting. Furthermore, the hospital was located on a large lake, and the patient's room overlooked a naval pier – perhaps accounting for his questions about Panama. Goldiamond suggested that the man's urinating inappropriately could be dealt with by rewarding him with cigarettes for urinating properly – in effect, controlling the behaviour by means of an external incentive. As anticipated, the technique worked.

The point of this story is to show that looking at the brain is not always the best way to understand behaviour. Goldiamond, like other behaviourists, preferred to look at the role that the environment plays in behaviour. There is no question that the man in the story had suffered brain damage. But it is equally clear that his behaviour could not be fully understood by looking *only* at the brain damage. In one sense, by placing such a heavy emphasis on *internal* events, the biological approach tends to give too little attention to the *external* context of behaviour – that is, the environment in which behaviour occurs. In this chapter, we will consider the role of environmental influences on behaviour, as seen from the behaviourist perspective.

INTRODUCTION

The *behaviourist approach* emphasizes the role of environmental stimuli in determining the way we act. In large measure, this means focusing on **learning** – changes in behaviour which occur as the result of experience. (By emphasizing experience, behaviourists exclude changes due to fatigue, injury or drug effects.) Behaviourism has added considerably to our understanding of learning through the study of what is called *classical* and *operant conditioning*. Before examining what has been discovered, let us look at the basic assumptions and methods which distinguish behaviourism from the other approaches.

learning a change in behaviour which occurs as the result of experience.

As with all of the approaches, the choice of focus is one of the factors which gives behaviourism its uniqueness. In this case, the behaviourist approach is commonly distinguished by its emphasis on the relationship between observable behaviour (*responses*) and environmental events (*stimuli*). Consider the simple interaction involved when a child reaches out towards a glowing fire, and then quickly draws back from the heat: first, the stimulus of the sparkling flame attracts their attention, so that they move their hand forward (stimulus of fire leads to response of reaching). Then, the heat of the fire leads to a reflexive withdrawal (stimulus of heat leads to withdrawal response). This, in turn, might lead the child to throw water on the fire, or take some other action. Thus, from the behaviourist perspective, human experience can be understood in terms of the interrelations between stimuli and responses.

Basic Assumptions of Behaviourism

Like other approaches, the behaviourist approach is defined not only by the kinds of data it emphasizes, but also in terms of its basic assumptions, which are closely related to its historical origins. At the turn of the twentieth century, psychologists tended to focus on either the experimental study of physiological processes, or the introspective analysis of experience (see Chapter 1 for a review). Physiological research was hampered by the limited technology available for studying the brain (for example, not even X-rays or EEGs existed), and introspectionism was proving limited due to problems of subjectivity in describing sensory experience. Consequently, both had serious limitations. As an alternative, William James argued that psychologists should focus on how behaviour relates to its purpose (called *functionalism*), but he was often better at framing the issues than at doing research to solve them. Thus, none of the available methods was achieving unequivocal success. It was against this backdrop that behaviourism arose. (It should be noted that 'behaviourism', like other approaches, can refer to a number of theories, each with some unique aspects. Nonetheless, it is possible to identify some common elements within the approach.)

While the temptation in discussing the behaviourist approach is to emphasize the type of data collected (the observable behaviours which give the approach its name), doing so ignores the broader assumptions which underlie the approach.

parsimony in the philosophy of science, the principle that one should always seek the simplest possible explanation for any event.

The most fundamental of the basic principles in behaviorism is the concept of parsimony. Sometimes called 'Occam's razor' after the English philosopher who first proposed it, **parsimony** favours seeking the simplest possible explanation for any event. If, for example, one can explain a person's eating a pastry without referring to a non-observable concept like 'hunger' or 'oral personality', then parsimony says one should avoid using such concepts. Behaviourists reacted against introspectionism in part because it seemed to invoke too many vague concepts, and thereby lacked parsimony. Instead, behaviourism focused on the use of *operational definitions* (defining concepts in terms of observable events) – and this led naturally to the focus on 'stimuli' and 'responses' (for further discussion of operational definitions, refer back to Chapter 1).

The second basic assumption of the behaviourist approach relates to the basis of behavioural change. Like functionalism, behaviourism tries to understand the conditions under which behaviour occurs. When does a particular behaviour occur? What conditions lead to it? What changes in the environment result from it? Since relatively few behaviours in human beings are genetically programmed, this leads to a focus on the role of experience, which is expressed through learning. It is easy to say that the way we act depends on our past experiences, but just *how* does learning occur? Since the time of Aristotle, the basic explanation has been that we learn by *association* – that is, by forming connections between ideas and/or events. For example, if the sound of an electric can opener leads a dog to run to the kitchen, we can speculate that the dog has formed an association between the sound of the can opener and being fed (canned food!). This concept of

associationism, which was also favoured by such English philosophers as David Hume and J. S. Mill, has had a fundamental influence on psychology, particularly for the behaviourists. As we shall see, behaviourist theories are essentially theories of how associations are formed.

associationism the doctrine, supported by Aristotle, Hume and others, that mental processes, particularly learning, are based on forming connections between ideas and/or events.

Taken together, parsimony and associationism formed the foundation from which behaviourism arose. Exactly how, and what the result has been for our understanding of psychology, will form the substance of this chapter.

The Pioneers of Behaviourism

Just over 100 years ago, an American named Edwin L. Thorndike was studying for his PhD in the newly-formed psychology department at Columbia University (he started at Harvard, under William James, but transferred for financial reasons). For his research, he studied problem solving in animals, using a series of puzzle-like tasks (such as confining a cat in a box, from which it could release itself by pressing against a lever). His dissertation, published in 1898, had the rather cognitive-sounding title of *Animal Intelligence* (Thorndike 1898).

Despite the title, Thorndike's research was basically concerned with trying to analyse the conditions under which animals changed their behaviour – that is, *learned*. In doing so, he focused on the relationship between a response and its consequences, resulting in what he called 'the law of effect'. Basically, the **law of effect** said that any response which leads to an outcome that the organism finds satisfying is likely to be repeated, and any response which leads to an unpleasant outcome is not likely to be repeated. This was a form of associationism, in that the organism (animal, person) was seen as making a connection between a response and its consequences. (This is technically called association by *contiguity*, in that it assumes the response and consequence must be closely linked in time and space.) While basically unoriginal – the idea that individuals respond to reward and punishment extends back to the ancient Greeks – Thorndike's version could be said to differ in that *it was supported by experimental data* (Robinson 1979). By framing the issue in experimental (and therefore scientific) terms, Thorndike paved the way for the behaviourist approach.

law of effect a principle of learning developed by Edwin Thorndike, stating that any response which leads to a satisfying outcome for the organism is likely to be repeated, and any response which leads to an unpleasant outcome is not likely to be repeated.

Key Thinker: **Edwin Lynn Thorndike**

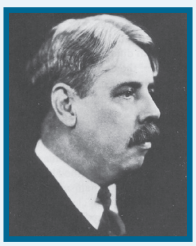

Edwin Lynn Thorndike (1874–1949) was born in Williamsburg, a small town in western Massachusetts. After receiving his bachelor's degree from Wesleyan University, he went to Harvard University to study psychology under William James, but was forced to transfer to Columbia University because of financial difficulties. In the newly-formed psychology department at Columbia, he studied under James McKeen Cattell, one of the most influential early American psychologists. For his research, he studied problem solving in animals, using a series of puzzle-like tasks (e.g., confining a cat in a box, from which it could release itself by pressing against a lever). His dissertation, published in 1898, had the rather cognitive-sounding title, *Animal Intelligence*. Thorndike is probably best known today for his 'law of effect', which foreshadowed Skinner's concept of reinforcement as a description of the role of consequences in learning. From 1899 he taught at Teachers College at Columbia, where he wrote prolifically on education as well as psychology. He died in New York at the age of 75.

Thorndike's law of effect, while significant, was not without problems; a key difficulty was that it was vague about what made something 'satisfying'. One way to resolve this might be to seek physiologically-oriented mechanisms for 'satisfaction'. However, given turn of the century knowledge of physiology, this often required resorting to non-observable concepts, which violated the principle of *parsimony*. A more radical approach, pioneered by John B. Watson, was to pare theorizing to the bone, restricting theoretical descriptions to factors which could be directly observed and measured.

John B. Watson was both gifted and provocative. As a student at the University of Chicago, he initially trained in introspectionism, but found its approach to psychology excessively vague, especially in its emphasis on mental processes. He began working with animals, and completed his PhD in three years – at that time being the youngest such graduate from the university. After teaching for only four years at Chicago, he was offered a full professorship at Johns Hopkins, and shortly after became chair of the psychology department there – an example of remarkable career advancement!

Watson can only be described as zealous in promoting his ideas. Reading his major work (*Behaviorism* 1930) today, one is struck by the scorn he heaps on William James and others, and by his willingness to test his ideas whenever possible (even using his own children). Confident that he was correct, he was willing to extend his claims even when he lacked experimental support, as in his famous remark, '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.' What is often omitted in this quotation is the statement which follows it: *'I am going beyond my facts and I admit it, but so have the advocates of the contrary* and they have been doing it for thousands of years' (Watson 1930, p. 104; emphasis added). Clearly, Watson did not shirk from confrontation in pursuing his ideas.

Key Thinker: **John Broadus Watson**

John Broadus Watson (1878–1958) was the founder of behaviourism. Educated in a one-room schoolhouse in the American farm belt (like many of his era), he went on to complete his PhD at the University of Chicago. After a brief exploration of the introspectionist approach, he continued under John Dewey and James Angell, two of the pioneers of the functionalist approach. Watson was both gifted and outspoken – two characteristics which played a major role in his career. He completed his PhD in only three years and began teaching at the university; four years later, he was offered a full professorship in psychology at Johns Hopkins University, and shortly after became chairman of the department. In 1913 he began publishing the first of a series of publications which outlined his behaviourist approach, which quickly gained him both fame and notoriety – his statement about shaping a healthy infant in any way desired (quoted in the text) is characteristic of his assertive style. While at Johns Hopkins, he met graduate student Rosalie Rayner, who became his second wife. After collaborating with Rayner on the case of 'little Albert', Watson became interested in human sexual behaviour; his activities in this regard (including participant observation) did not sit well with the prevailing moral views, and he was finally dismissed. At this point, he took a job with the J. Walter Thompson advertising firm. Not surprisingly, he did well in his new role, embarking on studies of consumer behaviour, writing psychology articles for the general public, and becoming a vice-president of the advertising firm in fewer than four years. Thus, Watson not only founded an entire approach to psychology, but perhaps was also the first psychologist to apply psychological theory to advertising and marketing! He died in New York at the age of 80.

Watson's writings and ideas were a lever that moved the world. In the following decades, until the mid-1950s, behaviourism became the dominant force in psychology, particularly in North America. The irony is that while the general approach became highly influential, most researchers never accepted Watson's extreme position, which is sometimes called *radical behaviourism*. Even B. F. Skinner, the spiritual heir to Watson's work, has commented, 'A shortage of facts is always a problem in a new science, but in Watson's aggressive program in a field as vast as human behaviour it was especially damaging' (Skinner 1974, p. 6). The greatest impact of Watson's ideas can be traced to three central elements:

1. the emphasis on observable responses and environmental stimuli;
2. the rejection of mentalistic concepts not grounded in direct observation; and
3. the focus on learning and experience as central to the understanding of human behaviour.

Some 70 years of research, both basic and applied, has demonstrated that in many ways, we are indeed what we learn.

Behaviourism, as already noted, has many variants; indeed, some would say it is more appropriate to speak of the behaviouristic approach than the behaviourist approach. Even E. C. Tolman, often regarded as one of the founders of the *cognitive* approach, considered himself a 'behaviourist' (Tolman 1932). However, what all behaviourists share is an interest in how behaviour is learned, and an emphasis on explanations based on observable events. In this chapter, we will see how this approach has been applied to a variety of situations. (Note that today, research which focuses on the relationship between environmental factors and behaviour is sometimes described as 'behavioural' – but this is a more general term than *behaviourism*.)

Try it Yourself

In everyday life, we all try to make sense of the behaviours we see in ourselves and others – but often we violate behaviourist principles by going beyond what is observable. For example, consider this scenario: you good-naturedly tease a friend about forgetting her wallet at home, and she snaps at you to mind your own business. You may think, 'She's feeling upset – she must have had a fight with her boyfriend.' This explanation, however, violates the basic assumptions of behaviourism: it explains your friend's behaviour in terms of something you can't see (her being upset with her boyfriend), while it neglects the role of observable events (her comment immediately followed your teasing). If you look at your friend's behaviour the way Watson recommended, you might conclude that her remark was in fact a reaction to your teasing. With the first explanation, you might disregard your own behaviour and instead focus on her (presumed) anger at her boyfriend. With the more behaviourist explanation, you might conclude that teasing a friend isn't always a good idea. The point here is that the behaviourist approach leads you to focus on observable aspects of the situation, and that can change your interpretation. Look at the following situations: are the interpretations you make using behaviourist principles the same or different from what you would normally conclude?

- A toddler hits another child in a school playground.
- A driver 'tailgates' your vehicle while driving on a highway.
- A classmate you encounter in the library offers to buy you a coffee.

Stimuli and Responses

Behaviourism, by focusing on *observable* events, sets its own limits on what can be studied. Thoughts, feelings and other inner mental states cannot be studied empirically, and so have no place in behaviourist theory. Genetic variation, while presumably contributing to differences among individuals, is also ignored, because traditionally it was not measurable (and is still largely inaccessible). By contrast, environmental conditions are relatively easy to measure and study.

Taken as a whole, the environment involves colours, shapes, smells, sounds and many other characteristics. Obviously, it is impossible in an everyday setting to measure every element of a typical environment. However, in most cases, this would be unnecessary, because there are many environmental elements that typically do not seem to enter our awareness, and consequently have little impact. (Recall the discussion of perceptual processes in Chapter 1.) Nonetheless, behaviourists recognize that in order to study environmental influences on behaviour, one must be able to rigorously define the environmental characteristics involved in a situation. In practice, this means that research often involves limiting the complexity of the environment, particularly in laboratory studies. It also means that one must be able to define terms clearly. With regard to the environment, sights, sounds and smells are all considered examples of stimuli. A **stimulus** (often abbreviated as **S**) is any event, situation, object or factor that is measurable and which may affect behaviour. Simple examples could include a red triangle, the ticking of a watch or a pinprick.

stimulus in general, any event, situation, object or factor that may affect behaviour; for the behaviourists, a measurable change in the environment.

For a behaviourist, an important element in understanding a particular behaviour is to identify the stimulus (or stimuli) involved. From the examples above, this would seem to be a fairly straightforward task. In reality, however, it can sometimes be quite difficult to define which environmental elements are involved as stimuli in a specific situation. For example, a mother approaches her 2-month-old infant, and the baby smiles. The mother seems to be the stimulus which elicits the baby's smiling. But is it the mother as a whole, or her face, or her expression, or her smell, or her touch, or some combination of these and other elements to which the baby is actually responding? In research, it would be necessary to identify the actual stimulus elements in order to understand the situation properly. (In this regard, the desire for *operational definitions* of terms, including 'stimulus', becomes understandable as a means of avoiding ambiguity.)

Similarly, it is necessary to describe clearly the behaviour being studied. Normally, the behaviour which is measured is called the **response** (often abbreviated as **R**). Again, this may seem very simple at first glance. For example, a person sits at a table, eating. 'Eating' is obviously a response; however, a moment's thought will show that there can be tremendous variations in the behaviour described as eating. A finicky child may pick reluctantly at a disliked vegetable. A hungry person may ravenously devour a favourite dish. While both are eating, there is clearly a large difference in their behaviour. Consequently, researchers must be careful to describe a response in terms that are meaningful to the situation. Often this will require specifying the rate, intensity and/or other characteristics of the response.

response in general, any reaction to a stimulus, whether overt or mental; for the behaviourists, a measurable change in behaviour.

One of the distinctions among responses that became evident to the early behaviourists was a distinction between reflexes and voluntary actions. **Reflexes** are unlearned responses that can be triggered by specific environmental stimuli. Examples of human reflexes include withdrawing the hand from a hot surface, or a baby's sucking on an

reflex an unlearned response that can be triggered by specific environmental stimuli, such as a baby's sucking on an object placed in the mouth.

voluntary response a response which is controlled by the individual (i.e., emitted) rather than being triggered (elicited) by specific stimuli the way relfexes are.

object placed in the mouth. By contrast, **voluntary responses** are emitted – that is, they are not triggered by stimuli in the way reflexes are; typically, they involve more complicated actions, which often require extensive practice. Thorndike, for example, in his studies of problem solving by animals, was looking at voluntary responses, and he found it took repeated trials for learning to occur. Voluntary behaviour can span a tremendous range, from simple actions like learning to use a fork, to complex behaviours like speaking a new language. Such differences led early behaviourists to the separate study of reflexes and voluntary responses. As we shall see, they discovered that the principles of learning seem to differ for the two types of behaviour. As we will also consider later in the chapter, the actual differences between the two types may be smaller than they initially appear.

CLASSICAL CONDITIONING

At the turn of the twentieth century, a Russian physiologist named Ivan Pavlov was engaged in a long-term project to understand the process of digestion. Beginning in 1879 and working primarily with dogs, his work earned him the Nobel Prize in 1904 (Windholz 1997). But sometime around 1902, he noticed a phenomenon which was to lead him in a new and unexpected direction. In order to study digestion, Pavlov measured a number of factors, including how much a dog salivated when it was given food. Then one day he noticed a phenomenon he labelled 'psychic salivation' – a dog would salivate before it was actually given food. Since Pavlov believed that digestion involved a series of reflexes, he set out to determine what controlled this anticipatory response. What he discovered became the basis for what is now commonly called **classical conditioning** – the study of learning which involves reflex responses. Essentially, classical conditioning explores how a new stimulus can come to elicit an existing reflex response due to learning.

classical conditioning the study of learning which involves reflex responses, in which a neutral stimulus comes to elicit an existing reflex response.

Key Thinker: Ivan Petrovich Pavlov

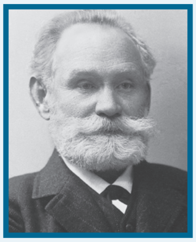

Ivan Petrovich Pavlov (1849 – 1936) was a Russian physiologist who pioneered the study of classical conditioning. Born in Ryazan, Russia, he initially began seminary studies, but then changed to St Petersburg University, where he graduated in natural science, and subsequently received his doctorate in physiology from the Military Medical Academy. After a few years spent in Germany, he went on to become a professor at the Military Medical Academy. His work on the physiology of digestion, begun in 1879, earned him the Nobel Prize in 1904. He first became aware of reflexes by reading Sechenov's work while still at seminary, but his own research on what became known as classical conditioning did not begin until about 1902. At this time, while still studying digestion in dogs, he noticed what he called 'psychic salivation' – a dog would salivate before it was actually given food. Since Pavlov believed that digestion involved a series of reflexes, he set out to determine what controlled this anticipatory response. Ultimately, his work on conditioning overshadowed the research which had earned him the Nobel Prize. He continued to be intellectually active, forming a genetics institute only a few years before his death at the age of 87.

