



CHAPTER OUTLINE

LEARNING OBJECTIVES

INTRODUCTION

CLASSICAL CONDITIONING

- Pavlov's dogs
- Other examples of conditioned responses
- Associative analysis
- The importance of classical conditioning

INSTRUMENTAL LEARNING

- Thorndike's cats
- The Skinner box
- The law of effect
- Control of performance
- The importance of instrumental learning

THE PRINCIPLES OF ASSOCIATION FORMATION

- Contiguity and predictiveness
- Selective association formation

NON-ASSOCIATIVE LEARNING

- Responding to a single stimulus
- Spatial learning
- Discrimination learning

FINAL THOUGHTS

SUMMARY

REVISION QUESTIONS

FURTHER READING

Learning Objectives

By the end of this chapter you should appreciate that:

- classical conditioning involves fundamental learning processes;
- the classical conditioning procedure forms certain kinds of specific associations;
- classical conditioning plays a role in psychological phenomena observed outside the laboratory;
- instrumental training procedures engage learning mechanisms to form particular associations;
- there are important factors that determine when instrumental learning will be expressed in behaviour;
- classical conditioning and instrumental learning depend on association formation;
- the associative principle does not apply to some forms of learning.

INTRODUCTION

We all know what 'learning' means. As we develop, we learn new motor skills, such as playing the piano or riding a bike. We acquire new cognitive skills, such as long division or computer programming. And we might learn a body of information, such as the dates of the kings of England or the words of a song.

In psychology, the term 'learning' covers all this and more. A wider, psychological definition might go something like this: 'Learning is the process whereby an organism interacts with its environment and becomes changed by the experience so that its subsequent behaviour is modified.' The acquisition of new information and new skills falls within this definition, but so do the following events:

- A snail experiences a brief jolt of the surface on which it is crawling and reacts by retracting its tentacles. Subsequent jolts, however, are found to be less effective in inducing withdrawal until the reaction finally disappears.
- The first conspicuous moving object seen by a newly hatched chick is a laboratory attendant. As a consequence, the chick develops an attachment to that person, approaching and following him or her, and tending to avoid other things.
- A rat is given access to a distinctively flavoured foodstuff that has been laced with a small amount of poison, enough to induce nausea but not enough to kill. On recovering from its illness, the rat will tend to shun the flavour, even if it is one that it liked beforehand.



Figure 4.1

As we develop, we learn new motor skills, such as riding a bike.

- A hungry pigeon is given a small amount of food each time it happens to make a turn in a particular direction. After experiencing a few rewards, the bird develops an increasing tendency to circle on the spot in the 'correct' direction.

These examples are all interesting in themselves, but that is not enough to explain why many psychologists should have chosen to concentrate on experimental studies showing how special training procedures conducted in rather constrained and artificial circumstances can produce changes in the behaviour of laboratory animals. The reason lies in their hope that, by focusing their attention on relatively simple examples of learning that are amenable to experimental and theoretical analysis, they will be able to discover basic principles of learning that can then be used to explain a wide range of complex learning phenomena.

Hence the interest in laboratory studies of learning in animals. Just as the geneticist has studied the genetics of the fruitfly in the laboratory to determine generally applicable laws of inheritance, so the psychologist has studied the behaviour of the rat in the maze in the hope of discovering equally general laws of learning. Whether psychologists have been successful in this endeavour, the rest of this chapter will tell.

CLASSICAL CONDITIONING

Our definition of learning mentions changes taking place. What kinds of changes are we talking about? The physical basis of the changes that constitute learning lies in the brain, and neuroscientists are close to discovering exactly what these changes are. But our concern in this chapter is with the psychological mechanisms of learning, rather than the physiological mechanisms.

association a link between two events or entities that permits one to activate the other (such as when a characteristic odour elicits an image of the place where it was once experienced)

Foremost among these is the concept of *association*. There is a philosophical tradition, going back at least 300 years, which supposes that, when two events (ideas or states of consciousness) are experienced together, a link, connection or association

forms between them, so that the subsequent occurrence of one is able to activate the other.

In the twentieth century the proposal was taken up by experimental psychologists, who thought that association formation might be a basic psychological process responsible for many, if not all, instances of learning. The first to explore this possibility in any depth was the Russian I.P. Pavlov with his work on *classical conditioning*.

classical conditioning learning procedure in which two stimuli are paired – one (the conditioned stimulus) usually presented shortly before the other (the unconditioned stimulus) to produce a conditioned response to the first stimulus (learning)

PAVLOV'S DOGS

Pavlov spent the first half of his long scientific career working on the physiology of digestion, turning to the study of learning in about 1900. He had noticed that dogs which salivate copiously when given food also do so in response to other events – for example, at the approach of the laboratory attendant who

supplied the food. This response was clearly acquired through experience. Pavlov (1927) took a version of this procedure into the laboratory, making it a model system that could be used to reveal basic principles of learning.

Pavlov's standard procedure involved a quiet, distraction-free laboratory, which gave the experimenter full control over events experienced by a lightly restrained dog. From time to time the dog was given access to food, and each presentation was accompanied (usually slightly preceded) by the occurrence of a neutral event, such as a flashing light. After several training trials (pairings of light and food), the dog would salivate at the flash of light, before any food had appeared.

unconditioned response (UR) evoked by a stimulus before an animal has received any explicit training with that stimulus

unconditioned stimulus (US) evokes an unconditioned response

conditioned response (CR) evoked by a conditioned stimulus as a result of classical conditioning

conditioned stimulus (CS) evokes a conditioned response as a result of classical conditioning

Salivation at the presentation of food is called an *unconditioned response* (UR), since it occurs automatically (unconditionally). The food is an *unconditioned stimulus* (US). The animal's tendency to salivate when the light flashes is conditional on the light having been paired with food, so this is referred to as a *conditioned response* (CR) and the event that evokes it as a *conditioned stimulus* (CS). The whole training procedure was labelled conditioning. As other forms of training, introduced later, have also been described as conditioning, Pavlov's version became known as classical conditioning.