Pavlov's original studies have become so well known as to be the object of jokes (like the psychologist who salivates when he hears the name 'Pavlov'!). In simple outline, Pavlov found that by ringing a bell and then immediately giving the dog some food, the bell came to evoke the same response as the food itself – salivation. To understand why this is remarkable, we need to consider the elements of the situation more closely. As Pavlov's lengthy studies of digestion showed, salivating at the presence of food is a basic neural reflex that requires no learning. For example, if you put a piece of chocolate in your mouth, you will salivate. A light shined in the eye will cause the pupil to contract. Reflex responses like these (and they exist in species from worms to humans) are referred to in classical conditioning as **unconditioned responses**. For any reflex, there is some stimulus which will trigger (*elicit*) the response (for example, food for salivating, light for pupil contraction); the stimulus which elicits an unconditioned response is called an **unconditioned stimulus**. Since 'conditioned' refers to *learned*, the term refers to the unlearned nature of reflexes. (Pavlov of course wrote in Russian, and actually used the term 'unconditiona*l*', but an early English translator erred, and the mistake has remained.)

unconditioned response a reflexive response produced by a specific stimulus, such as pupil contraction to bright light.

unconditioned stimulus a stimulus which elicits a reflexive (unconditioned) response.

If reflexes are unlearned, then what is the learning that occurs in classical conditioning? Pavlov noted that the learning is based on forming a connection between stimuli – in the dog's case, between the bell and the food. Ringing the bell initially had no effect on salivation – that is, with respect to the response of salivation, it was a **neutral stimulus**. (To be a stimulus, an environmental element must be something which the organism is aware of; normally this is demonstrated by the stimulus arousing attention, called an *orienting response*.) After repeated pairings with the food placed in the dog's mouth, the sound of the bell came to elicit drooling. At this point, the sound has become a **conditioned stimulus**, and the salivating which results is called a **conditioned response** (to distinguish it from the response to food alone, see Figure 3.1). Essentially, the conditioned stimulus has become *associated* with the occurrence of food (Pavlov 1927).

neutral stimulus a stimulus which initially produces no specific response other than provoking attention; as conditioning proceeds, the neutral stimulus becomes a conditioned stimulus.

conditioned stimulus a stimulus which by repeated pairings with an unconditioned stimulus comes to elicit a conditioned response.

conditioned response a response to a previously neutral stimulus which has become a conditioned stimulus by repeated pairing with an unconditioned stimulus.

In order to appreciate the significance of classical conditioning, we must examine its characteristics more closely. Since the response involved is essentially a pre-existing reflex, the learning which occurs does *not* involve a new response; instead, it consists of forming a connection (*association*) between two stimuli (the CS and UCS). In order for optimal conditioning to occur, the conditioned stimulus (CS) must occur a second or so before the unconditioned stimulus (UCS). If the two occur simultaneously, conditioning may occur, but is typically weaker. If the CS is presented *after* the UCS (sometimes called *backward conditioning*), then *no* learning occurs. What this tells us is that conditioning is closely linked to the ability of the CS to serve as a signal that the UCS is going to occur. This is further demonstrated by studies which show that conditioning is only likely when the CS *reliably* predicts the occurrence of the UCS (Rescorla 2000). In some sense, what makes classical conditioning a valuable process for the organism is the fact that *it allows one to anticipate environmental events*. This notion that classical conditioning helps in adapting to the environment is supported by research on a phenomenon called *blocking*. If a new stimulus is presented simultaneously with an existing CS, conditioning to the new stimulus does not occur, because the original CS is *already* an adequate signal (Kamin 1969). Flashing a light to signal food is unnecessary, if a bell already serves that purpose. Conditioning, then, seems to

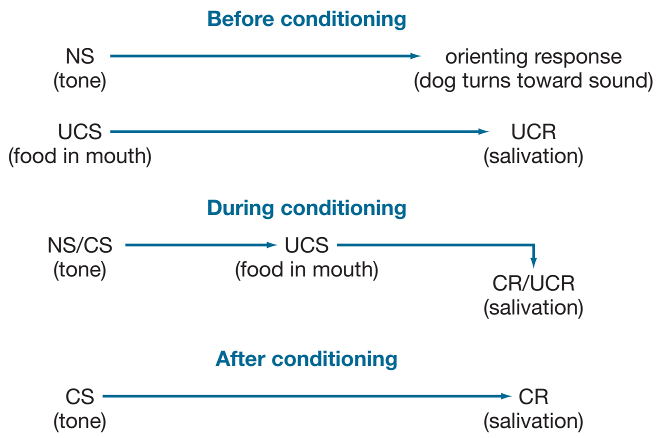

Figure 3.1 The Basic Classical Conditioning Procedure In classical conditioning, repeated pairing of a neutral stimulus (NS) with an unconditioned stimulus (UCS) results in conditioning, whereby the NS becomes conditioned (CS) to elicit the same type of response as the UCS did.

occur because such learning is often adaptive, enabling individuals to deal with the world more effectively.

Classical conditioning has been demonstrated in a wide variety of species, from worms to birds to primates. Thus, it appears to be a very fundamental form of learning. But given that much of human behaviour does not depend on reflex responses, it might seem that classical conditioning is of little significance in people. In fact, the reality seems to be just the opposite: examples of classical conditioning seem pervasive in our lives. For example, we respond to stimuli associated with food – smells, pictures in advertisements, words like 'chocolate cake' – in much the way that Pavlov's dogs reacted to the bell: by increased digestive activity. In these cases, the food cues, having been associated with food itself (the UCS), are conditioned stimuli. Such food cues are among the most reliable of conditioned stimuli, because the sight and smell of food always precedes the actual eating of it. In movies, directors will use sounds to enhance the emotional content of the story. For example, a particular theme may precede the repeated appearances of the villain in a horror movie; viewers then come to associate the theme to the moments of mayhem that follow. (Of course, film images of violence are *themselves* conditioned stimuli, associated with past experiences of actual injury (UCS). This relates to the process of *higher order conditioning*, which we will discuss later in the chapter.) Many people, when showering, develop a conditioned response of anxiety to the sound of a toilet flushing (CS), since

Try it Yourself

Pick up any magazine and look at the advertisements. You will probably see that in each advertisement, the product is displayed along with one or more attractive models. Given the basic principles of classical conditioning, why might this be so? Do you find the advertisements using attractive models to be more compelling than those that do not? Is that because you are focusing on the product or on the model? What about political advertising: although 'sex appeal' may not be used in the same way, do political ads seem to make use of classical conditioning principles? How?

it often results in a sudden increase in the temperature of the water (UCS)! Classical conditioning is thus a flexible process which allows us to anticipate biologically significant events (UCSs) by making an association to stimuli (CSs) which precede them.

Classical Conditioning Phenomena

Stimulus Generalization and Discrimination

Having established the basic elements of classical conditioning, Pavlov (and later, others) began to explore some variations of the original situation. One subject that interested him was the element of stimulus novelty: what would happen if a new stimulus was presented as a CS? Tests with unrelated stimuli quickly established that a neutral stimulus will not elicit a response which has been conditioned to a different stimulus (for example, flashing a light will not elicit a CR if the previous CS was the sound of a bell). However, what would happen if a stimulus similar to the CS were used (such as a different bell)? Tests of this type revealed a new phenomenon, called **stimulus generalization**: stimuli similar to the original CS would tend to elicit the same CR. Research has shown that the degree of response is related to the degree of similarity between the new stimulus and the original CS (see Figure 3.2).

stimulus generalization in classical conditioning, the tendency to produce a CR to both the original CS and to stimuli which are similar to it in some way.

This may not seem like a very surprising result, but it is very profound in its implications. In everyday life, we seldom encounter the same precise situation twice. For example, the traffic pattern on a road is never identical on two occasions, requiring us to pay attention each time we drive the same route. Even people change, as they wear different clothes, change their hair style, etc. Given this reality, it is generally desirable to be able to ignore these minor variations – in other words, to generalize across basically similar stimuli. This is precisely what the studies

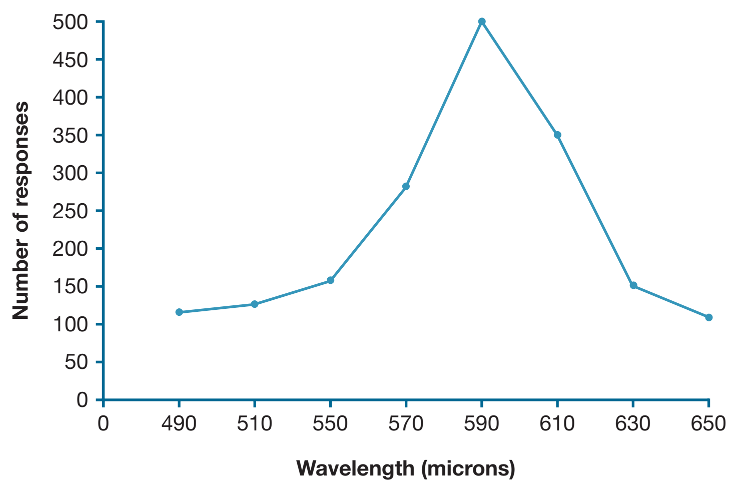

Figure 3.2 Stimulus Generalization Typically, an organism trained to respond to one stimulus will respond in the same way to stimuli which are perceived as similar. In the example, varying the wavelength (colour) from the original greenish stimulus leads to responding which decreases as the test stimuli become less similar. (These data, averaged from Peterson 1962, actually refer to operant conditioning of pigeons, but the basic phenomenon is the same in classical conditioning.)

of stimulus generalization in classical conditioning demonstrate. In practical terms, stimulus generalization results in responding to *a whole class of related stimuli*, after initial learning with a single stimulus. A child who has learned to withdraw after touching a glowing burner on a gas stove will tend to generalize this withdrawal to other stimuli that are similarly bright and hot – electric burners, open fires, etc. In this sense, stimulus generalization can enable organisms to adapt better to their environment – though it may not always be adaptive. For example, stimulus generalization has been noted in how people respond to brand names for products (Till and Priluck 2000). After using conditioning to establish favourable attitudes to imaginary brands, the researchers found that ratings carried over to products with the same brand in a different category. Hence, companies can 'build on the brand name' to market new products!

How, exactly, can we define or measure 'similarity'? Ideally, we should have some general procedure, operationally defined, to measure similarity for any stimuli. Considerable attention has been given to this problem, but as yet there is no universal standard to determine similarity. Lacking a clear general definition, one must resort to defining similarity by observing the outcome of experimental tests. Thus, if two stimuli elicit essentially identical results, they are highly similar; if CS_1 produces a strong conditioned response, but CS_2 elicits only a weak response, then they are not very similar.

The fact that there is no reliable way to predefine similarity may seem a serious weakness, but in fact it may actually tell us something about the nature of stimulus generalization. The typical experiment produces results like those in Figure 3.2, where response intensity drops off as the difference between stimuli increases. However, this is not the only possible outcome. Depending on experience during and prior to training, results can vary significantly. For example, when two stimuli are randomly mixed as the CS during training, generalization is basically equal to both stimuli (Grice and Hunter 1964). Even more interesting are the implications of experiments on stimulus generalization where the environment of the animals has been carefully controlled

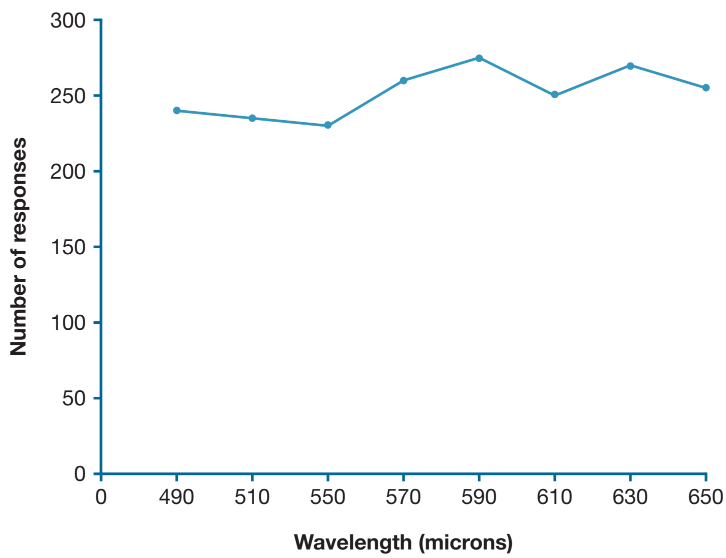

Figure 3.3 Pure Generalization When pigeons were raised under monochromatic light, testing for stimulus generalization with other colours of light yielded essentially no differences in responding – in effect, the animals generalized to any colour of light. (Data averaged from Peterson 1962)

prior to conditioning. For example, if pigeons are reared from birth with only yellow light for illumination, and then a coloured light is used as the CS, the pigeons will respond *equally* to *any* colour of light (Peterson 1962) (see Figure 3.3). In this situation, it seems that the absence of prior experience with colour as a stimulus characteristic leads to regarding all colours as similar. Babies show a somewhat comparable response, in that initially they smile at anyone who smiles at them, whether parent or stranger. Thus, it seems that the initial tendency of an organism is to generalize when encountering a new situation. (As a perceptual characteristic, this may also relate to such behaviours as *stereotyping*, which involve generalizing based on group membership.)

Why, then, do most experiments show the gradients seen in the first example? The answer seems to lie with another phenomenon which Pavlov studied. He noted that a dog conditioned to salivate to the presentation of a black square (CS) also salivated at the sight of a grey square – an example of stimulus generalization. Pavlov then ran a series of trials during which the black square was always followed by food (UCS), but the grey square was never followed by food. After a number of such trials, the dog reliably salivated to the black square, but no longer did when presented with the grey square. This was a demonstration of **stimulus discrimination**, whereby the organism is conditioned to *distinguish* between two stimuli (see Figure 3.4). Pavlov subsequently demonstrated that such discriminations can be remarkably precise, if training is continued with stimuli which become progressively more alike. What is notable is that *stimulus discrimination always requires training* – in the *absence* of such training, organisms tend to *generalize*.

stimulus discrimination selective responding to the CS, but not to stimuli which are similar in some way as a result of training.

If we then reconsider the puzzle of generalization gradients, what seems to emerge is the implication that 'typical' gradients reflect a *combination* of generalization and discrimination. In the everyday world, organisms learn that stimulus variations sometimes are significant, and sometimes are not. Pigeons, for example, may use their colour vision to determine when berries

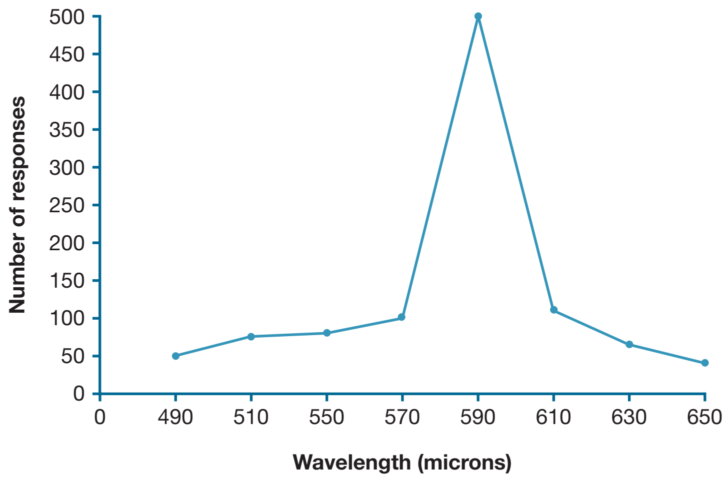

Figure 3.4 Stimulus Discrimination When pigeons are trained in discrimination, the generalization gradient becomes much steeper (compare to Figure 3.2). Taken with the two previous graphs, this suggests that a 'typical' generalization response actually reflects a degree of discrimination, unlike the pure generalization produced when there has been no prior experience of stimulus variation (Figure 3.3).

are ripe or not, but not to distinguish between berries and seeds. Typically, the smaller the colour difference, the less significant it would be. Only when they have *no prior experience* with colour variations (as in the unusual experiment described above) will they totally ignore colour differences. According to this analysis, trying to determine a universal standard of similarity is a hopeless task, since organisms will show varying response patterns depending on their past experience. In other words, learning based on classical conditioning is a cumulative process, with present behaviour being influenced by prior conditioning experiences.

While *what* we discriminate may depend on past experience, the *capacity* to discriminate seems to be inborn, and can often be crucial to adaptation. For instance, detecting the difference between food which is safe or spoiled often depends on discriminating particular odour cues. Some occupations are closed to individuals with colour blindness, because job performance requires discrimination based on colour (for example, certain types of electrical work, where wires are colour coded). Thus, in some circumstances, lacking the capacity to discriminate could seriously reduce our capacity to adapt, or even survive.

Extinction and Spontaneous Recovery

We have seen how classical conditioning, combined with stimulus generalization and discrimination, can lead to learning which is often highly adaptive. However, we have not said anything about how long the effects of conditioning last, or what happens if a conditioned response is *not* adaptive. A Russian researcher, W. H. Gantt, has commented on the possibility that conditioned responses, being persistent once formed, can turn an individual into 'a museum of antiquities' as time goes on. Many reactions would be based on particular past situations, and consequently might be either no longer useful, or even detrimental (Gantt 1966). For example, a person who broke an arm might continue to favour it (because of the pain associated with the original injury) long after healing had been completed. This would represent a form of persistent classical condi-

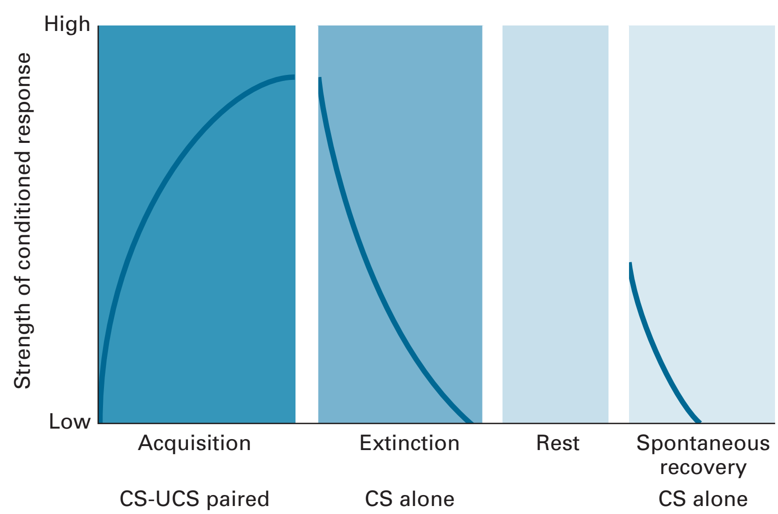

Figure 3.5 Changes in the Strength of a Conditioned Response As seen in the graphs, the strength of a conditioned response increases during acquisition, but drops dramatically during extinction; with the passage of time (represented by the rest interval), a previously extinguished response may reoccur (spontaneous recovery), but tends to be weak, and to extinguish relatively rapidly.

tioning which, as Gantt suggested, would be maladaptive. Clearly, if *all* conditioning persists indefinitely, then inappropriate responses become a serious possibility.

In part because of this possibility, Pavlov was also interested in the degree of permanence of classical conditioning. In order to test the limits, Pavlov and his colleagues first conditioned a dog to salivate at the sound of a bell. Once the response was well-established (by pairing the sound of the bell with a UCS of food placed in the dog's mouth), they continued to ring the bell, but no longer provided food. Under these conditions, the conditioned response (salivating) became weaker and weaker, and eventually ceased altogether. Pavlov referred to this cessation of responding when the CS is presented repeatedly *without* being paired with the UCS as **extinction**. Thus, extinction seems to suggest that what can be learned can be unlearned, and that conditioned responses are not necessarily permanent.

extinction in classical conditioning, the cessation of responding when the CS is presented repeatedly without being paired with the UCS.