Pioneer

I.P. Pavlov (1849–1936), born the son of a priest in Ryazan (250 miles south-east of Moscow), moved in 1870 to study natural science and medicine in St Petersburg. He spent the rest of his life there conducting scientific research, first on the physiology of the digestive system (for which he was awarded a Nobel prize in 1904) and later on conditioned reflexes. Although the study of conditioned reflexes was taken up mostly by psychologists, Pavlov insisted that his approach as a physiologist was far superior to that adopted by the comparative psychologists of his day. His demonstration of the salivary conditioned reflex in dogs, for which he is widely known, was just the start of an extensive body of work, in which he analysed the conditioning process in detail, revealing phenomena and suggesting learning mechanisms that are still being actively investigated today.

OTHER EXAMPLES OF CONDITIONED RESPONSES

Following Pavlov's pioneering work, the study of classical conditioning has been taken up in many laboratories around the world. Few of these have made use of dogs as the subjects and salivation as the response, which are merely incidental features of conditioning. The defining feature is the paired presentation of two stimuli – the CS and the US. The presentation of the US is often said to be contingent on (i.e. to depend on) the presentation of the CS.

Here are just a few of the wide range of training procedures that employ this contingency:

- **Conditioned emotional response** The experimental animal, usually a rat, is presented with a neutral cue, such as a tone sounding for one minute (the CS), paired with a mild electric shock (US) that occurs just as the tone ends. After several pairings (the exact number will depend on the intensities of tone and shock), the rat's behaviour changes. It begins to show signs of anxiety, such as freezing and other 'emotional responses', when it hears the tone before the shock has occurred. This is the CR.
- **Autoshaping** A hungry pigeon is presented with grain (US) preceded by the illumination for ten seconds of a small light (CS) fixed to the wall of the cage. After 50 to 100 trials, the bird develops the CR of pecking at the light prior to food delivery. It is as if the bird is predisposed to respond to the light even though the pecking does not influence whether or not it receives the grain.
- **Flavour aversion learning** Rats are given a novel flavour (e.g. saccharin is added to their drinking water) as the CS. This is followed by a procedure, such as the injection of a mild poison into their body, that makes them feel sick (the US). When it is subsequently made available, the rats will no longer consume the saccharin-sweetened water; they have developed an aversion (CR) to that flavour.

conditioned emotional response result of the superimposition of the pairing of a conditioning and an unconditioned stimulus on a baseline of operant or instrumental behaviour

autoshaping classical conditioning used with pigeons which results in pecking at an illuminated response key that has been regularly presented before the delivery of food, even though the delivery of the food does not depend on the pecking behaviour

flavour aversion learning classical conditioning procedure in which animals are allowed to consume a substance with a novel flavour and are then given some treatment that induces nausea, resulting in the flavour being subsequently rejected

This is clearly a very varied set of phenomena, but what they all have in common is the presentation of two stimuli, one contingent on the other. And, despite the fact that there is nothing in these training procedures that actually requires a change in

behaviour, in every case the animal's behaviour changes as a result of its experience.

In the autoshaping case, for instance, the experimenter simply ensures that the light reliably accompanies food. There is no need for the pigeon to respond to the light in any way, since food is delivered regardless of the bird's behaviour.

So why does behaviour change? Why are conditioned responses acquired? This puzzle must be dealt with by more detailed theoretical analysis.

ASSOCIATIVE ANALYSIS

When a dog trained by Pavlov's procedure sees the light (CS), certain neural mechanisms are activated. Without specifying what these mechanisms are, we can refer to this pattern of activation as constituting a representation of the CS. This is often referred to as the CS 'centre', implying that it is localized in a specific part of the brain, although this might not necessarily be the case (for the purposes of our current behavioural analysis, this does not matter too much). Eating food (the US) will also have its own pattern of proposed neural activation, constituting the US representation or 'centre'.

One consequence of the Pavlovian conditioning procedure is that these two centres will be activated concurrently. Pavlov suggested that concurrent activation results in a connection between the two centres, which allows activation in one to be transmitted to the other. So, after Pavlovian learning has taken place, presentation of the CS becomes able to produce activity in the US centre, even when the food has not yet been presented.

This theory therefore explains classical conditioning in terms of the formation of a stimulus–stimulus association between the CS centre and the US centre. (Given this framework, the fact that the presentation of the US provokes an obvious response is not strictly relevant to the learning process.)

Sensory preconditioning

If this account is correct, it should be possible to trigger classical conditioning using paired neutral stimuli that themselves evoke

sensory preconditioning pairing of two neutral stimuli prior to one of them being used as the conditioned stimulus in a standard classical conditioning procedure, leading to the other stimulus acquiring the power to evoke the conditioned response

no dramatic responses. Evidence that this can occur comes from a phenomenon called *sensory preconditioning*, first demonstrated by Brogden (1939) and confirmed many times since. In Brogden's experiment (see table 4.1), the animals in the critical experimental condi-

tion received a first stage of training consisting of paired presentations of two neutral stimuli, a light and a buzzer. If our theory is correct, an association should be formed between the central representations of these stimuli. The problem is to find a way to reveal this association.

Brogden's solution was to give a second stage of training in which one of the original stimuli (say the light) was given

Table 4.1 Design and results of the experiment by Brogden (1939) on sensory preconditioning.