However, this conclusion is not as straightforward as it seems. First, one should distinguish between active training in extinction, such as Pavlov used, and the persistence of conditioned responses in the absence of such training. Potentially, without active extinction, a conditioned response may simply remain dormant until the person encounters the CS again. For example, a fear response associated with the sound of a dentist's drill may persist despite lengthy intervals between visits, because no extinction training occurs. The effectiveness of extinction also depends on the *type* of conditioned response. Work by Gantt and others has indicated that internal responses like heart rate and blood pressure changes, which are frequently associated with stressful or emotional stimulus situations, are more persistent than simple muscle responses like withdrawing from a hot surface, or positive associations like salivation to food cues. This has implications for the conditioning of emotions, as we will see below.

Given that extinction occurs, one might ask what effect the passage of time will have on it. One might assume that the effects of extinction in 'erasing' the original conditioning would be as long-lasting as conditioning itself is in the absence of extinction. However, this is not really the case. Pavlov found that if he waited several hours after extinguishing salivation to the bell, ringing the bell tended to elicit the conditioned response again. While the response was weaker than when originally learned, and could in turn be re-extinguished, the most striking point was that it reoccurred *at all*. Pavlov called this return of the conditioned response **spontaneous recovery**, which is defined as the restoration of the response when the CS is presented after some time has elapsed since extinction training (see Figure 3.5). Spontaneous recovery of extinguished responses has been well demonstrated in a variety of species, sometimes after long time periods. This implies that, in terms of conditioning, what we learn is never really forgotten, but at best is simply overlaid with different experience. Instances where old fears re-emerge long after we thought we had conquered them (for example, fear of public speaking, fear of doctors, etc.) may reflect the enduring nature of conditioned behaviour. The results of research on extinction and spontaneous recovery suggest that conditioning is a 'one way street', whereby conditioned behaviour can be modified, but no conditioning is ever simply erased. Instead, extinction, and even new learning, are overlaid on earlier learning (Rescorla 2001).

spontaneous recovery in classical conditioning, the reoccurrence of the CR when the CS is presented after some time has elapsed since extinction training.

Higher Order Conditioning

We have seen that the principles of classical conditioning provide a mechanism whereby new stimuli can come to elicit a reflex response. Typically, the conditioned stimulus serves as a signal allowing anticipation of the UCS, which can be helpful to the organism. Sometimes, the sequence involves stimuli like food or water, which are beneficial to the individual. In other cases, the CS may signal something harmful, like heat or electric shock (for example, farm animals typically develop a fear of electric fences after a single experience of getting shocked). However, sometimes we encounter situations where the conditioned stimulus seems to have no direct connection to an unconditioned stimulus. For example, a child hears the word 'cake' and begins to salivate. How can this arise from the processes we have discussed?

Pavlov proposed a possible mechanism for such remote associations in terms of what he called **higher order conditioning**, where a previously-established conditioned stimulus is used as if it were an unconditioned stimulus to create conditioning to a new stimulus. While the description may seem complex, the process itself is easy to grasp. In Pavlov's original experiment, he first trained a dog to salivate to the sound of a buzzer (CS_1), using food as a UCS. Once conditioning was established, he introduced a new stimulus, a black square, which was repeatedly paired with the sound of the buzzer (but not food). After several such pairings, presenting the black square alone tended to elicit salivation (see Figure 3.6). Not surprisingly, the effect was rather weak, since each trial also functioned as extinction training for the original buzzer–food association. Pavlov called such conditioning *second order* conditioning, and tried to extend the sequence by using the black square as if it were a UCS, and attempting to link a new stimulus (*third order* conditioning). He found this was not possible when food was the UCS; however, he was able to create third order conditioning when conditioning leg withdrawal, with electric shock as the UCS. (This result may actually tell us more about the significance of aversive stimuli like electric shock than it does about higher order conditioning as such. As with extinction effects, the distinguishing element may be the use of stimuli associated with stress or negative emotions.)

higher order conditioning a form of classical conditioning in which a previously established conditioned stimulus is used as if it were an unconditioned stimulus to create conditioning to a new stimulus.

Dogs, of course, are not exactly like people, and the everyday world is not exactly like the laboratory. Without the controlled conditions of a laboratory, it can be difficult to prove the existence of higher order conditioning, but many researchers believe that analogues exist in human behaviour. For example, it is likely that the child who salivates to the word 'cake' has previously developed a conditioned response of salivating to the sight of cake (the original CS). Then, in the process of learning to talk, the word 'cake' became associated to the sight of the object it described; by repeatedly pairing word to the sight of cake, higher order conditioning became established.

It is tempting, certainly, to speculate on how language learning may be closely linked to higher order conditioning. Parents sometimes will verbalize (for example, 'that's bad!') while physically punishing a child; it is easy to understand how 'that's bad!' would then come to evoke fear and withdrawal, almost like a physical blow. If later, receiving a poor grade on a school report card leads to 'that's bad!', a poor grade could itself become a conditioned stimulus for fear and withdrawal. Precisely how significant this process is in everyday life, no one really knows. At the very least, we know that words *do* elicit emotional reactions, and such reactions are consistent with higher order conditioning. Rather than speculating further on this issue, let us examine some of the evidence for classical conditioning processes in everyday behaviour.

Applications of Classical Conditioning

While deceptively simple at first glance, classical conditioning seems to be a potent process for learning in a wide variety of species. In humans, salivation to food cues, fear arousal in the

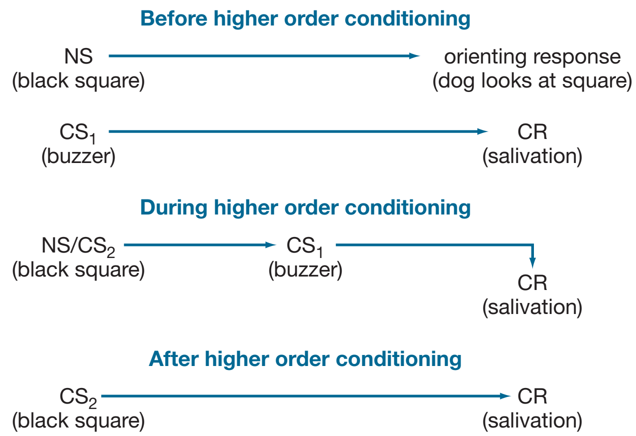

Figure 3.6 Higher Order Conditioning Pavlov found that once a conditioned response had been established, the CS could be used as if it were a UCS in order to develop further conditioning; compare this to Figure 3.1, and note where the UCS would be.

shower when the sound of a toilet flushing occurs, and a wide range of other instances show how many types of stimuli can come to elicit reflex responses. Sometimes, time itself can be a conditioned stimulus. For example, most pet owners notice that their pets are sensitive to the timing of daily routines, ranging from meal times to when a particular family member comes home. Since mammals (and a range of other species) have an internal biological clock, time of day can be an unseen stimulus for various responses, including hunger pangs. If you normally eat meals at a particular time of day, you have likely noticed that your stomach becomes active when that time approaches. This conditioned response enables your body to correctly anticipate the arrival of food (provided you don't skip a meal!). While these examples illustrate classical conditioning, note that some involve more than simply *motor reflexes*. To understand this more clearly, let us look at some other areas where the role of classical conditioning has been explored.

Conditioned Emotional Responses

Pavlov's work on classical conditioning became known relatively quickly – perhaps because he was already famous for his work on digestion. (By contrast, an American named Louis Twitmeyer, who discovered the same phenomenon almost simultaneously with Pavlov, died essentially unrecognized.) Among those who saw the importance of this new paradigm was John B. Watson, who saw in Pavlov's work a model for the behaviourist methodology he was trying to foster. One area where he saw potential was in the study of emotions, which had previously been the domain of the introspectionists. The introspectionists studied emotions, like other aspects of experience, by trying to describe the mental states involved. Watson instead believed that emotions represented observable responses, and proceeded to study the issue by attempting to create emotional responses experimentally. While he used a number of subjects, including his own children, the best-known case was a study done with Rosalie Rayner, using a toddler identified as Albert (Watson and Rayner 1920).

Albert was an 11-month-old boy who had been admitted to hospital for reasons unrelated to Watson's research. Watson initially observed Albert at play, and tested his responses to various stimulus objects, including blocks, a ball of cotton, some furry material and a white rat. The boy, like most children of his age, seemed curious about these objects, examining and playing with them. Then, Watson and Rayner began to systematically associate the white rat with the noise of a loud metal gong. On the first conditioning trial, Albert approached the animal without fear. Suddenly, the gong sounded behind him. The loud noise elicited a startle response (UCR), and also caused Albert to begin crying. Three times, the same sequence of events was repeated. Each time, Albert began crying at the sound of the gong. After a total of seven conditioning trials on two occasions, the white rat was presented without ringing the gong – and Albert began crying. Thus, a fear reaction had been classically conditioned to the rat, which previously had been a neutral stimulus. Watson called this fear a **conditioned emotional response**.

conditioned emotional response an emotional response such as fear which is established through classical conditioning.

About a week later, Watson and Rayner returned to test Albert again. This time, the experimenters showed Albert the objects from the original session. The toddler continued to show interest in the blocks and several masks. However, certain objects – balls of cotton, a white fur coat and a Santa Claus mask with a white beard – elicited the same crying and withdrawal as the white rat. In terms of classical conditioning, Albert had *generalized* his response to any white, fluffy stimulus!

Having established that fear could be classically conditioned, Watson and Rayner then sought to eliminate the fear response. To accomplish this, they used Pavlov's extinction procedure – presenting the white rat without pairing it with the sound of the gong. They tried this several times over a three-week period, but found that, contrary to their expectations, the fear did not extinguish. Unfortunately, before they could pursue the matter further, Albert was discharged from the hospital, ending the test.

With hindsight, we can recognize two factors that contributed to the failure of extinction. One is the fact that fear responses, like various other responses of the autonomic nervous system, are hard to extinguish (Gantt 1966). (Recall Pavlov's experiences with higher order conditioning – the conditioned fear produced by electric shock may account for his success in conditioning leg withdrawal to shock, but not salivation to food.) In addition, the occurrence of stimulus generalization, which is common for fear responses, tends to make extinction difficult, since a whole *range* of stimuli must be extinguished. Today, other techniques have been developed to deal with conditioned fear responses, since extinction training has such limited impact (see Chapter 9 for a discussion of such techniques, including *systematic desensitization*).

Before continuing our discussion of conditioned emotional responses, it is appropriate to consider the ethics of Watson and Rayner's study. Not surprisingly, they have been frequently criticized for the questionable ethics and potential harm of their test. Without attempting to second-guess past actions, it should be noted that the intent was not to permanently harm Albert; at the outset, Watson believed both conditioning *and* extinction would be successful. At the same time, it is clear that the procedure involved suffering for Albert, and it is unlikely that such a test would pass current ethical standards. (For a follow-up discussion of this study, see Harris 1979).

Watson and Rayner's demonstration, however questionable ethically, served to illustrate that emotional responses like fear could at least potentially arise from classical conditioning. In fact, most behaviourists would argue that phobias (a clinical category for irrational fears) can best be understood as conditioned emotional responses. Thus, anything from the fear of water to the fear of dogs could result from a traumatic episode in which the stimulus (water, dogs, etc.) was associated with a pain-evoking event.

However, human emotions extend well beyond fear. Could other emotions also be classically conditioned? Behaviourists would assert not only that such conditioning can happen, but also that it is responsible for most of the emotional richness of our lives. A new-born infant may instinctively respond to contact with the mother's body, but later this pleasurable response becomes associated to the mother's face, and still later to objects in the home, and maybe even to the home itself. Individuals who experience pleasure at hearing a favourite old song are experiencing emotions which have become associated to the conditioned stimulus of the music. Even when we go to the movies, conditioning is involved (probably through a higher order process) in our responses to heroes, villains and a variety of plot situations.

Words may even be the most refined of stimuli in terms of emotional conditioning. Words have a literal meaning and an emotional meaning; what is curious is that the two often do not correspond. For example, terms of endearment may range from the silly to the meaningless – 'little cabbage' or 'snuggie-poo'. Even more interesting is the emotional response to profanity. Generally, what are considered 'dirty words' varies from language to language – in English, most forbidden words relate to sexuality; in French, they usually relate to religion. Such differences relate not to the literal meaning of the words, but to the emotional significance of sexuality and religion in the respective cultures. When Shakespeare noted that 'a rose by any other name would smell as sweet', he recognized that the word is only a label – and labels depend on learning for their meaning. Without classical conditioning, it is likely that all language would be emotionally meaningless!

Try it Yourself

- When I (MH) was 3, a robin, protecting its nest, pecked me on the head. To this day, I have a fear of birds. How would Pavlov explain this? What were the UCS and the UCR? What are the CS and the CR today? How can you explain the fact that I have no fear of penguins or hummingbirds?
- *Jaws* is a classic movie depicting a huge shark killing swimmers in an Atlantic seaside town. When the movie first came out, many people who saw it became afraid to go swimming, even though they had never been attacked by even a small fish when swimming previously. How would classical conditioning principles explain this?
- Can you identify one fear which you feel affects you significantly? Can you recall a traumatic event that produced the fear (for example, a fear of dogs resulting from having been bitten as a child)? If not, do you think this invalidates the idea that phobias are based on conditioning?

Conditioned Drug and Immune Responses

As discussed in relation to the biological approach to psychology, the human body is a highly integrated system, involving neural, hormonal and immunological activity. Although we have not discussed the possible physiological mechanisms underlying classical conditioning, Pavlov believed that the mechanism was neural. Assuming this is true (and the available evidence supports this idea), one might still ask whether conditioning can influence other bodily processes, such as the response to drugs or disease. The exploration of such possibilities represents perhaps the most exciting area of conditioning research today.

Pavlov himself was interested in how drug reactions might be classically conditioned. In one study, the sound of a tone was repeatedly paired with a drug which induced vomiting (UCS); after several trials, the dog began to vomit at the sound of the tone alone. Similarly, diabetics taking insulin by injection sometimes show decreased glucose levels to cues associated with the injection. (Stockhorst *et al*. 2004) This suggests that stimuli present when a drug is administered may acquire the power to induce the drug's effects.

Interestingly, other work has suggested that with some drugs the conditioned response is the *opposite* to the primary effect of the drug itself. For example, rats were conditioned by giving injections of morphine in a specific environment. While morphine normally reduces sensitivity to painful stimuli, the rats after conditioning showed *increased* sensitivity to pain when placed in the conditioning context (Siegel 1976). This phenomenon of conditioning associated with drug use has been proposed as the basis of *tolerance* effects for addictive drugs, whereby repeated usage leads to lower response to the drug (MacRae *et al*. 1987; McDonald and Siegel 2004). Sometimes the cues involved are based on the location; for example, being in a pub can trigger both cigarette cravings and cardiac changes in smokers (Lazev *et al*. 1999). Internal bodily cues may also serve as conditioned stimuli. Thus, the physiological state or even the emotional state (such as anxiety) of the individual prior to using may trigger the desire for an addictive drug (Siegel 2005).

Why would conditioning mimic the effects of some drugs, and counteract others? At present we don't have a complete explanation, but the result seems to depend on the type of drug, and the body's response to it. For some drugs, like the vomiting agent used by Pavlov, the body reacts by showing a strong reaction, which gradually diminishes as the drug dissipates. By contrast, certain drugs, such as morphine, interact with the body's natural mechanisms for maintaining equilibrium (called *homeostasis*). In these cases, there is an initial reaction triggered by the drug, which is then followed by an opposite reaction, triggered by the body's homeostatic mechanisms. As a result, the conditioned stimulus becomes associated with the second reaction – which is opposite to that of the drug itself. Such *compensatory conditioned responses* have even been implicated in drug addiction fatalities: when drug addicts take drugs in an unfamiliar location (which doesn't trigger the offsetting conditioned response), the result can be an accidental (sometimes fatal) overdose (Siegel and Ramos 2002).

While drugs seem to interact with the body's own equilibrium processes, what about the immune system? Can classical conditioning affect the way our body reacts to disease? Research suggests this may be a real possibility. In one study, rats were given saccharin-sweetened water at the same time that they were injected with cyclosporine, a drug which inhibits immune system response. After several such pairings, tests showed that the taste of saccharin alone was able to suppress the immune system of the rats (Ader and Cohen 1975). Other studies have shown similar effects (see Ader and Cohen 1985). These studies imply that stimuli associated with low points of immune system functioning (for example, gifts received during a major illness, or objects associated with the death of a loved one) may continue to impair immune response at later times. Conversely, researchers have also been able to use conditioning to *enhance* immune system response (Alvarez-Borda *et al*. 1995; Gorcynski *et al.* 1982). While the practical implications have yet to be adequately tested, it may well turn out that conditioning effects can influence our long-term health (Ader 2003).

Research on conditioned emotional responses and drug/immune effects indicates that Pavlov's basic paradigm is still providing us with new insights on behaviour. The significance of classical conditioning is easily underestimated, since involuntary responses are often overlooked in our daily experience; this is partly because they *are* involuntary, and operate with no conscious intervention. Equally, conscious attempts at controlling reflexes have minimal success – as those who recognize that their fear responses are irrational can testify. While we cannot say at present that all issues related to classical conditioning have been resolved, neither have we reached the limit in terms of finding new applications and insights.

OPERANT CONDITIONING

As important as classical conditioning is, it must be recognized that it only deals with how new stimuli come to control existing involuntary responses. While reflexes and the 'gut-level' responses associated with emotions play a significant role in our everyday experience, most of our behaviour is self-generated, or *voluntary*. Behaviours like driving a car, working at a computer or calling a friend on the telephone are not elicited by conditioned stimuli. Instead, they are *emitted* – that is, generated by the individual as a way of influencing the surrounding environment. In order to understand the dynamics of such behaviour, we need to consider a different approach to learning, This approach, called **operant conditioning**, deals with how voluntary (emitted) responses change over time as a function of their consequences. For example, if Johnny climbs a tree (voluntary response) and gets hurt (consequence), he may not climb the tree thereafter.

operant conditioning in the behaviourist approach, the form of learning concerned with changes in emitted responses (voluntary behaviour) as a function of their consequences.

To understand the origins of operant conditioning, we need to go back to the work of Edwin Thorndike. Although Thorndike was studying animal behaviour and learning at roughly the same time as Pavlov, his approach was very different. Whereas Pavlov began with an interest in digestion and then became interested in conditioning, Thorndike was initially interested in problem solving and intelligence. Consequently, instead of focusing on simple reflexes, Thorndike studied situations where an animal was actively interacting with its environment. In a typical experiment, a cat was confined in a 'puzzle box', a cage-like structure from which it could free itself by pressing a lever. As any cat owner can testify, cats generally dislike confinement; hence the cats were typically eager to escape. However, Thorndike increased the incentive by depriving the animals of food prior to testing, and then placing food outside the puzzle box, where it was visible to the cat. Not surprisingly, the cats learned to press the lever, thereby escaping and obtaining the food (Thorndike 1898; see Figure 3.7).

Two primary conclusions emerged from Thorndike's work. The first was that if one measured *how long* it took a cat to escape, the time gradually declined with repeated trials. This improvement in performance represented a change in behaviour as a result of experience – in other words, *learning*. The other major conclusion concerned the relationship between the cat's behaviour and its consequences. Both escape and obtaining food appeared to be desirable to the cats, leading Thorndike to conclude that the satisfying outcome was what led to the behaviour being repeated. By contrast, flailing at the walls of the box, and other behaviours which did not lead to escape, declined. From observations like this, Thorndike formulated his *law of effect*: behaviour which leads to a satisfying outcome tends to be repeated, while behaviour that leads to an unsatisfactory outcome is unlikely to be repeated. Note that the law of effect makes no reference to reducing hunger, desire to escape or other mentalistic concepts. The observational nature of the principle was one of the factors which attracted the attention of Watson and later behaviourists.

Thorndike's research laid the foundation for the study of non-reflex behaviour. By emphasizing the connection between an action and its outcome, his law of effect provided a framework for studying such behaviour. In Thorndike's system, responses are initiated by the organism as part of dealing with its surroundings, not as a reflex triggered by an environmental stimulus. Depending on the consequences, a particular behaviour might or might not be repeated in the future. For example, a child who draws a picture and presents it to Mum may receive praise; this will encourage the child to draw more pictures in the future. On the other hand, if the child takes a cookie without permission and is scolded, they are less likely to try this again. In its simplest form, the law of effect reaffirms what might be considered 'common sense'. At the same time, by suggesting a simple framework for the study of non-reflex behaviours, it fits well with the developing behaviourist approach. Yet, while Thorndike's work was acknowledged by Watson, and stimulated a variety of subsequent research, it did not result in a coherent system comparable to Pavlov's paradigm until the work of B. F. Skinner in the 1930s.

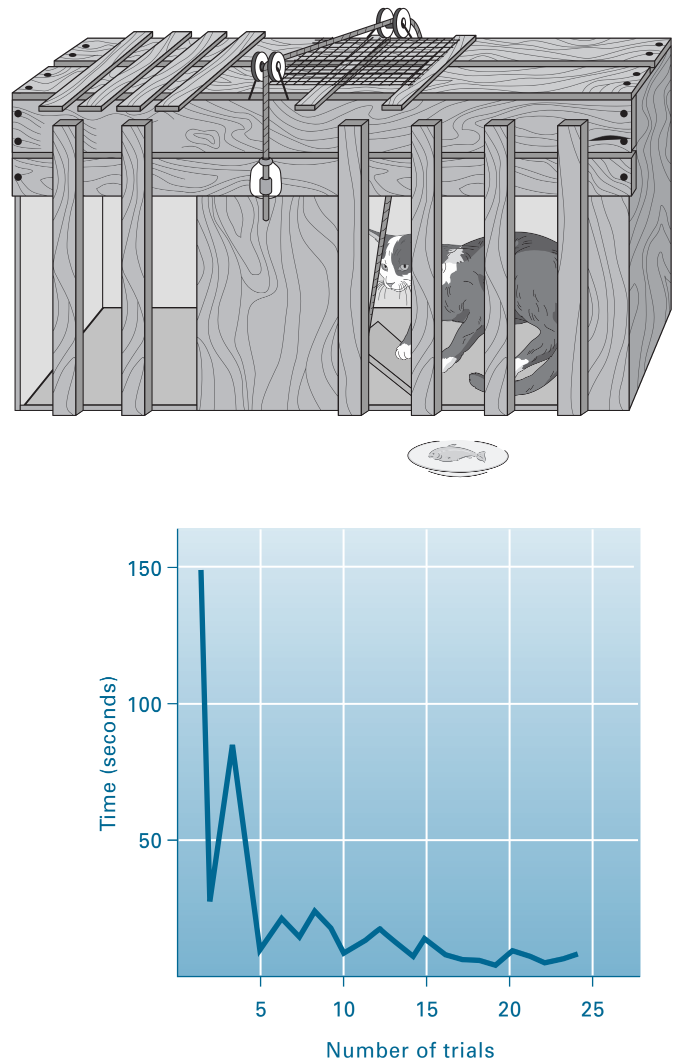

Figure 3.7 Thorndike and the Law of Effect Thorndike explored learning by placing cats in a box like that shown here; pressing on the pedal activated a mechanism which allowed the cat to get out of the box. As shown by the graph, over several trials, the cat learned to escape much more quickly; Thorndike explained this by saying behaviour changed as a result of the desirable consequence of escape.

Skinner and Operant Conditioning

Within behaviourism, B. F. Skinner occupies a position of influence equal to, and in some ways greater than, that of John B. Watson. As the pioneer of operant conditioning, he almost singlehandedly created a framework for the study of learned behaviour. Skinner's contributions are significant in terms both of research methods and conceptual analysis. To understand this, we need to consider the origins of his work.

Try it Yourself

'Satisfaction', like beauty, is in the eye of the beholder. We all differ in terms of what we consider to be satisfying, and sometimes what other people find satisfying surprises us. Consider, for example, the foods you like to eat and the clothes you like to wear. Clearly your preferences are not those of everyone else. Make a list of some things you would find 'satisfying' and some that you would find 'unsatisfying'. Ask a friend to do the same and compare your lists. Since friends often become friends because of their commonalities in what they enjoy, you will probably find many commonalities on your two lists. But you will undoubtedly find many differences. Ask an older person, a parent or grandparent perhaps, to make up a list as well. There are probably fewer commonalities between this list and the lists of you and your friend, and many of the differences reflect the age/generational differences between the list-makers. Keep the differences in lists in mind the next time you buy a gift for someone: we often select a gift thinking of what *we* would find satisfying instead of what the *recipient* would find satisfying!

While training as a graduate student at Harvard, Skinner was doing studies of animal behaviour somewhat similar to Thorndike's. Influenced by Watson, he found himself frustrated that so much of the vocabulary of psychology seemed clouded by the ambiguities of everyday language. In particular, he felt that terms referring to mental states (for example, 'drive', 'belief', 'intent') were both vague and unnecessary to the understanding of behaviour. For Skinner, the inner workings of the mind (and the body) were a 'black box', inaccessible to direct observation. This point of view was shared by all behaviourists, but Skinner went further, arguing that even if thoughts and other mental states *could* be studied, they would have no real value in explaining behaviour. Instead, the environment in which a response occurs, the response itself, and the response's consequence are all that are necessary to understand behaviour. By insisting that mental states are both inaccessible to study *and irrelevant to understanding behaviour*, Skinner was advocating a point of view which has come to be called **radical behaviourism**.

radical behaviourism a position adopted by Watson and Skinner which argues that mental states are both inaccessible to scientific study and irrelevant to understanding behaviour.

Given his concerns about amibiguities in language, one of Skinner's first goals was to develop new terms for describing and analysing behaviour. He began by coining the term *operant conditioning* to replace Thorndike's 'instrumental learning'; similarly, he renamed classical conditioning as 'respondent conditioning'. He referred to emitted behaviours as operant responses, arguing that 'voluntary behaviour' implies undesirable notions about free will. By developing this new vocabulary, he attempted to purge the study of behaviour of all excess conceptual baggage (Skinner 1987). (Skinner even went so far as to say radical behaviourism is not a part of psychology, but rather an approach to understanding certain issues both inside and outside psychology.) Operant conditioning has in fact become established as a major form of behaviourism, as we shall see. While Skinner often argued that his approach is pragmatic, not theoretical (Skinner 1950), his critics have disagreed. In fact, Skinner's framework is generally regarded as a *meta-theory* – that is, a theory about what makes a good theory of behaviour. In this sense, the apparent simplicity of his ideas can sometimes be deceptive.

Key Thinker: Burrhus Frederic Skinner

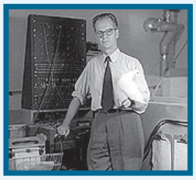

Burrhus Frederic Skinner (1904–90) is probably the best-known American behaviourist, and the founder of operant conditioning. His early years were rather peripatetic – educated at Hamilton College and then Harvard (receiving his PhD in 1931), he taught at the University of Minnesota and Indiana University. During World War II, he did research with a military flavour, including a programme designed to teach pigeons to direct missiles to targets while flying in the nose-cone; the technique was never implemented. In 1947 he returned to Harvard to deliver the annual William James Lectures; in 1948 he was appointed as a full professor at Harvard. Skinner's development of operant conditioning began while training as a graduate student at Harvard. Although his initial work on animal behaviour was somewhat similar to Thorndike's, he became influenced by Watson's ideas, and began a systematic attempt to purge psychology of mentalistic concepts and language. While his theories have remained controversial, the practical applications of operant conditioning have been widespead. Skinner died in Cambridge, Massachusetts in 1990.

A second key aspect is the interaction between the *concepts* of operant conditioning and the *procedures* used – that is, the methods of observing and measuring behaviour. As a graduate student doing research with rats, Skinner found that the typical learning tasks, like mazes or Thorndike's puzzle boxes, required extensive labour. For each trial, the researcher had to put the animal in the box, record behaviour, retrieve the animal after the trial, etc. In order to simplify this process, Skinner developed an apparatus which would allow running continuous trials, with behaviour automatically recorded. He called this apparatus (which resembled a small box with a lever within) an 'auto-environmental chamber', but it became known (to Skinner's lasting dismay) as a Skinner box! While it accomplished Skinner's basic goal of automating Thorndike's approach, it also led to other consequences. The most notable of these involves the way behaviour was measured. Since pressing the lever in the Skinner box could be considered analogous to the cat pressing the lever in Thorndike's puzzle box, counting the number of presses (that is, the frequency of response) became the standard measure of operant learning. In some respects, this is unfortunate, since it has led to operant conditioning mostly considering *only* the frequency of behaviour; as a result, aspects of behaviour such as intensity, duration or quality of responses have been largely ignored. (Consider, for example, the many different aspects of a response like hitting a tennis ball – especially if one is comparing an amateur and a pro player!) While the focus on frequency was a practical consideration, it eventually became part of the overall conceptual framework as well.

Although research based on analysing the frequency of behaviour has often been highly productive, it should be noted that in everyday life, frequency is not always the most meaningful aspect of behaviour. (For example, should we judge the quality of an artist by *how many* works they create, or should we look at the *content* of their work?) Thus, operant conditioning, while claiming to be a pragmatic analysis unencumbered by theory, in fact has evolved out of a unique set of assumptions about both theory and methodology. As with other approaches to psychology, recognizing the foundations of behaviourism should help in comprehending where it has led in the understanding of behaviour.

Reinforcers and Reinforcement

One of the first issues which Skinner attempted to address was Thorndike's law of effect. While it is intuitively obvious that a response which leads to a satisfying consequence is likely to be repeated, Skinner was bothered by the vagueness of 'satisfying'. To avoid this, he coined a new term, 'reinforcer'. A **reinforcer** is a stimulus which, when it follows a response, results in a change in the probability of the response recurring. Thus, unlike notions of satisfaction, a reinforcer becomes an *observable environmental event*. **Reinforcement** is the process by which a reinforcer *increases* the probability of a response. (Note that in talking about probabilities, one is implicitly describing how *often* a response occurs – i.e., frequency.)

reinforcer a stimulus which, when it follows a response, alters the probability of the response recurring.

reinforcement the process by which a reinforcer increases the probability of a response.

The most basic reinforcers are those which are related to survival, such as food or water. Such reinforcers are described as **primary reinforcers**, since they have an innate biological significance. For example, a baby cries because it is hungry. When it receives food, this serves as a reinforcer for the response of crying. As a result of this reinforcement process, the baby is more likely to cry the next time it is hungry. (Hunger, of course, is a reference to an internal state which can't be directly observed. To avoid such terms, Skinner would talk about the length of time since the baby was last fed. Any parent who has monitored a baby's feeding schedule can testify that this is a reasonably accurate gauge of hunger!) While food and water are the most common primary reinforcers, many other items (including clothing when it's cold, air to breathe, and drugs such as nicotine or opiates) also seem to function as primary reinforcers. Note that primary reinforcers also typically elicit some form of reflex response – that is, they are *also* unconditioned reinforcers, in terms of classical conditioning. Food, for example, is a positive primary reinforcer, but also elicits salivation. This dual nature underlines the fact that primary reinforcers seem to have direct biological significance.

primary reinforcer a stimulus whose capacity to act as a reinforcer is based on an innate biological significance, such as food, water or electric shock.

By contrast, there are a large number of environmental events which seem to act as reinforcers, but are not based on biological survival. For example, attention, praise, money and trophies can all act as reinforcers. Reinforcing stimuli like these are described as **conditioned reinforcers**. As the name suggests, conditioned reinforcers are stimuli which assume reinforcing properties because they have been reliably associated with a primary reinforcer (this is actually a form of classical conditioning, with the conditioned reinforcer and primary reinforcer related as CS and UCS, respectively). For example, Skinner has argued that attention becomes a conditioned reinforcer in early infancy because it precedes (and is therefore associated with) receiving primary reinforcers: the baby cries, an adult gives attention by coming to see what is wrong, and then the adult provides a primary reinforcer like food or a dry nappy. A young child may receive praise for a particular action, and receive a cookie; soon, praise itself becomes a reinforcer, because it is associated with the cookie. Later, other conditioned reinforcers may develop as stimuli are paired with existing conditioned reinforcers. For example, the adult often smiles while giving praise to the child, and the smile alone becomes reinforcing on its own. (Recall how higher order conditioning allows new stimuli to become linked to existing conditioned stimuli.) Since conditioned reinforcers are based on learning, not innate factors, the potential range of such reinforcers is virtually unlimited. Perhaps the most powerful conditioned reinforcer in our society is money, which can be used to obtain a wide range of other reinforcers, both primary and conditioned.

conditioned reinforcer stimuli which act as reinforcers but are not based on biological survival, such as praise, money or criticism.

Reinforcement, the process of increasing the frequency of a response by means of a reinforcer, is at once both simple and subtle. One element which is important for proper reinforcement is contiguity – that is, the reinforcer should immediately follow the response. If a child does something desirable, then praise should be given *immediately*; if not, one runs the risk that the reinforcer will influence a subsequent response. For example, 2-year-old Johnny uses the toilet instead of wetting his pants. Half an hour later, Johnny's mother realizes what has happened, and praises Johnny – who is now engaged in pulling books out of the bookcase. In this situation, the positive reinforcer is unlikely to be strongly associated with going to the toilet, and may in fact reinforce the less desirable current behaviour! Parents sometimes misunderstand the significance of contiguity, but it is a powerful factor in operant conditioning. When the wrong response is reinforced, the tendency is to assume that the principles don't work. But Skinner once commented that if an experiment doesn't turn out the way we expect, there is a temptation to tell the animal, 'Behave properly!' However, in such cases, the animal *always* behaves – the error is in our understanding (Skinner 1967).

The world today: **Video Games**

In the past 20 years, a new form of entertainment has entered our lives: video games. Some people say they are good, and some say they are bad. Good or bad (and there may be elements of each), there is no question that many people are spending hours each day playing these games, some even to the extent where they seem to be addicted. Why do people spend so much time on these games? Skinner would answer, 'Well obviously, they are being reinforced for playing.' By Skinner's definition of positive reinforcement, it *is* obvious: positive reinforcement increases the behaviour. And video games are highly reinforcing. In the first place, for most video games, feedback (which in itself is reinforcing) comes immediately: information that one has responded correctly comes immediately after the response and is often accompanied by points, flashes of light, screen congratulations, advancement to a higher level of game-playing, etc. The feedback is not the only reinforcing part of video games, however. As in any other situation, what is reinforcing to one person may not be reinforcing to another. For example, reinforcement for successful playing may be the admiration of peers, mastering the game, achieving a 'personal best' and so on. Richard Wood and his colleagues (2004) (some of whom are 'gamers' themselves) conducted an online survey of 382 university students who played a variety of video games, asking them about what they found important (i.e., reinforcing) in the structure of various games. Here is what they found:

- Good quality, realistic sound and graphics were important for both males and females, although the exact form of the graphics in particular revealed gender differences. Males prefer realistic graphics with competitive, action-based, fast-moving events, such as simulated battles and sports contests, while females prefer gentler, cartoon-like graphics with fantasy-based themes.
- The ability to compete with or form alliances with other players online was considered important/reinforcing for males, although less so for females. (Can you speculate on why this might be?)
- Both males and females found the ability to get into the game and to advance in the game quickly to be important.

Not surprisingly, success at playing video games makes people feel good and induces them to play more (Chumbley and Griffiths 2006). No wonder it's so easy to get hooked!

Misunderstanding how reinforcement works can lead to difficulties in trying to change behaviour. In everyday life, people tend to invoke rationalizations for failing to change, particularly when it comes to bad habits. A would-be dieter, or a heavy smoker, might say, 'I lack willpower.' Skinner would argue that such phrases have no real explanatory value, and in fact obscure the actual dynamics of the situation. A smoker, for example, may worry that they will develop lung cancer some years hence, and thus wishes to quit. When the attempt fails, the person concludes that it is because of personal weakness. In fact, they are ignoring the actual reinforcers – the nicotine, the conditioned reinforcement of smoking being associated with enjoyable moments, possibly peer approval, and so forth. Compared to these *immediate* reinforcers for the act of smoking, the perceived value of better health some time *in the future* has much less effect on the response of quitting. Thus, both *what* the reinforcers are, and *when* they are received, is crucial to understanding the dynamics of behaviour. For example, the reinforcing value of video games is discussed in *The World Today*.

As noted previously, reinforcers are stimuli that alter the probability of a behaviour. While this is correct, it should be noted that, thus far, we have been talking as if all reinforcers are desirable things. However, just as Thorndike distinguished between 'satisfying consequences' and 'unsatisfying consequences', Skinner recognized that not all reinforcers are alike. In fact, the examples we have been discussing all represent **positive reinforcers**, which Skinner defined, not by intuitive notions of satisfaction, but by the effect on behaviour: if a reinforcer which follows a response makes that response more likely in the future, it is a positive reinforcer. In everyday terms, a positive reinforcer is similar to what is often called a 'reward' – but of course, Skinner would reject such words as being too imprecise for scientific purposes. By contrast, a reinforcer which makes the behaviour *less* likely in the future Skinner called a **negative reinforcer**. This would be equivalent to Thorndike's dissatisfying consequences; in everyday terms, something unpleasant or aversive. (Note: some researchers refer to 'punishers' instead of 'negative reinforcers'; however, this discussion will follow Skinner's usage.) Like positive reinforcers, negative reinforcers can be either primary or conditioned: negative primary reinforcers are consequences which cause harm or threaten survival, such as a physical blow or electric shock. Negative reinforcers may also be conditioned, by being associated with negative primary reinforcers. For example, a child may learn to associate criticism or the word

positive reinforcer a stimulus which when it follows a response serves to increase the probability of the response in the future.

negative reinforcer an aversive stimulus which when it follows a response serves to decrease the probability of the response in the future.

	Primary	Conditioned
Positive	Food Water Shelter	Attention Praise Money
Negative	Physical blow Burning heat Electric shock	Being ignored Criticism Fines

Figure 3.8 Types of Reinforcers Both positive and negative reinforcers can be either primary or conditioned, as the chart shows. Note that the reinforcers mentioned are examples, not a complete list of possibilities; Skinner would also argue that defining any reinforcer can only be done with reference to a specific situation!

'bad' with being hit; criticism then becomes a negative conditioned reinforcer. (For an overview of the various types of reinforcers, see Figure 3.8.)