Condition	Phase 1	Phase 2	Test (leg flexions to B)
Experimental	A + B	A → shock	9.5
Control	no training	A → shock	0.5

A and B represent a light and a buzzer.

orthodox conditioning, being paired with a US until it came to evoke a CR (in this procedure, a response of flexing the leg). A final test showed that the buzzer was also able to evoke the leg flexion response, even though the buzzer had never previously been paired with the US. This result is what might be expected on the basis of the stimulus–stimulus association theory. The light evokes the CR by virtue of its direct association with the US, whereas the buzzer is able to do so 'by proxy' because its association with the light allows it to activate the representation of that stimulus.

Why and how does the CR occur?

What remains to be explained, once the stimulus–stimulus association theory has been accepted, is why the CR should occur and why it should take the form that it does. Pavlov's dogs might 'know', by virtue of the CS–US link, that light and food go together, but this does not necessarily mean that the animal should start to salivate in response to the light. The most obvious explanation is that activation of the US (food) centre will evoke a given response, whether that activation is produced by presentation of the US (food) itself or, via the learned CS–US (light–food) connection, by presentation of the CS (light).

An implication of this interpretation is that the CR and the UR should be the same, and this is true for the case just considered: the dog salivates (as a UR) to food and also comes to salivate (as a CR) to the light that has signalled food. In other examples of conditioning, however, the CR and UR are found to differ. In the autoshaping procedure, for instance, the UR is to approach and peck inside the food tray, whereas the CR that develops with training is to approach and peck at the light. In this case, the CR appears to be a blend of the behaviour that activation of the US (food) centre tends to evoke and the behaviour evoked by the CS (the light) itself.

So we cannot say that the CR and the UR are always the same. There is, however, a simple rule that describes the relationship between them for most cases of conditioning, in that, as a result of classical conditioning, the animal generally comes to behave toward the CS (the light in these examples) as if it were the US (food). In other words, the CS (light) appears to take on some of the properties of the US (food) and to serve as an adequate substitute for it. So the unconditional response of a hungry animal is to approach food, probably salivating as it does so, and then to consume the food (by pecking, if the animal is a pigeon). The CR consists of directing these behaviour patterns toward the CS, in

Everyday Psychology

How are phobias acquired?

phobias intense and seemingly irrational fears

Phobias can be very debilitating and distressing phenomena. Many of us know someone who is anxious about enclosed spaces, needles or spiders. Watson and Rayner (1920) speculated that the complexity of emotional responsiveness in adults might be explained by the conditioning of children's simple emotional reactions when they are exposed to new stimuli in their youth. They decided to test this proposal by attempting to establish a conditioned emotional response in a child under experimental conditions.

The study was done with a 'stolid and unemotional' infant boy, Albert B. At eight months of age he was exposed to a number of stimuli, including a white rat, a rabbit and a monkey, and showed no signs of fear to any of these stimuli. The fear reaction could be produced, however, by a sudden loud noise (produced by striking a steel bar with a hammer). On the first presentation of this noise, the child 'started violently'; on the second occasion, the 'lips began to pucker and tremble'; and on the third, the child 'broke into a sudden crying fit'.

At 11 months, Albert was subjected to the conditioning trials. The CS was a white rat, which Albert was allowed to play with, and the US was the loud noise. On six occasions over the course of a week, Albert was presented with the rat and at the same time he was subjected to the noise produced by striking the steel bar.

Evidence for emotional conditioning came from a test trial in which the rat was presented alone (i.e. not accompanied by the noise). Here is Watson and Rayner's description of the result: 'The instant the rat was shown, the baby began to cry. Almost instantly he turned sharply to the left, fell over on [his] left side, raised himself on all fours and began to crawl away' (1920, p. 5). Subsequent tests show that the fear response generalized to other stimuli, both animate and inanimate. A rabbit and a fur coat both produced a strong response, although the responses to a dog and to cotton wool were less marked.

This experiment establishes three important points:

1. It confirms, in a very vivid way, that classical conditioning processes work for humans as well as non-human animals.
2. It establishes that conditioning can influence whole patterns of emotional responding (in addition to the simple reflexes that had been most commonly studied up to this point in the learning literature).
3. Watson and Rayner note the parallel with the clinical condition of phobia – intense and seemingly irrational fear of intrinsically harmless objects or events. They suggest that phobias present in adults may be the product of a traumatic conditioning episode occurring earlier in life. Although it is not always possible to obtain evidence of such an episode, the general proposition that phobias derive from conditioning is widely accepted, albeit these days with some caveats related to the notion of preparedness and the apparent selectivity of some learning phenomena.

Watson, J.B., & Rayner, R., 1920, 'Conditioned emotional reactions', *Journal of Experimental Psychology*, 3, 1–14.

stimulus substitution when the conditioned stimulus comes to acquire the same response-eliciting properties as the unconditioned stimulus

so far as the physical properties of the event used as the CS will allow this. This rule is sometimes referred to as the *stimulus substitution* hypothesis.

1. Although the behavioural consequence of conditioning may appear to be merely the development of an anticipatory reflex, the underlying process is fundamental to learning about the relationship among environmental events. Sensory preconditioning tells us that when neutral stimuli co-occur, an association forms between them. Presumably, the informal equivalent of sensory preconditioning will be occurring all the time as an animal goes about its normal everyday business. Simply moving through the environment will expose the animal to sequences of events that go together, and the associations that form among them will constitute an important piece of knowledge – a 'map' of its world.
2. As a laboratory procedure, classical conditioning is important because it allows exploration of the nature of associative learning. The observed CR (salivation, pecking, or whatever) may not be of much interest in itself, but it provides a useful index of the otherwise unobservable

THE IMPORTANCE OF CLASSICAL CONDITIONING

If classical conditioning were simply a procedure that allows a reflex response previously elicited solely by a particular US (such as food) to come under the control of another stimulus (such as the presentation of a light), then perhaps there would be no reason to regard it as fundamentally important to our understanding of learning. But three features of our analysis give us reason to believe that it is fundamentally important:

formation of an association. Researchers have made extensive use of simple classical conditioning procedures as a sort of ‘test bed’ for developing theories of associative learning. Some of these will be described in a later section of this chapter.

3. As a mechanism of behavioural adaptation, classical conditioning is an important process in its own right. Although the CRs (such as salivation) studied in the laboratory may be trivial, their counterparts in the real world produce effects of major psychological significance. Here are two examples from the behaviour of our own species.

Illness-induced aversion learning

Experiencing illness after consuming a given flavour will induce an aversion to that flavour, not just in rats, but in people too. Informal surveys of undergraduate students reveal that about 50 per cent report having an aversion to a particular flavour (usually a novel alcoholic drink). Most can clearly remember the occasion on which they tasted that flavour and subsequently became ill.

More significant are the aversions that can develop with the severe nausea that sometimes results from chemotherapy used to treat cancer. Chemotherapy patients sometimes find that strongly flavoured foods eaten prior to a session of treatment start to develop aversive properties. The need to change our eating habits is an inconvenience. But more worrying is the phenomenon known as anticipatory nausea and vomiting. Some patients (up to 50 per cent for some forms of treatment) develop an aversion to the clinic in which treatment is given, so that, after a few sessions, they begin to feel nauseous and even vomit as soon as they walk in. This reaction can be so severe that the patient refuses to continue treatment, with obvious life-threatening consequences.

Research on the nature of association formation has suggested ways of limiting this clinic–illness association (Hall, 1997).



Figure 4.2

Rats can learn to avoid a food associated with illness.

Emotional conditioning

The conditioned emotional response was first demonstrated not in rats, but with a human participant. In what may well be the most famous and influential experiment in psychology (although not one that would survive the scrutiny of a modern-day ethics committee), Watson and Rayner (1920) set out to establish that Pavlovian conditioning procedures would be effective when applied to a human infant. See *Everyday Psychology* and chapter 1 for more on this.

INSTRUMENTAL LEARNING

THORNDIKE'S CATS

At about the time that Pavlov was beginning work on classical conditioning in Russia, E.L. Thorndike, in the United States, was conducting a set of studies that initiated a different tradition in the laboratory study of basic learning mechanisms.

Thorndike was interested in the notion of animal intelligence. Motivated by an interest in Darwinian evolutionary theory, comparative psychologists of the late nineteenth century had investigated whether non-human animals can show similar signs of intelligence to those shown by humans. Thorndike took this endeavour into the laboratory. In his best-known experiment, a cat was confined in a ‘puzzle box’ (figure 4.3). To escape from the box, the cat had to press a latch or pull a string. Cats proved able to solve this problem, taking less and less time to do so over a series of trials (figure 4.4).

Cats solved the problem not by a flash of insight but by a gradual process of trial and error. Nevertheless, here was a clear

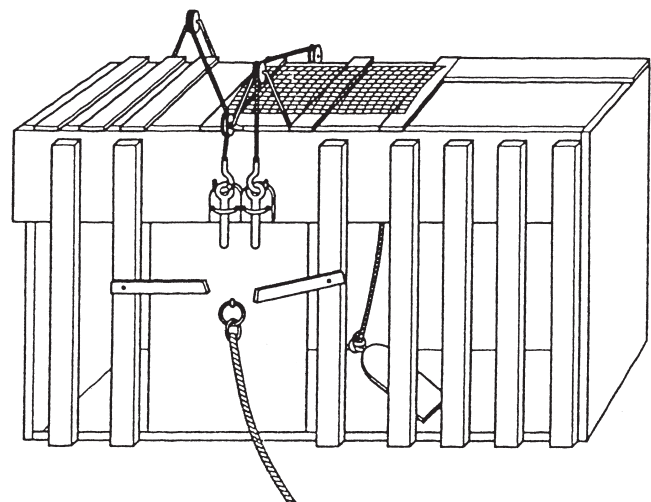


Figure 4.3

One of the ‘puzzle’ boxes used by Thorndike (1898) in his studies of instrumental learning in the cat.




Figure 4.4

Time taken by a cat to escape from the puzzle box over a series of 20 successive trials. Source: Based on Thorndike (1898).



Figure 4.5

A rat in a Skinner box. The animal pictured has his nose next to a lever; when it depresses the lever, a food pellet can be delivered from the container outside the chamber on the left. In normal use, the apparatus is enclosed in a sound- and light-attenuating outer shell.

instrumental learning the likelihood of a response is changed because the response yields a certain outcome (a reward or punishment) (also called operant conditioning)

from classical conditioning, in which the animal's response plays no role in determining the outcome.

Subsequent researchers who took up the analysis of this form of learning include the Polish physiologist Konorski (1948), who called it Type II conditioning (as distinct from Pavlov's Type I conditioning). Another investigator interested in this type of conditioning was Skinner (1938) in the United States, who named it operant conditioning (Pavlov's version of learning being referred to as *respondent conditioning*).

respondent conditioning alternative name for classical conditioning

behaviour (or response) and a subsequent state of the environment (the effect or outcome).