At first glance, talking about positive and negative reinforcers may not seem much different from Thorndike's satisfying and unsatisfying consequences. Indeed, some critics have suggested that both the law of effect and the definition of a reinforcer are circular, since we cannot determine the value of a stimulus until we observe how it affects behaviour. For example, we only know something is a positive reinforcer when we see that it results in an increase in the probability of behaviour. This issue is still a source of debate, but it should be noted that Skinner went much further than Thorndike in analysing how the relationship between a response and a reinforcer affects behaviour, as we will see in the next section.

Contingencies of Reinforcement

In general, operant responses are freely produced by the individual, but the likelihood of making a response is determined by its consequences on previous occasions. For example, if Tim is given a cookie for having finished his peas at dinner, he is likely to eat his peas in the future. Thus, there is a relationship between the behaviour (the response of eating peas) and its consequence (the cookie as positive reinforcer). In Skinner's terminology, the relationship between a response and a reinforcer is called the **contingency of reinforcement**. (A 'contingency' describes how something depends on another event.) As he realized, identifying the contingency is a powerful tool for understanding changes in behaviour.

contingency of reinforcement a description of the relationship between a response and a reinforcer.

One type of contingency is *reinforcement*, as already mentioned. Reinforcement *always* results in an *increase* in the likelihood of a response. In the example above, it is easy to recognize that a biscuit is a positive reinforcer, and that the likelihood of eating peas will increase. Thus, when a response is followed immediately by a **positive reinforcer**, the response becomes more likely; Skinner called this process **positive reinforcement**, because it is reinforcement using a positive reinforcer.

positive reinforcement a process of increasing the probability of a response by immediately following the response with a desirable stimulus (a positive reinforcer).

A second possible contingency is when a response is immediately followed by a *negative* reinforcer. Since this represents an aversive consequence, do you think that it would make the response more likely in the future? Obviously, the answer is no – in fact, the response would become *less* likely. For example, 3-year-old Sally pokes a pin in an electric outlet and receives a shock. In the future, Sally is not likely to repeat this action! Since the probability of the response does not increase, this cannot be termed a case of *reinforcement*. In everyday life, one would likely call it 'punishment' – and that is the term used in operant conditioning as well. (It is one of the rare instances where Skinner adopted a term with obvious everyday meaning!) He defined **punishment** as a process whereby a response is followed by a *negative reinforcer*, which results in a decrease in the probability of the response.

punishment a process whereby a response is followed by a negative reinforcer, which results in a decrease in the probability of the response.

The distinction between positive reinforcement and punishment gets to the heart of Skinner's framework. In order to understand the dynamics of learning, one must be able to identify the contingency which is involved. For example, parents and teachers often react to a child who is misbehaving by scolding the child. The intent, of course, is to decrease the undesirable behaviour – that is, to use *punishment*. What sometimes happens, however, is that the child continues to misbehave, and may even become more disruptive. The frustrated adult exclaims, 'I don't know what's wrong with Johnny! The more I punish him, the worse he behaves!' Skinner would respond by looking at the situation from the child's point of view (since the child is the one receiving a reinforcer). Given that the disruptive behaviour *increases*, Skinner would say that obviously the child is receiving reinforcement, and so the reinforcer (scolding) is actually a positive reinforcer for the child! At first glance, this may seem silly, but in fact, scolding requires paying attention to the child – and attention is a powerful positive reinforcer, especially for a child who feels neglected. What typically happens is that a busy adult ignores the child who plays quietly, but immediately responds to misbehaviour; so, in order to get attention, the child misbehaves more and more. The moral of this example is that the *organism* (i.e., the child) determines the significance of the reinforcer, *not* the environment which delivers the reinforcer (the adult, in this case).

Reinforcement and punishment represent the most common contingencies in operant conditioning – they are roughly equivalent to the old notion of using 'the carrot and the stick' to train a mule. However, reinforcers can also be related to behaviour in other, more indirect, ways. For example, it is also possible to produce an increase in behaviour by *terminating* or *withholding* a negative reinforcer (an aversive stimulus); this process is called **negative reinforcement**. For example, a teenager is nagged by a parent to clean up a messy bedroom. In this situation, the nagging is unpleasant – a negative reinforcer. When the teenager eventually (albeit reluctantly) cleans up the room, the parent stops nagging. If we look at this from the viewpoint of operant conditioning, the desired response is cleaning up the room. As long as the response is *not* made, a negative reinforcer is presented (the nagging). When the response is finally made, the negative reinforcer ceases!

negative reinforcement a process for increasing the probability of a response in which a response immediately leads to termination or withholding of an aversive stimulus (negative reinforcer); note that since the response increases in frequency, it is not equivalent to punishment.

In the example just given, the teenager reacts to eliminate the nagging – that is, to escape from an existing negative reinforcer. In the future, they might respond at the first hint of parental displeasure, before the nagging actually begins. In this case, they would be responding before the reinforcer is given, in order to *avoid* it. (In this case, making the response leads to the withholding of the negative reinforcer.) Thus, negative reinforcement actually has two variations, *escape* and *avoidance*. Normally, initial learning requires presenting the negative reinforcer until the response is made (that is, escape); later, the individual anticipates the sequence, and responds before the negative reinforcer is presented (that is, avoidance). Experiencing a 'sigh of relief' after getting out of an unpleasant situation (such as leaving the dentist's office) is characteristic of *escape*. Similarly, the anticipatory fear that you feel in some situations (for instance, if a large, unkempt stranger approaches you on an isolated street) can trigger a response (such as crossing to the other side of the street) in order to *avoid* an anticipated unpleasant situation (a confrontation with a hostile stranger).

If you review the foregoing discussion of contingencies, it may occur to you that there is a fourth possibility, based on terminating or withholding a positive reinforcer. How would you react if your behaviour led to losing a positive reinforcer? For example, a teenager comes home very late, and loses their driving privileges as a result. As you can imagine, when a response leads to terminating or withholding a positive reinforcer, the behaviour becomes less likely. (This contingency is generally called **omission**, or sometimes 'postive punishment'). Thus, the effects of omission, in terms of reducing the likelihood of a response, are similar to the effects of punishment (see Figure 3.9 for a summary of the four contingencies).

omission a process whereby a response is followed by terminating or withholding a positive reinforcer, which results in a decrease in the probability of the response.

Type of reinforcer	Desired change in behaviour: Increase response	Desired change in behaviour: Decrease response
Positive reinforcer	Positive reinforcement	Omission (withholding possible reinforcer)
Negative reinforcer	Negative reinforcement (escape, avoidance)	Punishment

Figure 3.9 Contingencies of Reinforcement Skinner argues that in order to understand how operant responses change, one must look at both the type of reinforcer, and its relationship to the response. (See text for definitions of forms of reinforcement.)

To recap what has been said about the process of reinforcement, one can understand the dynamics of behaviour by identifying the *contingency of reinforcement* involved. To do so, one must identify the response and the reinforcer, *and* how they are related. In doing so, one must remember that the value of the reinforcer is determined by the organism, not the environment. An example might help to clarify this.

Imagine that you are offered a chocolate milkshake if you will sing a song. Assuming you like milkshakes, you will likely sing; thus, the milkshake is a positive reinforcer. Based on your rousing performance, you are offered a second shake if you sing another song. You do so, but drink the second shake more slowly. When you finish, you are offered a third shake for another song. At this point, the prospect of consuming another milkshake is very unappealing, and you refuse to sing. Thus, what started out as a *positive* reinforcer has now become a *negative* reinforcer. The shakes haven't changed, but their value to the organism has – and that is the crucial point. In order to understand behaviour, one must look at how the behaviour changes in order to identify the contingency involved. As Skinner said, the organism always behaves, it is our understanding that is sometimes wrong.

Try it Yourself

In order to understand the contingencies of operant behaviour more fully, consider the following situations:

- You have worked very hard and your employer wants you to keep on working hard. How might he or she do this? Would you prefer a raise in pay, or movement to a better office where there is less noise and fewer distractions? Would both be positive reinforcers for you, even if differing in value? If one was actually a negative reinforcer, how would this affect your working?
- Unthinkingly, you said something that hurt your friend's feelings. Your friend now has certain options. For example, they may respond angrily to you, or they may stop speaking to you for some period of time. With both of these options, you might learn to stop hurting your friend's feelings (i.e., this behaviour would decrease). Are they the same contingency? Which would you prefer? Why?

Operant Conditioning Phenomena

Shaping and the Learning Process

In all the examples we have discussed, the reinforcer was used to alter the likelihood of an existing response. While this shows the power of reinforcement, it also poses a problem. Since one cannot reinforce a response that doesn't occur at all, how do new behaviours arise? And how can operant principles explain the development of complex behaviours? For example, how does a child learn to walk, or an adult learn to play the piano? One factor to recognize is that complex behaviours do not suddenly emerge fully formed. Instead, they tend to be formed out of a series of simpler behaviours, which can then be combined. A child learns to crawl, and to pull itself upright, before taking its first steps. Piano playing involves a whole set of responses, from learning how to position the hands and body at the piano, to identifying written musical notes, to controlling the pedals while playing. Thus, complex behaviours can be thought of as a series of simpler responses that are combined as a sequence, which is then treated as a single response in terms of getting reinforcement.

While this description can account for complex responses, it still doesn't explain how new responses arise. Skinner explained variability and originality as forms of 'behavioural drift'. That is, operant behaviour, in the absence of reinforcement for a specific response, tends to vary somewhat over time. Much of this variation is simply random, but the fact that drift occurs means that sometimes new responses will occur – and therefore may be reinforced. This means that desired new behaviours can be encouraged through a process called shaping.

shaping the process of guiding the acquisition of a new response by reinforcing successive approximations to the desired response.

Shaping is defined as the process of reinforcing *successive approximations* to a desired response. The process assumes that someone (an experimenter, a parent, etc.) has in mind a behavioural goal, and can control the delivery of a reinforcer accordingly. For example, most operant research involves animals (such as a white rat) pressing a lever in a Skinner box. While rats are capable of pressing a lever, it is not a natural response in the wild. Consequently, the rat must be shaped to acquire the response. (As graduate students working as lab assistants, we had many experiences of doing such training.) Typically, when one places an untrained rat in the Skinner box, it begins to explore this new environment, looking around and sniffing at everything. In one corner of the chamber, there is a food dispenser which can deliver pellets of dry food. At first, one simply dispenses pellets, one at a time, until the rat associates the click of the mechanism with arrival of a food pellet. Once this pattern is established, a food pellet is given only when the rat turns towards the lever (a first approximation to the desired response of pressing the bar). After eating, the rat will likely turn back towards the lever. If it moves closer, or lifts a paw towards the lever (a closer approximation to the desired response), another pellet is given. Finally, the pellet is only given when the rat actually presses the lever – and at this point, the desired shaping has been achieved.

Shaping is a simple concept, but very powerful. Although they might not use the terms of operant conditioning, generations of animal trainers have applied the same principle in their work. There is even a story, probably apocryphal, about a class of psychology students who shaped their professor to stand in the corner. The students used writing in their notebooks as the reinforcer (professors tend to find this reinforcing, since it implies both paying attention and being interested). Whenever the professor, who tended to walk about, moved towards one corner, the students all wrote furiously. When he moved in the opposite direction, they all put down their pens. By the end of the class, the story goes, the poor professor was wedged into the corner!

The process of shaping has also been applied to more serious purposes, including assisting children with language learning and helping accident victims to reacquire basic skills. Anyone who has watched trained animals perform, whether at a circus or in a movie, has also witnessed the power of shaping. Having said this, it is also worth noting that shaping is easier to grasp as a concept than to apply in practice. I (WG) can still recall my first attempts to shape a rat to press the lever in a Skinner box – it took nearly an hour. Later, after gaining more experience, I could typically do it in 15 minutes! The difference was one of judgement and timing – deciding when a new approximation was good enough to merit a reinforcer. If I rewarded too often, the rat spent more time eating than learning; if I rewarded too infrequently, the rat lost interest, or seemed to forget what the last reinforced response had been. In this sense, shaping, while clearly consistent with operant principles, is not simply a mechanical process.

Extinction

Shaping uses the variability inherent in responding in order to produce a desired response. Once a response occurs, it can be reinforced, as we have seen. But what happens if the reinforcement is discontinued? Consider a rat that has been reinforced with a food pellet each time it presses the lever, or a child that has been praised each time they pick up their toys. If the situation changes so that reinforcers are no longer given, what will the organism do? An intuitive answer, which is supported by research, would be that the behaviour might continue for a short time, but once it is recognized that reinforcers are no longer forthcoming, the behaviour will decrease in probability. This drop in responding when reinforcement is discontinued is called **extinction**. (One can see a parallel to extinction in classical conditioning, where the CR disappears when the CS is no longer paired with the UCS.) Note that while both extinction and punishment in operant conditioning produce decreases in responding, they do so in very different ways: extinction can be considered a passive process, in that it diminishes the value of the response by eliminating the reinforcer which supported it. Punishment, on the other hand, uses an aversive stimulus to actively suppress the (undesired) behaviour.

extinction in operant conditioning, a drop in responding when reinforcement is discontinued.

One implication of extinction seems to be that, in order to be effective, reinforcement must be continuously given for every response. Unfortunately, this seldom occurs in the real world. Students study regularly, but receive reinforcement only after the occasional test. People go to work every day, but may be paid only weekly or even monthly. How, then, can operant conditioning be said to apply to such behaviours?

Schedules of Reinforcement

Early in his research, Skinner recognized that in everyday life we rarely experience either true extinction (no reinforcement at all) or **continuous reinforcement** (every response reinforced). Instead, what we tend to encounter is something in between – some responses get reinforced, and some don't. Skinner coined the term **partial reinforcement** to describe situations where reinforcement is given only intermittently. In order to understand what happens under such circumstances, he began a series of studies looking at various forms of partial reinforcement. In order to distinguish various types of intermittent reinforcement, Skinner coined the term **schedules of reinforcement** (see Ferster and Skinner 1957). In general, a reinforcement schedule describes when a reinforcer is given, in much the same way that a train schedule describes when a train departs.

continuous reinforcement a reinforcement schedule in which every response is followed by a reinforcer; equivalent to a FR 1 schedule.

partial reinforcement a contingency of reinforcement in which reinforcement does not follow every response.

schedule of reinforcement a description of the conditions which determine when a response will be followed by a reinforcer.

The most straightforward schedule, of course, is continuous reinforcement, since every response receives a reinforcer. By contrast, partial reinforcement can occur under an essentially infinite number of variations. However, a surprising variety of situations can be described in one of two ways: according to the number of responses made before a reinforcer is given, or the amount of time that elapses between reinforcers. Schedules which depend on the number of responses made are called *ratio schedules*; those which are time-dependent are called *interval schedules*. In addition, such schedules may be very regular (for example, every third response, every 15 seconds), or somewhat unpredictable. Regular schedules are called *fixed schedules*, while those which are more unpredictable are called *variable schedules*. Let us look at the different types more closely in order to see how they affect behaviour.

fixed ratio schedule a reinforcement contingency defined by the number of responses the organism must make in order to get a reinforcer; the ratio is measured as FR *x*, where *x* is the required number of responses.

Fixed ratio schedules are the simplest to understand: the ratio is measured as FR *x*, where *x* is the number of responses the organism must make in order to get a reinforcer. (Continuous reinforcement is actually equivalent to FR 1, since each response leads to a reinforcer.) Thus, FR 5 means that every fifth response would receive a reinforcer (see Figure 3.10a). In everyday life, this is analogous to piecework, where a person will be paid according to the number of responses made. An example might be a seamstress: each shirt may have seven buttons to be sewn, and completing a shirt earns $1; this would be FR 7. Skinner found in his experiments that increasing the ratio tended to increase the rate of responding; the animals sought to maintain the total amount of reinforcement received, regardless of the schedule (in the same way that a seamstress might try to work faster if the rate of pay per shirt declined; the faster the seamstress works, the more shirts she completes and she is able to maintain her previous rate of pay). In the laboratory, pigeons would peck at a key 50 times or more to get one reinforcer! Eventually, however, as fixed ratios increase, behaviour slows, and may even cease. Essentially, if the ratio is too high, the organism reacts as if it were an extinction situation (for example, if it takes 150 responses to produce one reinforcer, it may seem as though reinforcement has ceased!).

Interestingly, when the ratio is made less predictable, performance tends to be better. A **variable ratio schedule** is defined in terms of the *average* number of responses required to receive a reinforcer (for example, VR 10 means on average every tenth response is reinforced) (see Figure 3.10b). This means that the ratio is predictable in the long run, but in the short run the number of responses required varies. This is much like a slot machine, which pays out on a predefined percentage of plays – but one cannot predict precisely when the next pay-out will be. This analogy actually works surprisingly well, for animals on variable ratio schedules tend to perform very steadily, much like the gambler who plays a slot machine for hours, hoping that the next play will be the big pay-out. In both cases, the fact that reinforcement occasionally occurs after only a few responses tends to maintain the behaviour over the sequences when no reinforcement occurs. Many other activities also seem to be based on variable ratio reinforcement. In sales jobs, for example, a sales agent might have to make a varying number of client contacts before closing a sale; even so, the occasional sale sustains the behaviour. Sport fishermen typically don't succeed on every outing, and cannot predict when they will catch a 'big one'; even so, they continue to try, knowing the next outing may lead to success. Thus, while both variable and fixed ratios tend to produce very steady responding, variable schedules are slightly better overall.

variable ratio schedule a reinforcement contingency defined in terms of the average number of responses required to receive a reinforcer; thus, VR 8 means that on average every eighth response is reinforced.

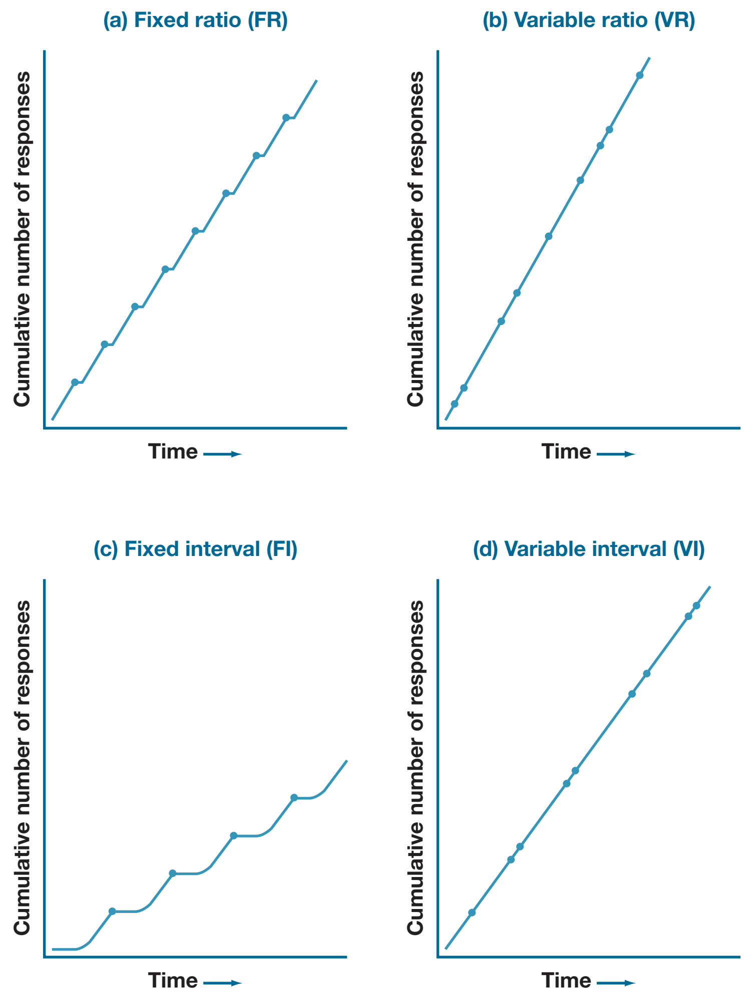

Figure 3.10 Schedules of Reinforcement Using cumulative frequency graphs (where total responses are counted), one can see the differences in response patterns using different types of schedules. Note in particular the lag in responding after a reinforcer (shown by dots) under fixed interval schedules, and the relatively fast response rates (steep curve) under variable ratio.

fixed interval schedule a reinforcement contingency defined by the amount of time that must pass since the previous reinforcer was given, before a response will receive a reinforcer; thus FI 5 min. means a fixed interval of five minutes.