THE SKINNER BOX

The Skinner box soon replaced Thorndike's puzzle box in the laboratory study of instrumental learning. In the version used for the rat, the Skinner box consists of a chamber with a lever protruding from one wall and a nearby food cup into which food pellets can be delivered by remote control (figure 4.5). Pressing the

example of learning. Its characteristic feature was that the animal's actions were critical (instrumental) in producing a certain outcome. In this respect, *instrumental learning* is fundamentally different

lever operates an electronic switch and automatically results in food delivery. So there is an instrumental contingency between the lever-press (the response) and the food (the effect or outcome). A rat exposed to this contingency presses the lever with increasing frequency. The Skinner box is similar to Thorndike's puzzle box, but instead of using escape from the box as a reward, the animal stays in the box and the reward is delivered directly to it.

This is an example of rewarded, or appetitive, instrumental learning, but the same general techniques can be used to study aversive instrumental learning. There are two basic aversive paradigms, *punishment* and *avoidance*.

punishment an aversive event as the consequence of a response to reduce the probability of the response

avoidance instrumental training procedure in which performing a given response brings about the omission of an aversive event that is otherwise scheduled to occur

Punishment The event made contingent on the response is aversive. For example, the habit of responding is first acquired. Subsequently, occasional lever-presses produce a brief electric shock through a grid floor fitted to the box. Unsurprisingly, the rate of responding declines. (It is worth adding that, although the effect may not be surprising, it still requires explanation. It often happens in psychology that the basic behavioural facts seem obvious; but when we try to explain them, we realize how little we really understand them.)

Avoidance A signal occurs from time to time, accompanied by a foot shock. If the rat presses the lever, the shock is cancelled. So there is an instrumental contingency between the response and the omission of a given outcome. By behaving appropriately, the animal can avoid the shocks. In fact, rats are rather poor at avoidance learning when the response required is a lever-press; they respond better when they are required to jump over a hurdle. So the apparatus usually used is a two-compartment box, with a hurdle separating the two parts. Rats readily acquire the habit of jumping the hurdle in response to the warning signal.

Training procedures that inflict pain (however slight) on the animal should obviously be employed only for good reason. Studies like this are justified by the insights they have provided into the nature of human anxiety disorders and neuroses (see below).

THE LAW OF EFFECT

Thorndike's studies of cats in the puzzle box led him to propose the following interpretation of their behaviour: 'Of several responses made to the same situation, those which are accompanied or closely followed by a state of satisfaction to the animal

Pioneer

B.F. Skinner (1904–90) developed the framework of radical behaviourism, focusing on establishing laws of behaviour (empirical relationships between environmental events and behaviour). This framework was based on the intensive observation of a single subject under carefully controlled experimental conditions. His approach, the experimental analysis of behaviour, investigated 'operant' behaviours – so-called because they 'operate' on the subject's environment. Skinner's classic work involved the study of bar pressing (or pecking) by rats (or pigeons) in a 'Skinner box' that was constructed to eliminate all extraneous stimuli. A hungry animal was placed in the box and allowed to explore it. Sooner or later the animal would accidentally press a lever that released a food pellet. The food acted as a reinforcing stimulus (or reinforcer) for the bar-pressing behaviour, increasing the probability of its future occurrence.

will, other things being equal, be more firmly connected with the situation, so that, when it recurs, they will be more likely to recur' (Thorndike, 1911, p. 244). This is the *law of effect* as applied to appetitive instrumental learning.

Thorndike also put forward (and later retracted) a negative counterpart for the case of punishment, which proposed that certain effects ('annoyers') would weaken the connection between a response and the training situation. In modern terminology, Thorndike's 'satisfiers' and 'annoyers' are called reinforcers and punishers.

Thorndike's presentation of the law of effect has two major features:

1. What is learned is a stimulus–response (S–R) association.
2. The role of the effect produced by the response is to determine whether this association will be strengthened or not.

Both of these propositions are debatable and, as we shall shortly see, this theoretical version of the law of effect has not stood up well to further experimental analysis. As an empirical generalization, though, the law seems much more secure. Everyone accepts that the likelihood of an animal responding in a particular way can be powerfully controlled by the consequence of that response.

Partial reinforcement

Skinner, who completely rejected the theoretical law of effect, devoted several years of research (e.g. Ferster & Skinner, 1957) to exploring and demonstrating the power of the empirical law. He worked mostly with pigeons, trained in a Skinner box to peck a disc set in the wall for food reinforcement. Skinner investigated the effects of *partial reinforcement*, in which food was presented after some responses but not all. Animals will usually respond well in these conditions, and with some *schedules of reinforcement* the rate of response can be very high indeed. If, for example, the animal is required to respond a certain

number of times before food is delivered (known as a fixed ratio schedule), there will usually be a pause after reinforcement, but this will be followed by a high frequency burst of responding. Other ways of scheduling reinforcement control different but equally systematic patterns of response. There is a clear parallel here between the pigeon responding on a partial reinforcement

law of effect Thorndike's proposal that reward will strengthen the connection between the response that preceded it and any stimuli present when it is delivered, or more generally, the principle that the consequence (effect) of behaviour will determine how likely it is to recur

partial reinforcement the delivery of a reinforcer in operant conditioning is scheduled to occur after only a proportion of the responses rather than after all of them (continuous reinforcement)

schedules of reinforcement rules that determine which responses will be followed by a reinforcer in operant conditioning (see **partial reinforcement**)

schedule and the human gambler who works persistently at a one-armed bandit for occasional pay-outs.

The punishment paradigm

For a while, doubts were raised about how reliable the negative version of the empirical law of effect was. It soon became clear, however, that early studies failed because the punishment (such as the presentation of white noise) was too weak. Subsequent work using more intense punishments, such as shock, confirmed the effectiveness of the procedure in suppressing behaviour. What remained to be shown was that the shock had its effect by way of the instrumental contingency.

The following study conducted by Church (1969) investigated this question. Three groups of rats were all trained to lever-press for food. One group then began to receive occasional shocks contingent on lever-pressing (contingent group). A second group received the same number of shocks but these occurred independently of lever-pressing (noncontingent group). The third group of rats was given no shocks (control group). Church found that simply presenting shocks in the apparatus, with no contingency on behaviour, was enough to produce some response suppression. So the threat of shock (an effective Pavlovian unconditioned stimulus or US) is enough in itself to suppress behaviour to some extent. But powerful suppression of the response was seen only in the contingent group, demonstrating that the instrumental contingency between the response and the outcome is effective in producing pronounced learning (see figure 4.6).




Figure 4.6

Results of an experiment by Church (1969) on the punishing effects of shock. The contingent group, which received shock when it responded, came to show response suppression. (A ratio score of 0.5 means no suppression; a score of zero means complete suppression.) The noncontingent group received shocks independently of its behaviour and showed less suppression. The control group received no shocks and showed no suppression.

When learning becomes habit

According to the theoretical version of the law of effect, the only function of the *reinforcer* is to strengthen a connection between the response (R) that produced that reinforcer and the stimulus (S) that preceded the R. It

reinforcer an event that, when made contingent on a response, increases the probability of that response; also another term for the unconditioned stimulus in classical conditioning

follows that an S-R learner does not actively know what the consequence of the R will be, but rather the response is simply triggered based on previous contingencies. In other words, the rat in the Skinner box is compelled in a reflex-like fashion to make the R when the S is presented and it is presumed to be as surprised at the delivery of the food pellet after the hundredth reinforced response as it was after the first. Not only is this an implausible notion, but experimental evidence disproves it.

The evidence comes from studies of the effects of reinforcer reevaluation on instrumental performance. The results of one such study are summarized in figure 4.7. In a first stage of training, rats were allowed to press the lever in a Skinner box 100 times, each response being followed by a sugar pellet. Half the animals were then given a nausea-inducing injection after eating sugar pellets – a flavour-aversion learning procedure. As you might expect, these rats developed an aversion to the pellets, so the reinforcer was effectively devalued.




Figure 4.7

Results of an experiment by Adams (1982) on the effects of reinforcer devaluation on instrumental responding. Scores are from a 20-minute test session in which rats were allowed to respond by depressing a lever without consequences. In initial training some animals had received 100 reinforced responses, others 500. For half the animals in each condition the reinforcer was then devalued by being associated with illness.

In the subsequent test phase, the rats were returned to the Skinner box and allowed access to the lever (although no pellets were now delivered). The researchers found that rats given the devaluation treatment were reluctant to press the lever, compared with the control animals. This result makes common sense – but no sense in terms of the theoretical law of effect. According to the strict interpretation of the law of effect, an S–R connection would have been established at the end of the first stage of training by virtue of the reinforcers that followed responding, before the nausea-inducing injection was administered. Subsequent changes in the value of this reinforcer (which, according to the theory, has already done its job in mediating a ‘state of satisfaction’) should have been of no consequence.

These results suggest that the critical association in instrumental learning is not between stimulus and response, but between representations of a) the response and b) the reinforcer (or more generally, between the behaviour and its outcome). The stronger this association, assuming that the outcome is valued, the more probable the response will be. But an association with an aversive outcome (i.e. a devalued foodstuff or a punishment) will lead to a suppression of responding.

This does not mean that S–R learning can never occur. Often, after long practice, we acquire patterns of behaviour (habits) that have all the qualities of reflexes. In other words, they are automatically evoked by the stimulus situation and not guided by consideration of their consequences. The results shown in figure 4.5 may be an experimental example of this. One group of rats was given extensive initial training in lever-pressing (500 rather than 100 reinforced trials) prior to the reinforcer-devaluation treatment. As the figure shows, these animals continued to press the lever in the test phase. One interpretation of this result is that with extensive training, behaviour that is initially goal-directed (i.e. controlled by a response–outcome association) can be converted into an automatic S–R habit.

When next you absent-mindedly take the well-worn path from your home to the college library, forgetting that on this occasion you were intending to go to the corner shop, your behaviour has been controlled by an S–R habit rather than the response–outcome relationship – just like the rats!

CONTROL OF PERFORMANCE

If an animal has acquired an S–R habit, then we can predict that the R will occur whenever the S is presented. But what controls performance if learning is the result of a response–outcome association?

A rat can be trained to press for food or jump to avoid shock only in the presence of a given stimulus (called a *discriminative stimulus*) which signals that food or shock are likely to occur. Presumably the response–outcome association is there all the time, so why is it effective in producing behaviour only when the stimulus is present? How does the presentation of the discriminative stimulus activate the existing instrumental association?

discriminative stimulus signals whether or not a given response is likely to produce a particular outcome

Classical conditioning and motivational control

For instance, a rat trained on an avoidance task, in which the sounding of a tone indicates that shock is likely, will, at least before the avoidance response has been fully learned, experience some pairings of the tone and the shock. As well as acquiring a response–outcome association, the rat can also be expected to form a tone–shock association. In other words, classical conditioning will occur, as a sort of by-product of the instrumental training procedure.

This Pavlovian (S–S) association, it has been suggested, is responsible for energizing instrumental responding. By virtue of the S–S link, the tone will be able to activate the shock representation, producing in the animal both an expectation of shock and the set of emotional responses that we call fear. The state of fear is presumed to have motivational properties, so that the presentation of the tone could effectively boost the supply of energy that causes the animal to behave.