In ratio schedules, since getting a reinforcer depends on making the specified number of responses, performance tends to occur at a relatively high rate. The situation is somewhat different with *interval* schedules, where time is the crucial factor. In a **fixed interval schedule**, a timer determines how long it has been since the previous reinforcer was given, and only a response made *after* the required time interval has elapsed will receive a reinforcer (of course, if no response at all is made, no reinforcer is given). Thus, in a FI 15 sec. schedule, at least 15 seconds must pass between delivery of reinforcers (assuming a response is made at the end of the required interval, see Figure 3.10c).

It may seem odd to consider reinforcement as being time-dependent, but in fact there are many situations where this type of schedule applies. For example, if you are waiting for a bus on a rainy day, you must go out at the appropriate time in order to get the bus. It does not matter how often you go out to check (i.e., make multiple responses) – only going out at the *appropriate* time will get reinforced (by actually getting the bus). Another example is the administering of painkillers in hospitals. In many cases, a patient will only receive medication for pain (the reinforcer) if they request it (the response). However, to avoid overdoses, the medication will only be given after a certain time interval since the previous dose, such as four hours (the fixed interval); any requests made before this time has passed will be ignored. Work which is paid on an hourly rate rather than piecework basis may also be considered a type of interval schedule. In this situation, there is an assumption by employers that responses are being made (i.e., that work is being done), but technically the response which is reinforced is being present at work! (You get paid according to how many hours you work, not the number of things you accomplish.) Thus, fixed interval schedules are actually a fairly common form of reinforcement.

One special characteristic of fixed interval schedules is that they only *require* a single response in order to receive reinforcement, provided the required time interval has elapsed. The result is that fixed interval schedules tend to produce rather low rates of responding compared to ratio schedules. In fact, even in laboratory studies with animals, as the animal becomes familiar with the situation, one finds a distinctive pattern emerges. Typically, few responses are made immediately after a reinforcer is delivered. Instead, there is a lull, and then a few tentative responses, and finally a brief surge in responding clustered near the end of the time interval. Essentially, the organism comes to recognize that premature responses are wasted, and tries to determine the end of the time interval. (While not as precise as a real clock, most species have a biological clock which provides a sense of time.) On a cumulative record of responses, this clustering of responses produces a distinctive 'scallop' in the graph (see Figure 3.10c). An inventive demonstration of this phenomenon is a study by Weisberg and Waldrop (1972), which found that the number of bills passed in sessions of the US Congress showed this FI scallop – the closer it was to the end of a session, the more bills were passed!

If the intention is to encourage steady responding, a fixed interval schedule is inappropriate, because of the clustering of responses which it produces. Given this difficulty, is there any time-based alternative schedule that can remedy this problem? The answer is yes; the trick is to vary the time interval. In a **variable interval schedule**, the schedule is defined in terms of the average time interval required over the long term, much as a variable ratio is defined by an average number of responses required. Thus, on a VI 15 sec. schedule, one interval might be only 5 seconds, the interval after another reinforcer might be 20 seconds, and so on – only over a long period would the average duration be 15 seconds (see Figure 3.10d). From the point of view of the organism being reinforced, this variability means that the availability of a reinforcer is no longer predictable; consequently, the only way to determine if the interval has elapsed is to make a response. The result is that variable interval schedules result in steady behaviour – albeit at much lower rates than ratio schedules. Since the schedule is still time-based, very fast response rates don't accelerate the process of getting a reinforcer; instead, the purpose of responding regularly is simply to check if a reinforcer is available.

variable interval schedule a reinforcement contingency defined by the average time interval which must elapse since the last reinforcer before a response will be reinforced; thus, on a VI 15 sec. schedule, over a long period the average duration would be 15 seconds.

An example which provides an analogy may help clarify the nature of behaviour under variable interval schedules. Imagine someone who works as a quality control inspector on an assembly line; the person's role is to watch items as they pass by, and pick out any which appear defective. In this situation, a person who ignores the task (for example, to read a newspaper) may not be detected for several months, until complaints about defective products are received from customers. Consequently, the foreman checks the person every two hours, on the hour. Since this is equivalent to a fixed interval schedule, a person intent on bunking off could simply begin working just before the foreman arrives, and then return to loafing after the foreman leaves (this is equivalent to a fixed interval scallop). In order to avoid this, the foreman varies the time of his visits – in effect, a variable interval schedule. Now, since the visits are unpredictable, the worker must work steadily, or run the risk of being caught unexpectedly. The result would be typical of behaviour under variable interval schedules – steady responding! (Compared to variable ratio schedules, where there is a direct incentive for fast responding, variable interval schedules tend to produce slow but steady response rates.)

While these four types of schedules are only a small sample of the possible types of partial reinforcement, they show that intermittent reinforcement can be used to sustain behaviour, and that the behaviour tends to reflect the specific requirements of the schedule. In this sense, both people and animals are adept at recognizing the demands of their environment and adjusting their responses to fit the situation. Beyond that, it is worth noting some other characteristics of behaviour using partial reinforcement. The most obvious difference between continuous and partial reinforcement is that under partial reinforcement, an organism does more work to get a reinforcer. This gap increases as the size of the ratio or length of the interval increases. One might expect that this 'more work, less pay' regimen would result in resistance (reduced responding), yet behaviour tends to occur at a *higher* rate with partial reinforcement than with continuous (except with very high fixed ratios, as discussed previously). In addition, variable schedules (ratio *or* interval) tend to produce greater response rates than equivalent fixed schedules. Overall, variable ratio schedules seem to be the most effective in maintaining high rates of behaviour. Skinner, of course, would never speculate about *why* this occurs, since that would require discussing non-observable events. Still, the dynamics may involve the fact that primary reinforcers have survival value, and organisms must meet their survival needs, regardless of the effort involved. Early hunters, for example, likely received only partial reinforcement for their efforts, and still had to persist in hunting when game was scarce. Studies of foraging behaviour have in fact led to the suggestion that animals (and presumably people) look at long-term costs, as well as the immediate consequences, when seeking food (Collier *et al.* 1997).

Since partial reinforcement tends to produce higher rates of responding than continuous reinforcement, what happens under extinction conditions (when no reinforcement at all is given)? Generally speaking, *behaviour acquired under partial reinforcement is much more persistent than behaviour acquired under continuous reinforcement*. The standard measure is to count how many responses are made once reinforcers are no longer available; this is called 'resistance to extinction', where a larger number reflects more persistent behaviour. Interestingly, one finds that resistance to extinction roughly parallels the hierarchy found when reinforcement is available; that is, variable schedules are more resistant to extinction than fixed schedules. One example of this is gambling, which tends to produce very persistent behaviour, in spite of the fact that individuals are assured of losing in the long run (for example, slot machines, roulette and craps all fit this description). Again, Skinner does not speculate as to *why* this is true, but a number of researchers have noted that the change in contingency (from reinforcement to extinction) is simply harder to detect with partial reinforcement. That is, with continuous reinforcement, it is immediately obvious if reinforcement ceases, since a reinforcer should follow every response.

By contrast, with any form of partial reinforcement, the organism has learned (expects) to make responses that go unreinforced; the more variable the schedule, the longer it would take to determine that reinforcement has definitely ceased. Consider an analogy: if you play a slot machine for some time without winning, when would you conclude that the machine was broken or crooked, as opposed to your simply being on a losing streak? In a study with children who used machines that either paid out every time or on a variable ratio, a change to extinction was quickly recognized by the first group, while the second group continued playing for lengthy periods without ever winning (Lewis 1952). Thus, partial reinforcement produces higher rates of responding during reinforcement, and greater persistence during extinction.

Discriminative Stimuli

At this point, it should be clear that contingencies of reinforcement are very diverse, and that individuals seem to be capable of adapting to the requirements of different situations. But how do they *know* what the requirements are in a particular situation? Consider a simple example: 5-year-old Manny likes to eat spaghetti with his fingers. His mother dislikes this behaviour, and slaps his hand when she catches him eating that way. By contrast, his father is relatively indifferent to the behaviour, neither scolding nor praising it. Thus, depending upon who is present, the consequences of Manny's behaviour differ significantly. Obviously, Manny is likely to make the connection between who is present and what consequence occurs, and adapt his behaviour accordingly – eating with his fork when his mother is present, but using his fingers when his father is present! In this situation, the parents are **discriminative stimuli** – stimuli which signal the contingency of reinforcement available. In the above case, Manny's mother is a stimulus which signals *punishment* for eating with his fingers, while Manny's father signals *positive reinforcement* (i.e., Manny presumably eats with fingers because it is reinforcing). In general, discriminative stimuli arise when elements of the environment are associated with a particular contingency.

discriminative stimulus a stimulus which signals the contingency of reinforcement available.

In principle, it is possible for any contingency to become associated with a discriminative stimulus. Researchers have found, for example, that rats will use discriminative stimuli to decide when to press a bar – pressing it when the contingency is highly reinforcing (continuous reinforcement), and not responding when the contingency is not reinforcing (extinction). In another case, pigeons were trained to discriminate between cubist paintings by Picasso and impressionist paintings by Monet (Watanabe *et al.* 1995)! In everyday life, we all make distinctions based on the perceived contingency in the situation. For example, most people will exceed the speed limit when no police are visible (in which case speeding is not punished), but will immediately slow down when a police car is spotted (a discriminative stimulus that speeding will be punished). Children may react differently with each parent, as in the case of Manny, above. We act differently at a party than at work, because the situation signals that different behaviours will be rewarded in each case. Thus, discriminative stimuli, by indicating the potential consequences of behaviour, tend to influence the responses we make. The behaviour is still operant, not reflex – it is up to the individual what response is produced. For example, some people would rather park a car illegally and risk a ticket, than use a car park where it is certain they must pay a fee. Our capacity to recognize discriminative stimuli, and to modify our response, makes it easier to adapt to a changing environment.

Non-contingent Reinforcement

In all of the situations we have been examining, there has been a clearly identified relationship between a response and a reinforcer, described by the contingency. Depending on the type of contingency, behaviour increases or decreases systematically. But is this really a fair description of what happens in the real world? Is reinforcement always clearly dependent on behaviour, or do consequences sometimes occur randomly? Most people would quickly grant that some events in life are random, at least in terms of our ability to control them through our actions. That means that sometimes reinforcement is also random. For example, if you find money on the pavement, is it likely a result of someone seeking to reinforce you? And if it is actually a chance event, can we say anything about how organisms react to such random consequences?

Skinner considered this question, and described such random consequences in terms of **non-contingent reinforcement**, which means that the presence of the reinforcer is unrelated to the occurence of the response. Using pigeons as subjects, Skinner did some inventive studies of the issue. In the typical situation, a pigeon that was already familiar with a Skinner box apparatus would be placed in a chamber, and a timer would provide a food pellet every 15 seconds, regardless of what the pigeon did. (Note that since *no* responses are required, this is *not* the same as a fixed interval schedule; in the extreme case, if the pigeon went to sleep, the feeder would still keep dropping food pellets!) After a period of time, Skinner and another observer would return to see what was happening. According to his description (Skinner 1948a), six out of eight pigeons had developed elaborate, stereotyped response sequences. Since these behaviours actually had no effect on the availability of reinforcement, Skinner called such behaviours 'superstitious'. All that was happening, he argued, was that responses were reinforced *by coincidence*, and then the organism maintained the response that was reinforced.

non-contingent reinforcement a situation where reinforcers sometimes occur independently of any specific response; chance forms of reinforcement.

Skinner went on to suggest that non-contingent reinforcement has similar effects on people. Superstitious behaviours seem to arise in situations where behaviour is only inconsistently reinforced, and where the behaviour has no real influence on the outcome (note that this is *not* the same as partial reinforcement, where the behaviour *does* influence the outcome). For example, many people have particular rituals for trying to make elevator doors close – tapping the edge of the door, holding the 'close' button, even jumping to momentarily increase the load on the elevator! Unfortunately, most modern elevators operate on a programmed cycle, and so people are really acting like Skinner's pigeons, engaging in a ritual which makes the time go by. Similarly, many sports rituals, like having a lucky shirt, seem to be based on superstitious behaviour.

The idea that superstitious behaviour can develop due to non-contingent reinforcement has been demonstrated experimentally – though, interestingly, not every participant developed a superstitious ritual (Ono 1987). Also, it may be that people are more likely to develop superstitions when trying to avoid negative reinforcers than when trying to obtain desirable outcomes

Try it Yourself

Do you feel uneasy if you break a mirror, half-expecting to have seven years of bad luck? Do you avoid walking under ladders? Do you always take your 'lucky pen' with you when you write a test or examination? Can you think of an instance of superstitious behaviour in everyday life (for example, your own behaviour, or someone you know)? Can you identify the reinforcer that seems related to the behaviour? Can you tell whether it is contingent or non-contingent? Is it possible that performing the superstitious behaviour makes you feel better and more confident, and that this is what actually contributes to a positive outcome? Does that seem to support or contradict Skinner's interpretation of superstitious behaviour?

(Aeschleman *et al.* 2003). That is, people are more superstitious about trying to 'ward off bad luck' than about trying to 'attract good luck'. At the very least, non-contingent reinforcement shows how sensitive organisms are to environmental consequences.

APPLICATIONS AND IMPLICATIONS OF CONDITIONING

One of the striking things about behaviourism is the strong pragmatic element which underlies it. Behaviourists are typically very interested not only in trying to understand behaviour, but in *applying* their understanding in the real world. Watson's claim about raising children, cited at the start of this chapter, is one example of this impulse. Skinner was often outspoken concerning his ideas for reshaping society, including writing a utopian novel, *Walden Two* (Skinner 1948b). In more limited ways, behaviourist methods have been applied to many aspects of human behaviour; these applications are commonly referred to as **behaviour modification**. While more specifics about the application of conditioning principles to therapy will be given in Chapter 9, it is appropriate to consider here some general issues related to conditioning.

behaviour modification the application of conditioning techniques to altering human behaviour, particularly those behaviours identified as abnormal.

Negative Reinforcers and the Aversive Control of Operant Behaviour

As Thorndike noted, not all behavioural outcomes are alike. While positive reinforcers like praise or money are welcomed by individuals, negative reinforcers like criticism or physical punishment are unpleasant. Therefore, positive reinforcement is more appealing as a means of modifying behaviour than are punishment and negative reinforcement, which depend on the use of negative reinforcers. Although Skinner would shun descriptions like 'positive reinforcement is more appealing' as being vague and subjective, researchers have found a number of ways to examine the differences between the use of positive and negative reinforcers.

The use of negative reinforcers is often referred to as *aversive control of behaviour*, because of the way organisms react to negative reinforcers. As noted previously, there are two ways in which aversive control is used: *punishment* is used to reduce the frequency of a response, whereas *negative reinforcement* is used to increase a response (that is, the response leads to escaping or avoiding the negative reinforcer).

Punishment is probably the most used, and misunderstood, method of dealing with undesirable behaviour. Parents resort to scolding when disciplining their children. Employers will criticize or threaten hapless employees. Even our legal system is based on punishment for breaking society's rules. Unfortunately, punishment has several limitations, which are worth noting.

First, as with *any* operant reinforcement, the use of punishment depends on **contiguity** between the response and the reinforcer – for effective learning, the reinforcer must immediately follow the response. Punishment which is delayed will be ineffective in controlling the response, or worse, may become associated with a different response. For example, 5-year-old Sarah breaks a plate, and then later tells her mother what happened. The mother, upset about the broken plate, scolds Sarah – but in fact this negative reinforcer will tend to be associated with *telling* about the accident, *not* breaking the plate. In the future, Sarah may still break things, but may not be forthcoming about admitting it! In the same

contiguity the principle that a reinforcer must occur immediately after a response in order for learning to occur.

way, imagine if your dog soiled the carpet, and two hours later, when you discovered the damage, you yelled at the dog, who was now quietly resting on his blanket. Ignoring the importance of contiguity can seriously hamper the effectiveness of punishment. (Similarly, one might note that our legal system, with its typically long delays between crime and punishment, is also poorly designed in terms of contiguity.)

A second limitation of punishment is that it tends to encourage avoidance behaviours. A child in school may be scolded by a teacher for giving a wrong answer. Obviously, the teacher's intention is to get the child to study harder, and thereby give more correct answers, but the real result may be rather different. The child may refuse to answer at all, or even skip the classes taught by that teacher. Because of classical conditioning, any stimulus associated with a negative reinforcer (the teacher, the classroom) may itself become aversive, and trigger avoidance. The avoidance responses may be directed at the situation, or simply at finding ways to avoid the negative reinforcer itself. A criminal with a history of robbing banks was once confronted by a prison official, who asked, 'After all these years, don't you know robbing banks is wrong?' The criminal's response was, 'Actually, I don't see anything wrong with robbing banks; it's getting caught that I don't like!' Research by criminologists has indicated that most people obey laws because they feel it is the right thing to do (feeling virtuous can be a form of positive reinforcement); it is only a minority of people (mostly criminals) who focus on the punishments for transgressing. As anyone aware of crime statistics knows, having laws which specify punishments for criminal acts does not in itself deter criminal behaviour, since those so inclined will simply seek to avoid getting caught.

This issue of punishment encouraging avoidance is a major concern, because punishment tends to *suppress* behaviour, not *extinguish* it. Any operant response occurs because there is some reinforcer supporting it; applying punishment to suppress the response simply pits one reinforcer against the other. Worse, whereas positive reinforcement can be highly effective with only partial reinforcement, *punishment must be continuous in order to suppress behaviour effectively*. Thus, in situations where it is possible to avoid the punishment even occasionally, punishment will not be fully effective in suppressing the undesired response. (Consider what this means in terms of our legal system, where arrest and conviction rates rarely approach 100 per cent.) Instead, it would be better to identify the factor which encourages the undesirable behaviour, and try to alter the environment to eliminate that positive reinforcer (Lerman and Vorndran 2002). For example, children sometimes misbehave because it quickly draws parental attention (a positive reinforcer); to a child starved for attention, the scolding which may follow is less significant than the attention. The result is an increase in disruptive behaviour, followed by more attention (and scolding), etc. The way out of this may be for the parents to *ignore* (extinguish) misbehaviour, and *also* attempt to offer attention when the child is playing quietly or otherwise desirably engaged. (In cases like this, the reinforcing of appropriate behaviour is as important as the extinguishing of undesirable behaviour.)