The expectation evoked by the tone also gives value to the outcome. In avoidance learning, the outcome associated with the response is the absence of an event (the omission of shock). The absence of an event would not normally be reinforcing in itself, but it could certainly become so given the expectation that something unpleasant is likely to occur.

This account of avoidance learning is a version of *two-process theory*, so called because it acknowledges that classical and instrumental learning processes both play a part in determining this

two-process theory emphasizes the interaction of instrumental and classical conditioning processes in producing many types of behaviour

type of behaviour. Although the theory was first elaborated in the context of avoidance learning, there is no reason to suppose that it applies only to this procedure. We have already seen how classical conditioning might contribute to the response suppression generated by the punishment procedure (see the earlier discussion of the experiment by Church, 1969, and figure 4.4). In the appetitive case, stimuli present when an animal earns food by performing an instrumental response can be expected to become associated with the food. These stimuli will then be able to evoke a positive state (an ‘expectation of food’, a ‘state of hopefulness’) that parallels the negative, fearful, state produced in aversive training procedures.

Conditional learning and stimulus control

Although the ability of the discriminative stimulus to evoke a (conditioned) motivational state is undoubtedly important, this still does not fully explain how it controls instrumental responding.

It is difficult to believe that a rat that receives food for lever-pressing in the presence of a tone is insensitive to the conditional nature of the task – in other words, that it fails to learn that the response yields food only if the tone is on. But the version of two-process theory just described proposes only that the rat will form two simple associations – stimulus–food and response–food.

There is no room in this account for the learning of a conditional relationship of the form ‘only lever-pressing in the presence of the tone results in the presentation of food’.

This issue has been addressed experimentally in recent years, and several researchers have demonstrated that animals are capable of conditional learning. The stimulus control of performance revealed by these experiments cannot be explained in terms of standard two-process theory, in which discriminative stimuli

have their effects solely by virtue of orthodox associations with reinforcers. Instead, it shows that animals are capable of learning the conditional relationship between a stimulus and a particular response–reinforcer relationship. So, discriminative stimuli exert their effects because they are able to trigger not just the representation of the reinforcer but also the more complex, response–outcome representation produced by instrumental training. This represents the learning of a conditional relationship.

Research close-up 1

The hierarchical structure of instrumental learning

The research issue

If lever-pressing is rewarded only in the presence of a tone, a rat will learn to respond only when the tone is sounding. How is such stimulus control achieved? Colwill and Rescorla (1990) accepted that the formation of a classically conditioned association between the tone and food could play a role (as outlined in our discussion of two-process theory on p. 82). But they also suspected that some other process might be involved. So they devised a training procedure which attempted to establish stimulus control that could not be the consequence of direct stimulus–food association.

Design and procedure

A slightly simplified version of their experimental design is presented in table 4.2. The apparatus offered two possible responses (R1 and R2): the rats could press a lever or pull on a chain hanging down from the ceiling. Two different reinforcers (rf1 and rf2) were also available: the rats could be trained to make one response to receive a standard food pellet and the other response for a small amount of sugar solution. But which response produced which reinforcer depended on the stimulus conditions. Two stimuli (S1 and S2, a light and a noise) were used. In the presence of S1, the lever produced food and the chain sucrose; in the presence of S2, the lever produced sucrose and the chain food – so both S1 and S2 became associated with both the reinforcers. This is a conditional discrimination, in that the outcome of a given response differs according to which background stimulus was present.

In Phase 2, instrumental training was discontinued. Rats simply received ‘free’ deliveries of one of the reinforcers. After consuming the food or sucrose, the rat received a nausea-inducing injection, so that this particular reinforcer became devalued. In the final test phase, the rats were given access to the lever and the chain again. The light and the noise were each presented eight times and the rate of response in the presence of each stimulus was noted.

Results and implications

Figure 4.8 shows the rate of response for each test trial and for two categories of responding – one in which the stimulus signalled that the response chosen might lead to the devalued reinforcer, and one in which the stimulus signalled that the response chosen might lead to the valued reinforcer. In fact, no reinforcers were presented during the test (hence the steady decline in responding).

Table 4.2 Design of the experiment by Colwill and Rescorla (1990) on stimulus control.

Phase 1 (discrimination)	Phase 2 (devaluation)	Test
S1: R1 → rf1; R2 → rf2	rf1 → illness	S1: R2 vs. R1
S2: R1 → rf2; R2 → rf1		S2: R1 vs. R2

R1 and R2 represent two different responses, pulling a chain, and pressing a lever; S1 and S2 represent two different discriminative stimuli, noise and light; rf1 and rf2 represent two different reinforcers, food pellets and a sucrose solution. In the test, both responses were available and the rate at which they were performed in the presence of each stimulus was measured. No reinforcers were presented during the test.

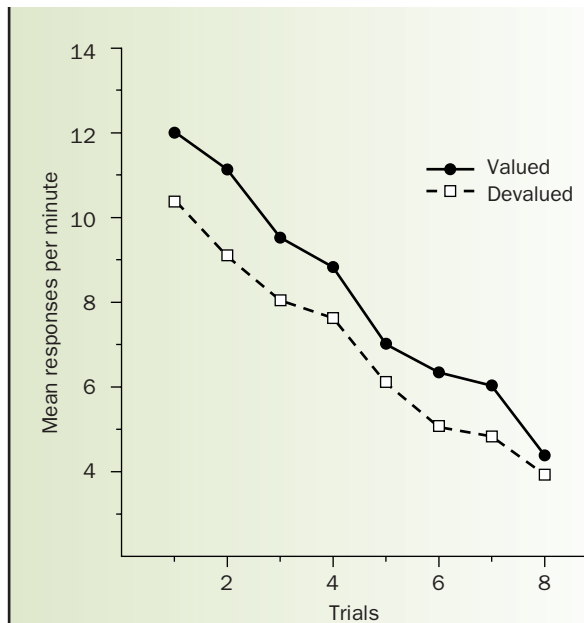


Figure 4.8

Test performance in the experiment by Colwill and Rescorla (1990). Each of the two possible responses might lead to a devalued reinforcer or to a valued reinforcer, according to the stimulus presented. The higher rate of response in the valued condition, compared with the devalued condition, indicates that the rats were less likely to make the response that (given the stimulus present) might be expected to produce the devalued reinforcer.