Because of the way the contingency operates, negative reinforcement is often more effective than punishment as a form of aversive control. Since the focus is on *increasing* a desired response, one does not encounter the problems of suppression associated with punishment. In addition, the acquired fear which can lead to escaping or avoiding the punishment situation, in negative reinforcement tends to *sustain* the desired behaviour. For example, a child may refuse to dress for school, whereupon the parent scowls, and begins yelling at the child, who finally gets dressed. In the future, a scowl may be sufficient to induce the child to dress. Because it is designed to increase a response, not suppress it, negative reinforcement is probably preferable to the use of punishment. In addition, because making the desired response is directly linked to removal of the aversive stimulus, negative reinforcement is more effective than punishment in signalling the desired behaviour (punishment simply indicates what is *not* desired, nothing about what is desired).

Unfortunately, there is one consequence of aversive control which is associated with *any* use of negative reinforcers. By their nature, negative reinforcers represent aversive stimuli, and no organism readily tolerates such events. The use of aversive control therefore tends to promote anxiety, resentment and even aggression, in addition to the problems identified above (Azrin and Holz 1966; Berkowitz 1983). This means that depending on aversive control to regulate behaviour is going to produce a whole range of secondary problems, which may even be less desirable than the original behaviour.

Because of the difficulties outlined above, behaviourists would say that the use of aversive stimuli, in any form, should never be a preferred choice. Despite this, it is clear that punishment is still frequently used, by parents, and by society (Gershoff 2002; Lerman and Vorndran 2002). By contrast, Skinner once suggested that a well-designed society would depend on a combination of positive reinforcement (for desired behaviours) and extinction (for undesired behaviours), not aversive control (Skinner 1948b). It is unlikely that we will ever achieve Skinner's ideal of a punishment-free society – as Vollmer (2002) has noted, aversive outcomes happen, and sometimes they have nothing to do with social interaction (for example, turning the wrong knob on the shower, or slipping on a patch of ice). Nonetheless, better recognition of both the limitations of aversive stimuli and the alternatives to their use is potentially worthwhile.

Try it Yourself

My (MH) adult niece relates the story of being a physically small child of 8, and being tormented one winter by several children who were younger than she, but physically much larger. One day, when returning home from school, the younger children chased her, pelting her with rocks covered with snow. Unable to outrun her tormentors, she threw some snow at their feet to try to keep them at a distance. The next day, she told her teacher about the episode, and was severely reprimanded and punished for having 'thrown snowballs at younger children'. My niece still flushes with resentment when she recounts the story and says, 'Even today, I don't know what it was I was supposed to do!' Given what you know about ways that learning principles can be used to modify behaviour, what do you think she should have done? What should the teacher have done? How can learning principles best be used to modify the behaviour of bullying children?

Interrelationships of Classical and Operant Conditioning

For most of this chapter, we have been discussing classical and operant conditioning as if they were totally separate aspects of behaviour. However, it should not be surprising to find that there are interconnections between the two: after all, organisms are constantly producing *many* responses, both reflex and operant. In this sense, the distinction between the two types of learning is partly a way of simplifying the analysis of behaviour, by breaking it into reflex and operant components. In the real world, both processes can be occurring simultaneously. One striking example of this is *negative reinforcement*. You may recall that negative reinforcement utilizes a negative reinforcer in order to increase the probability of a response. One form of this is escape, where a negative reinforcer is presented, and is only removed after the organism makes the desired response. In this circumstance, the removal of the aversive stimulus is effectively like a reward, so the behaviour becomes more likely (hence, reinforcement). For example, a dog given a mild shock through an electrified floor grid will learn to jump to another chamber to escape the shock. Now, if a light flashes before the start of the shock, the dog will soon anticipate the shock, and jump before the shock begins. This becomes *avoidance* – the dog is jumping in order to avoid the negative reinforcer. This leads to an interesting problem: since the dog jumps before the shock, there is no longer any experience of the original reinforcer – a circumstance that would lead to *extinction* of the response if one were looking at positive reinforcement. So why does the dog keep jumping each time the light goes on? The light, of course, has become a discriminative stimulus, enabling the dog to respond before the shock occurs. Still, why should the dog persist in jumping without at least an occasional experience of shock? The answer seems to be that, through classical conditioning, the light has become a CS associated with the UCS of shock – which is a perfect scenario for creating a conditioned fear. Thus, the dog continues to jump, not to *avoid* the *shock*, but to *escape* from the feared *light*! (Mowrer 1956; Rescorla and Solomon 1967).

Recognizing that the two processes (operant and classical conditioning) are occurring together also adds to our understanding of conditioned fears. Watson, in his demonstration with little Albert, discovered that conditioned fears do not readily extinguish. The reason for this seems to be that the feared stimulus (the CS) *also* triggers operant escape behaviour. This escape response removes the individual from the situation *before* there is an opportunity to determine if the UCS will follow or not – thereby preventing the conditions necessary for extinction. (The same mixture of classical and operant responses happens in the shower when we hear the toilet flush: while we *fear* the sound, we also tend to *jump* away from the water spray to avoid being scalded.) The fact that fear stimuli can evoke an operant response is a very significant point, in terms of those everyday fears, which are called phobias. If, as Watson argued, such fears are based on classical conditioning, then it is also likely that the fears persist long after the original experience, because we *avoid* the situations that *elicit* the fear. As a result, there is no opportunity to find out if our fear is realistic or not. For example, a person who is afraid of flying will be reluctant to fly, and therefore has no chance to find out that flying is safe, and that there is nothing to fear. In essence, until we face the fear situation, there is no opportunity to extinguish the fear response.

Another type of interaction can occur in which conditioned behaviours are also sustained by reinforcement. For example, a phobia may arise through classical conditioning, but the individual may also be positively reinforced by attention and sympathy from other people. In such circumstances, the individual may be unlikely to try to change.

Try it Yourself

If you have a phobia or fear yourself, how do you cope with it? If you tend to avoid the fear-arousing situation, do you think this reaction is adaptive for you in the long run? For example, many students with a fear of public speaking avoid taking courses in which they will be required to give oral presentations. Can you suggest a technique based on learning principles that might help in dealing with such fears?

Autonomic Conditioning and Biofeedback

Consider the following proposal: we will give you $20 if you raise the temperature of your left hand while simultaneously lowering the temperature of your right hand. Short of getting a blanket and ice pack, this may seem like an impossible goal. After all, body temperature is an involuntary (reflex) function. How could we expect to control it with a reinforcer like money? At one time, psychologists would have agreed that such a task was impossible – after all, reflexes are the domain of classical conditioning, not operant conditioning. Even the evidence that shows the two types of learning can occur together (as in avoidance learning) does not challenge the fundamental distinction of reflex vs. operant responses. In fact, however, we now recognize that the boundaries are more ambiguous than the traditional view suggests.

The change was triggered by the work of Neal Miller, a noted researcher in the field of learning, and Leo DiCara, then a graduate student working with Miller. Miller and DiCara wondered if it would be possible to use operant reinforcement with so-called involuntary responses. Although this seemed far-fetched, data on phenomena like meditation suggested that under some circumstances individuals *could* deliberately alter these responses. While the details of the original procedures were rather complex (involving partially paralysed rats, with electrical stimulation of the brain as a positive reinforcer), the implications of the results were quickly apparent: involuntary responses *could* be operantly controlled!

Key Thinker: **Neal Miller**

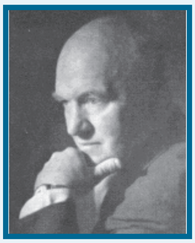

Neal Miller (1909–2002) contributed to many aspects of psychology over a lengthy career. Born in Milwaukee, Wisconsin, the son of a psychologist, Miller got his BS at the University of Washington in 1931, an MS at Stanford a year later, and a PhD in psychology from Yale in 1935. Miller then spent a year as a research fellow at the Institute of Psychoanalysis in Vienna. Though he began as a Freudian, Miller gradually became interested in focusing on more measurable aspects of behaviour. Over the course of his career, he worked on developing selection procedures for aircrew during World War II, collaborated with John Dollard on a behaviourist translation of Freudian concepts, and did pioneering work in learning and neuroscience. Perhaps the most significant of his contributions stemmed from his work on autonomic conditioning. While the breakthrough studies were done with his student Leo DiCara, the foundation of this work went back to Miller's interests in motivation, behaviour and the brain, starting in the 1930s. The work of DiCara and Miller was initially greeted with scepticism by many, but biofeedback procedures are now a standard form of medical treatment. From his early Freudian days, Miller had evolved into a rigorous researcher whose contributions, particularly in learning, were recognized by many awards, including the presidency of the APA in 1961. Miller remained active throughout his career, in later years serving on the APA Board of Scientific Affairs. He died in March 2002, at the age of 92.

autonomic conditioning (also called 'learned operant control of autonomic responses') the conditioning of changes in autonomic (involuntary) responses (such as heart rate or blood pressure) by means of operant reinforcement.

Miller described the process as 'learned operant control of autonomic responses', or **autonomic conditioning** (Miller 1969). By providing reinforcement which was based on changes in autonomic (involuntary) responses, it was possible to alter behaviours such as heart rate, blood pressure – even the temperature of various limbs (by changes in blood flow). To understand what is involved in autonomic conditioning, it is necessary to consider how operant responses normally function. For all voluntary muscle movements, our brain receives information, called *proprioceptive feedback*, about the execution of the movement. It is proprioceptive feedback which tells you the position of your arm even when your eyes are closed, for example. But for involuntary functions (involving the autonomic nervous system), there is little or no proprioceptive feedback. Consequently, there is typically no direct awareness of autonomic responses like blood pressure. To circumvent this limitation, DiCara and Miller used sophisticated equipment to monitor these hidden processes, and thus determine when to deliver a reinforcer. Since most autonomic functions show natural fluctuations (for example, heart rate varies slightly even when sitting), the procedure amounted to a process of shaping a desired response.

biofeedback a general term for applications of the process of autonomic conditioning; the name refers to the fact that in humans reinforcement is based on providing an individual with information (feedback) about physiological processes (bio) which are normally not observable.

Thus, the essential element of autonomic conditioning is the ability to measure the response. While the recording of physiological activity dates back to the 1930s, using such techniques in the context of operant conditioning is much more recent. Today, applications of the process are frequently referred to as **biofeedback**, since the process provides an individual with information (feedback) about physiological processes (bio) which are normally not observable (see Mercer 1986).

Basically, any biofeedback procedure requires equipment to monitor the response of interest, and a means of conveying information to the individual about changes in their response (unlike the original animal studies, in applications with people, informational feedback is often a sufficient reinforcer). For example, if one is interested in muscle relaxation, one would use a device called an EMG (electromyograph), which measures the electrical activity in the motor neurons which control the muscles. If the heart is of interest, one would use an ECG (electrocardiograph), and so on. The means of providing feedback might be a buzzer, or a light which flashes, when the desired response occurs (see Figure 3.11).

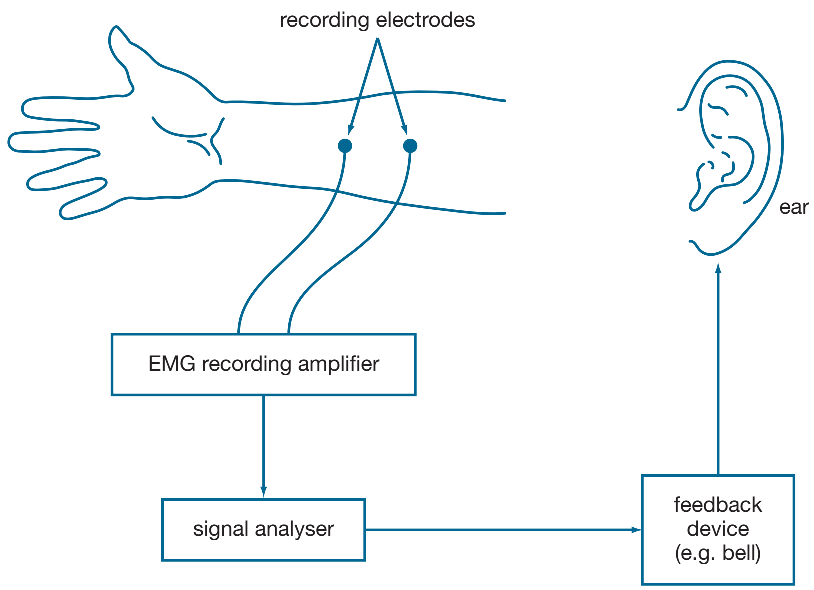

Figure 3.11 Biofeedback Basically the procedure for biofeedback requires some means of monitoring the physiological response (in this case a muscle twitch, recorded using an electromyograph), and a means of making the individual aware of changes that occur (in this case, a bell that rings whenever the target muscle twitches). Using this type of procedure, while monitoring several muscles with several bells, subjects actually can learn to play tunes!

The most significant applications of biofeedback are for medical treatment. It has proved very effective for the relaxation of voluntary muscles (for example, arms, legs, neck), which is often utilized for tension control. In fact, Grazzi (2007) reports that biofeedback treatment of migraine and tension headaches is just as effective as treatment with medication, with the advantage that biofeedback does not entail side effects or complications, as is often the case with medication. Biofeedback has also proved moderately effective for reducing hypertension (high blood pressure), and for regulating the rhythm of the heart, but not very effective for reducing overall heart rate (McGrady 1996). One of the more interesting uses has been to train individuals to control epileptic seizures, by teaching them to produce a brain-wave pattern which seems to inhibit seizure activity (Sterman 1978). Biofeedback has also been successfully used to help people with physiologically-based behavioural disorders, such as voice disorders related to the removal of the larynx. In this situation, biofeedback has been found to be effective in helping people improve the quality of their voice and their voice performance, and thereby the quality of their lives (Maryn *et al*. 2006).

From a more theoretical point of view, autonomic conditioning demonstrates that the processes of classical and operant conditioning are more closely intertwined than was once believed. While it does *not* imply that Pavlov was wrong about the formation of CS–UCS associations, it raises questions about the definition of operant behaviour. Skinner originally defined operants as emitted responses, in contrast to elicited reflexes. This made sense, in view of how such responses can be used to alter one's environment. Now, however, it seems that operant conditioning can be applied to almost any response, *provided that there is a clear contingency of reinforcement* (i.e., connection to consequences). Clearly, this would suggest a widening of the boundaries for operant behaviour. Thus, the study of autonomic conditioning has opened up new areas for operant research. Ironically, at just about the time that research on autonomic conditioning was broadening the boundaries, other research was suggesting new limitations of operant learning.

Before leaving the topic of autonomic conditioning, there is a sad side-note to the original discovery that bears mentioning. The initial study by DiCara and Miller was very complex, as noted. When it was first reported, it was viewed as so remarkable and unexpected that other researchers immediately set out to duplicate the results – and failed. Miller himself eventually tried, and also could not reproduce the original findings. While no one today doubts that

Try it Yourself

While you probably don't have access to a machine that will measure and signal changes in your autonomic arousal, you can still try to control some of your autonomic responses yourself. Lie down comfortably and imagine that you are on the beach or in a meadow. Imagine the sun streaming down, warming your whole body. Imagine that your stomach is becoming very warm with the sunlight. After a few minutes of imagining this, many people feel an actual warmth in their abdominal region. In fact, this is a popular technique in physical relaxation training that is often used for stress management. Try it a few times. Does your stomach feel warm? Do you feel more relaxed? Do you think this effect might be enhanced if a machine told you when your surface body temperature was increasing?

If you know anyone who has had biofeedback, talk to them about their experiences. Would you consider biofeedback in preference to medication if you experienced severe headaches? Why or why not?

autonomic conditioning is a real phenomenon, and Miller maintained that there was no evidence of fraud in DiCara's work, the younger researcher was discredited by the controversy. Today, the original study goes largely uncited, and DiCara has become a forgotten pioneer. In science, reputation can sometimes make or break a career.

Biological Constraints on Learning

When most people read about conditioning processes for the first time, one of the first things to strike them is the apparent artificiality of the experimental situations. After all, what does a ringing bell, or manipulating a lever, have to do with an animal obtaining food in the real world? What can the study of such arbitrary acts tell us about ordinary behaviour? In fact, the situations *are* rather artificial, and deliberately so. Behaviourists would argue that laboratory environments provide a high degree of experimental control, and behaviour is still behaviour, in the lab or out. Consequently, many would see the success of their methods as demonstrating just how powerful the basic principles are. However, as with all research, there are inevitably assumptions made when one generalizes from the laboratory to the real world.

equipotentiality premise an assumption made by some behaviourists which states that the principles of conditioning should apply equally to all behaviour, in any species.

One of the assumptions commonly made is called the **equipotentiality premise**. Essentially, this premise says that the principles of conditioning should apply to any response, and any species. (Interestingly, Skinner himself never endorsed such generalizations. He might be willing to generalize from observations of one pigeon to another, but would hesitate to generalize from one species to another.) For many years, studies of a wide variety of species in the laboratory seemed to demonstrate that the equipotentiality premise was correct.

ethology the study of the behaviour of animals in their natural environments.

species-specific behaviour behaviours which are characteristic of all members of a particular species; these response patterns (sometimes popularly called 'instincts') apply to behaviours such as mating, finding food, defence and raising offspring.

However, quite independently of behaviourism, other researchers, coming from a tradition of biology rather than psychology, have studied the behaviour of animals in their natural environments. This approach is called **ethology**. The founder of ethology is often regarded as Konrad Lorenz, a German researcher who did pioneering studies of species ranging from fish to wolves (Lorenz 1967). Ethologists like Lorenz have tended to study **species-specific behaviours** – behaviours which are characteristic of all members of a particular species. These response patterns (sometimes popularly called 'instincts') apply to behaviours such as mating, finding food, defence and raising offspring. Typically, the behaviours seem to be genetically shaped, but also responsive to environmental demands. For example, the young of many species identify their parents through a process called *imprinting*, whereby they attach themselves shortly after birth to the nearest moving stimulus. (Lorenz at one point had a group of ducklings who followed him around as though he were their mother!) Unlike simple reflexes, species-specific behaviour can involve complex sequences of responses, such as the ritual fighting in some species of tropical fish.

For many years, ethologists and behaviourists pursued their interests separately, but in recent years there has been increasing dialogue between the approaches. One of the prime areas of interest has been the interaction of hereditary and environmental influences on learning. The ethologists have tended to assume that much of behaviour is governed by the genetic make-up of a species, while many behaviourists have viewed behaviour as completely malleable, based on the principles of conditioning. As with many such issues, the truth seems to be something in between: environmental circumstances affect the expression of many species-specific behaviours, and biological constraints limit the process of learning. Contrary to the equipotentiality premise,

critical period in development, the concept that there are optimal periods for the learning of certain behaviours.

not all learning is the same. Some behaviours are easier to acquire than others (for example, learning to clap your hands is easier than rubbing your stomach and tapping your head simultaneously). Also, some behaviours appear to be learned best at particular times during development, called **critical periods**. For instance, a child who is not exposed to language prior to the age of 6 will generally have great difficulty learning to speak later.

One way to make sense of these variations in ease of learning is the preparedness dimension developed by Martin Seligman (Seligman 1970). **Preparedness** refers to the degree to which physiological structure influences the occurrence of behaviour (see Figure 3.12). Some behaviours seem to develop with little or no specific experience required; and Seligman refers to these as *prepared* behaviours, because the organism seems physiologically structured to produce the behaviour. Species-specific behaviours would belong to this category; the pouncing behaviour of cats is an example – cats don't have to be taught to pounce! Dogs, by contrast, do not have a preprogrammed pounce response; instead, pouncing must be learned. In the same way, most human behaviour does not stem from 'pre-wired' origins. For example, while we are capable of the balance and coordination required to ride a bicycle, the specific responses involved are not inborn, but must be learned. Seligman refers to behaviours which must be acquired through experience as *unprepared*, since there is no hereditary predisposition involved. At the other extreme, there are some types of complex patterns of behaviour that we find very difficult, if not impossible, to acquire. In these cases, it seems that the physiological structure is not intended to cope with these situations. Seligman says that we are *contraprepared* to acquire such behaviour patterns. For instance, cats are prepared to lick themselves after eating, but they are contraprepared to use licking as an operant response to obtain food. Their physiology is structured for a 'food, then lick' sequence, not 'lick, then food'. (Rats, which have no 'food, then lick' predisposition, *can* use licking as an operant response for food.)

preparedness a concept developed by Martin Seligman to describe the degree to which physiological structure influences the occurrence of behaviour.

Traditionally, behaviourists focused on those response patterns which depend on learning (that is, unprepared behaviours). In doing so, they developed the view that all behaviour was alike, as expressed by the equipotentiality premise. As time went on, however, evidence arose that called this concept into question. While unlearned species-specific behaviours are clearly prepared, there are other behaviours which are learned so quickly as to seem prepared as well. One example concerns food-avoidance learning: Food preferences are well known in both people and animals – but what leads to rejecting a particular food? (See Box 3.1 for why we prefer certain foods.)

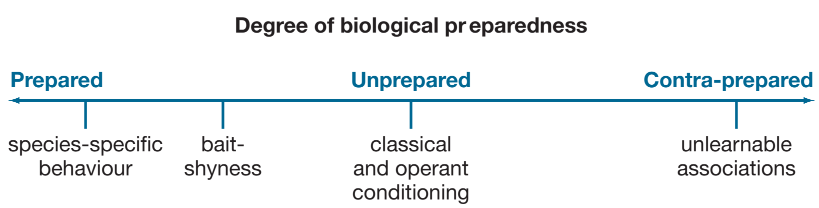

Figure 3.12 Seligman's Preparedness Continuum Martin Seligman has suggested that not all responses are the same, and that various behaviours can be understood in terms of how physiological factors prepare us to learn easily (prepared) or not at all (contraprepared).

Box 3.1 Taste Preferences: Experience and Culture

People in all cultures learn to salivate in anticipation of tasty food, but what do they find tasty? Does experience determine what appeals, or are there genetic differences? Certain aspects of taste, such as our capacity to detect differences in sweetness and saltiness, appear to be innate (Laing *et al*. 1993). However, how much sweetness or saltiness we *prefer* in our food seems to be the result of past experiences with food. For example, foods in both Syrian and New Orleans Cajun cooking tend to be highly salted. It is no wonder, then, that people who grew up with these types of food find Swedish cooking (which uses little salt) to be bland and tasteless (Beauchamp 1987).

As Garcia's work on bait-shyness demonstrated, there is an adaptive value to developing an aversion to the taste of something which makes you sick. But why do people show differences in what they prefer? Research on the origins of taste preferences suggests that they are based on traditional classical and operant conditioning, not biological predispositions. Infants, for example, will make active food choices based on what they like and dislike; in operant terms, different tastes differ in their reinforcing value (Mennella and Beauchamp 1997). Interestingly, however, these preferences seem to be influenced by exposure to flavours in the mother's breast milk, and even in the amniotic fluid prenatally – in other words, the infant comes to like the flavours of food the mother eats! Essentially, a process of classical conditioning is going on, whereby the flavours are associated with the nourishment of the food, but also with positive emotions elicited by contact with the mother during feeding. (Studies of contact comfort would imply that similar emotional conditioning would occur with bottle feeding – but, of course, the caregiver's diet would not affect the flavour of the milk!)

The effects of reinforcement on taste preferences are clear in many cases, too. As discussed in the chapter, food is an important reinforcer, and caregivers not only use food as a positive reinforcer (for example, a sweet as a reward), but also will typically encourage children to eat less attractive foods (for example, vegetables) by using praise, encouragement and other reinforcers. (A friend recalls as a child not being allowed to leave the table until his plate was empty!) Eating is also a social activity, and most people grow up associating certain foods with loved ones (for example, Grandma's special soup) and special occasions (for example, holiday desserts). These foods then tend to elicit emotions of well-being and happiness, due to classical conditioning. This is why we regard some foods as 'comfort foods'.

While food preferences depend on the individual's experience, culture also plays a role, producing commonalities within a culture and differences between cultures. Thus, in Toronto, one restaurant lists the heading of 'Comfort Foods' on its menu, with offerings such as macaroni and cheese, and chicken noodle soup. In Newfoundland, cod tongues might evoke similar appeal. In China, monkey brains were traditionally considered a delicacy, while on the prairies of Canada some people would feel the same way about the testicles of a bull. Whether one prefers potatoes, pasta, couscous or rice with meals likely reflects one's cultural background. In the end, what we choose to eat may not determine who we are, but it says a lot about our past experience!

The answer to the puzzle of food aversions came while a psychologist named John Garcia was studying the effects of exposure to X-rays. Among the effects of large doses of X-rays was that animals became sick to their stomachs several hours later; if they had eaten earlier, they would subsequently avoid whatever the food had been. Garcia became interested in this behaviour, in much the way that Pavlov moved from the study of digestion to the exploration of classical conditioning. In a series of insightful studies, Garcia and his colleagues demonstrated that if a rat gets ill after eating a distinctive-tasting food, it will avoid that food in the future (Garcia *et al*. 1974). Garcia recognized that while it was the X-rays (which the rats could not directly detect) that produced the sickness, the rats associated it with the *food* instead. Garcia called this behaviour *bait-shyness* (based on fishermen's belief that a fish nearly hooked on a particular bait won't strike it again).

To understand the bait-shyness effect, imagine a person who, while coming down with the stomach flu, happens to go out for dinner and eats something out of the ordinary (for example, curried chicken). Later, they feel sick because of the flu – but the next time they order curried chicken, the reaction of nausea reoccurs. Subsequent research, as well as anecdotal evidence, has confirmed that this phenomenon is genuine. For example, cancer patients often develop a wide range of food aversions as a result of the nausea produced by chemotherapy (Bernstein 1991).

The exploration of bait-shyness led to the conclusion that many species, including man, have developed in such a way that getting sick is very readily associated with the last food eaten (Logue 1988). The link between food and sickness often occurs after a single experience, despite the long delay between eating and getting sick. This long delay violates the basic behaviourist principle of *contiguity* – imagine trying to teach a rat to press a bar by giving it a positive reinforcer several hours after it makes the correct response! Garcia suggested that this capacity to learn even with long delays between eating and the consequence evolved because it was adaptive: an animal which ate something harmful would do well to avoid it in the future. Through mutation and natural selection, a neural mechanism was created to link taste with stomach upset. Thus, bait-shyness represents a form of prepared behaviour.

Over time, research has shown that food-avoidance learning is only one example of such prepared behaviours. For example, migratory birds are biologically predisposed to learn landmarks on their route (Shettleworth 1972). It has also been suggested that some types of human fears are more easily conditioned than others. Consider fear of the dark: humans are basically daytime creatures. To our ancestors, who depended (as we still do) on vision more than smell or hearing, the night world of darkness was a place of invisible dangers. Consequently, natural selection may have 'prepared' us to be afraid of the dark. (Note this does not mean that *all* people automatically have a strong fear of the dark; it simply implies that very little experience is needed to *develop* such a fear.)

Studies of classically-conditioned fears suggest that not all stimuli are equally likely to elicit conditioned fear responses. A few years after Watson's study of little Albert, another researcher found that an infant could be conditioned to fear a rat, but not wooden blocks or pieces of cloth (Bregman 1934). More recently, studies with adults have shown that fear is more easily conditioned to pictures of snakes or spiders than to pictures of flowers or houses (Ohman 1986; Ohman and Soares 1998). One way of interpreting such differences in the frequency of occurrence of different phobias is to assume that some fears (like snakes and spiders) are biologically prepared. This would make sense in evolutionary terms, since being fearful of creatures which are potentially poisonous or otherwise dangerous can be seen as adaptive. While other factors (such as cultural influences) could also be involved, evidence that other primates are similarly predisposed to fear snakes supports the evolutionary interpretation (Ohman and Mineka 2003).

It is harder to evaluate the possibility of there being *contraprepared* behaviours in humans. One of the difficulties is that, by definition, what is contraprepared is unlearnable. So, if there is some behavioural pattern that we cannot learn, would we even be able to recognize it? Thus far, no one has identified a clear example of contrapreparedness in people, despite many known examples in other species (such as the lick-then-food sequence in cats). Since there is no logical necessity that states that there must be such behaviours in humans, the absence of examples tells us little. So, at present, the issue is unresolved.

Overall, there seems little doubt that genetic and physiological factors play a role in human behaviour. However, such factors do not seem to play as significant a role in human activity as they do in many other species. This is indicated in part by the long infancy/childhood of human development; whereas many animals may be fully developed at birth, human infants require care and assistance longer than any other species. While this is a disadvantage in terms of survival, it allows for maximum flexibility of behaviour, based on experience (Bjorklund 1997). Learning takes time, and with time, the helplessness of the infant becomes the diverse and complex behaviour of the adult.

The study of biological constraints, by both ethologists and behaviourists, has enriched our understanding of behaviour, even as it has limited the range of application of conditioning principles. One of the significant lessons seems to be a recognition of the limitations of laboratory research. Whenever one enters the laboratory to study behaviour, one trades the advantages of control for the disadvantages of an artificial situation. It is well recognized that people often react differently when they know they are being studied (see Chapter 1), and even with animals, the laboratory setting may give a distorted perspective. This is not to say that the years of research on operant and classical conditioning are invalid; indeed, both frameworks have added a great deal to our understanding. However, just as each approach to psychology has its limits, so too it seems that a full understanding of behaviour cannot come from the laboratory alone. By exploring biological constraints, it can be argued that behaviourists have enriched their approach, rather than weakened it.

Try it Yourself

Have you ever developed a sudden aversion to a particular food? Do you recall the circumstances? Does Garcia's work on bait-shyness help you to understand your own taste preferences? In what ways?

CONCLUSION

The behaviourist approach is rooted in the assumption that science must be based on the study of observable events. In terms of behaviour, this means looking at the interactions between an organism and its environment. In adopting this stance, behaviourists forgo attempts to study consciousness and internal subjective states. As Skinner has pointed out, behaviourism is a method of analysis rather than simply a theory (Skinner 1987). Critics say that treating the organism like a 'black box' means that one ignores the mental processes that are central to human behaviour. Skinner says that such events are scientifically unknowable, and in any case, do not *cause* behaviour: thinking about something before doing it is simply correlated with the observable behaviour. For example, if a Freudian theorist suggests that adult behaviour can best be understood by looking at childhood experiences, Skinner agrees – but suggests that the connections are based on the reinforcement history of the person, not some vague concept of 'conflicts between id and ego'.

Ultimately, the best criterion for judging any approach, including behaviourism, is not our theoretical preference, but the extent to which it helps us to make sense of behaviour. While many have criticized the restrictions of behaviourism, the reality is that the study of classical and operant conditioning has added to our overall understanding in psychology. The appeal of behaviourism is reflected in the fact that for many years it was the dominant force in North American psychology. It is interesting, and in some ways ironic, that behaviourism has influenced the attitudes and methods of many psychologists, and has even contributed to the success of the cognitive approach. For example, Edward Tolman, regarded as one of the founders of the cognitive approach, considered himself a behaviourist – though not a radical behaviourist like Skinner. In addition, the study of many cognitive issues, such as observational learning (imitation) and the use of hypotheses in problem solving, began with similar behaviourist studies of animals. Where the introspectionists failed in their attempts to make sense out of mental processes, the behaviourists have pointed the way to new possibilities for a scientific psychology of the mind. More broadly, while radical behaviourism is no longer very influential, there is an active interest in exploring environmental influences on behaviour – what is sometimes called 'behavioural' psychology.

Although the behaviourist approach has contributed significantly to our understanding of behaviour and has led to some highly effective therapies for some behavioural disorders, it no longer occupies the pre-eminent position it once did within psychology. In part, this reflects changes in the discipline, and in part, the limitations of the approach. One major weakness is that research by ethologists and others has shown that the principles of conditioning are not as universal as was once asserted. This failure of the 'equipotentiality premise' restricts the generality of behaviourist principles in important ways. Beyond that, interest in mental processes has not diminished simply because the behaviourists have refused to address the issues. Instead, researchers have found new ways to study mental processes, resulting in new interest in the cognitive approach – as we shall see in Chapter 4.

Putting It All Together – Behaviourist Approach

There are several ways in which the behaviourist approach can shed light on Sam's behaviour. First, it seems likely that he had developed a conditioned fear to writing tests. Given that Sam was always a student who had to work hard to attain good marks, through classical conditioning, he probably came to associate schoolwork with anxiety. While the case study does not specify the circumstances, it is likely that on one or more previous occasions, Sam did badly on tests. The poor grade was likely associated with scolding or disapproval from his parents, or embarrassment from friends (UCS). The fear and anxiety produced (UCR) by these stimuli then became associated with the test situation, so that now a test is a CS for anxiety (CR). (Note that this process may reflect higher order conditioning, in that the anxiety/fear produced by criticism was likely itself acquired through prior conditioning.) Similarly, Sam may have had experiences in which bright, independent girls rejected him (UCS). The anxiety of embarrassment he felt (UCR) became associated with bright, independent girls, so now, even though he is attracted to Vanessa, she is a CS for anxiety (CR).

Sam's procrastination represents a form of operant conditioning. Because his schoolwork is highly aversive (conditioned negative reinforcer), he avoids it by procrastinating. Since procrastinating produces removal of the aversive stimulus, this is negative reinforcement, so he continues to avoid doing his schoolwork. At the same time, socializing with his friends, doing household chores and watching TV have some positive reinforcement value, even though this may be limited by the guilt he feels (a negative reinforcer) while engaging in these pastimes. Note that the reinforcement for procrastinating is immediate, as opposed to the delayed reinforcement he would get from completing his schoolwork (for example, passing his courses), so it has greater strength. The books in this situation represent a discriminative stimulus – as noted in the chapter, behaviour under a contingency of negative reinforcement usually begins as escape (for example, from the anxiety associated with poor school results) and progresses to avoidance of the discriminative stimulus (the books) which signals that negative reinforcer. (It is also worth noting that in this situation, Sam's behaviour is being affected by several types of reinforcers and contingencies of reinforcement. In the real world, behaviour is seldom as simple as in laboratory situations.)

Despite his obvious struggle with anxiety and procrastination, Sam has managed to pass his courses in the past, so he has received some positive reinforcement for being in school. Because he does not seem to enjoy the process of studying, it appears that the main reinforcers for his school behaviour are his grades. Given that his academic success does not appear to have been consistent, and the reinforcers (passing grades) only intermittent, this represents a form of VR schedule. Thus, despite the limitations of his last-minute techniques, he is being positive reinforced for continuing this pattern of behaviour. (As noted in the chapter, a VR schedule is highly resistant to extinction.)

Sam seems to generalize his behaviour to all his school subjects, but the chances are he shows discrimination in his procrastinating by recognizing that not studying just before work is due will result in an immediate negative reinforcer (a failing grade) – that is, the due date is a discriminative stimulus for a change in contingency.

Vanessa is presently a friend of Sam's, so she provides a positive reinforcer for him, as he does for her. He is not sure, however, whether this positive response of hers would generalize to a more romantic relationship, or whether she would discriminate between Sam the friend and Sam the lover. He fears that Sam the lover would receive negative reinforcers from her, so he shows discrimination in his behaviour – he demonstrates only the behaviour guaranteed to bring him a positive reinforcer and does not ask her on a date. That is, he avoids the possibility of anxiety through rejection.

Behaviourists would say that Sam, like all individuals, has a complex history of past learning. But to deal with his anxiety and both his academic procrastination and his avoidance of a romantic overture to Vanessa, it is more important to focus on changing his current behaviour by means of behaviour modification.

CHAPTER SUMMARY

- The behaviourist approach emphasizes the study of *observable responses*, and rejects attempts to study internal processes like thinking.
- In doing so, behaviourists focus on *learning* as the primary factor in explaining changes in behaviour. Depending on the type of response, this involves either *classical conditioning* or *operant conditioning*.
- Classical conditioning is concerned with how *conditioned stimuli* come to elicit *conditioned responses* – reflex responses which are normally elicited by *unconditioned stimuli.*
- Classical conditioning can be applied to a number of aspects of human behaviour, including *emotional responses* like fears, and even activity of the *immune system*.
- Operant conditioning is concerned with how the probability of a *voluntary* 'operant' *response* changes as a function of the environmental consequences (*reinforcer*) which follow the response.
- This process of *reinforcement* can be analysed in terms of the *type of reinforcer*, the *contingency of reinforcement* and the *schedule of reinforcement*.

- The application of operant conditioning to everyday behaviour is commonly called *behaviour modification*; related research includes the effects of *aversive control* and methods of altering behaviour by *biofeedback*, among other uses.
- Recent research has indicated that while conceptually distinct, classical and operant conditioning are interrelated in actual behaviour. In addition, research on *biological constraints on learning* has suggested that there are limits to the generality of conditioning principles, as illustrated by the concept of *preparedness*.

Key terms and concepts

learning
law of effect
classical conditioning
unconditioned response
unconditioned stimulus
conditioned stimulus
conditioned response
stimulus generalization
stimulus discrimination
extinction (classical)
spontaneous recovery
higher order conditioning
conditioned emotional response
operant conditioning
reinforcer
reinforcement
primary reinforcer
conditioned reinforcer
positive reinforcer
negative reinforcer
positive reinforcement
punishment
negative reinforcement
omission
shaping
extinction (operant)
continuous reinforcement
partial reinforcement
schedules of reinforcement
interval schedules
ratio schedules
discriminative stimulus
behaviour modification
aversive control
biofeedback
species-specific behaviour
preparedness

Online Learning Centre

When you have read this chapter, log onto the Online Learning Centre website at ***www.openup.co.uk/glassman*** where you will find answers to these Test Yourself questions and suggested answers to the Try it Yourself activities, plus amny more learning resources to help you study psychology.

Test yourself questions

1 What is the meaning of the following terms?
unconditioned stimulus
conditioned stimulus
unconditioned response
conditioned response
conditioned emotional response
2 What are the differences and similarities between classical and operant conditioning?
3 Why is punishment not recommended in reducing misbehaviour in children?
4 How does biology constrain learning?

SUGGESTIONS FOR FURTHER READING

- **For the reader interested in a more detailed discussion of the principles of learning** (both classical and operant), Chance's *Learning and Behavior* (2008) is a very readable account. For an assessment of the impact of behaviourism, see Kunkel's 1996 article.
- **B. F. Skinner has maintained a distinctive position within behaviourism, not least for his outspoken comments on changing society. One of his clearest presentations of his views on society is** *Beyond Freedom and Dignity* (1971).
- **One significant influence of behaviourism has been the practical application of conditioning principles to everyday behaviour. Martin and Pear's** *Behavior Modification: What It Is and How to Do It* (2006) provides a good overview of such applications. For a more specific discussion of the use of biofeedback techniques in clinical applications, see *Biofeedback: A Practitioner's Guide* (2005), edited by Schwartz and Andrasik.
- **For an account which shows how the ethological approach differs from laboratory studies of behaviour, Jane Goodall's** *Through a Window* (1990) provides a highly readable beginning.