The important result is that the rats were unwilling to perform the response that might produce the devalued reinforcer. They were behaving selectively, showing an unwillingness to perform R1 in the presence of S1 and to perform R2 in the presence of S2. This selectivity cannot be explained in terms of classically conditioned associations between stimuli and reinforcers, because both stimuli have been associated with both reinforcers, so their classical associations are equivalent. Devaluing one of the reinforcers might reduce the animal's general readiness to respond, but there is no reason to predict that this would occur selectively to one particular response in the presence of a particular stimulus.

Colwill and Rescorla concluded that the rat is capable of learning the conditional relationship between a stimulus and a given response–reinforcer relationship – for example, that R1 produces food in the presence of S1 but produces sucrose in the presence of S2. This implies the existence of a hierarchy of associations in instrumental learning. Not only do the animals form the response–outcome association already discussed, but they also form a link between a stimulus and the response–outcome association that has been formed in its presence. This higher-order association allows the stimulus to activate the lower-order link, contributing to the stimulus control phenomenon.

Colwill, R.M., & Rescorla, R.A., 1990, 'Evidence for the hierarchical structure of instrumental learning', *Animal Learning and Behavior*, 18, 71–82.

THE IMPORTANCE OF INSTRUMENTAL LEARNING

As we have seen, classical conditioning allows an animal to learn about the relationship between events in the environment and so anticipate what will happen next on the basis of stimuli currently present. If there are grey clouds in the sky, then it will probably rain; if the light is presented, then food may well follow. Instrumental learning is the process by which an animal learns about the relationship between its behaviour and the consequences of that behaviour. And it serves a complementary but equally important function in allowing the animal to control (at least partially) the occurrence of environmental events – in other words, to bring about a desired event or to avoid an aversive event by responding in a particular way.

Instrumentally trained responses are not entirely elicited by identifiable stimuli. Instead, they are controlled by their consequences, becoming more likely when they produce a positive result and less likely when they lead to an aversive outcome. As Skinner emphasized, this sort of control is the characteristic feature of what we call 'voluntary' behaviour. So the study of instrumental learning and performance is important for what it tells us about the nature of voluntary, goal-directed behaviour.

On the other hand, instrumental learning processes can also play a role in establishing and maintaining behaviour that seems, at first sight, anything but voluntary. Patients with the clinical condition known as *obsessive–compulsive disorder* (OCD) suffer from persistent, intrusive, unpleasant thoughts (obsessions) and feel compelled repeatedly to carry out certain acts (compulsions) that they know are senseless but which appear to provide some relief (see chapter 15). OCD can be quite disabling. One patient, who believed that contact with everyday objects contaminated her in some way, felt compelled to shower at least six times a day and to wash her hands very systematically every 20 minutes. With hands rubbed raw and half her working day taken up in these activities, her ability to lead a normal life was severely curtailed. OCD patients tend to feel a build-up of extreme anxiety prior to performing the compulsive ritual, which dissipates as the ritual is enacted. This has been measured both by patients' own reports and by objective indices such as heart-rate (Hodgson & Rachman, 1972).

obsessive–compulsive disorder (OCD) characterized by intrusive unwelcome thoughts (obsessions) and the need repeatedly to perform certain patterns of behaviour (compulsions), such as hand-washing

A parallel can be drawn between such cases and a trained rat 'compulsively' responding to the presentation of a tone by jumping

a hurdle, and continuing to perform this apparently senseless act for a large number of trials in the absence of any obvious reward. Although this behaviour appears senseless, it becomes understandable when the rat's training history is known – when it becomes clear that the tone evokes fear by virtue of its initial association with shock and that the response avoids a shock that would otherwise occur.

In the same way, the rituals performed by OCD patients may well be avoidance responses that are reinforced and maintained because they reduce the sufferer's state of anxiety. Of course it remains to be explained why the patient has acquired such a fear of dirt, or whatever, in the first place. Nevertheless, this illustration demonstrates the relevance of the analysis of basic instrumental learning processes to an understanding of interesting and important aspects of human behaviour.

THE PRINCIPLES OF ASSOCIATION FORMATION

Classical conditioning and instrumental learning both depend on the formation of associations. An association will be formed

principle of contiguity the proposal that events must be experienced close together in time and space for an association to be formed between them

between a pair of events (two stimuli, or a response and a stimulus) that occur together (in contiguity). This *principle of contiguity* is clearly important, but it has some limitations.

CONTIGUITY AND PREDICTIVENESS

Asymptote – where learning stops

When a CS (e.g. a light) and a US (e.g. food) occur together, an association appears to be established between their central (i.e. neural) representations. And the more often they occur together, the stronger this association becomes. This is revealed by the growing strength of the CR (e.g. light-induced salivation). But this growth does not go on forever. With repeated CS–US pairings, the increment in the strength of the CR (and also, we deduce, the underlying association) becomes progressively smaller until there is no observable increase in its strength. At this point – referred to as asymptote – contiguity between the CS (light) and US (food) is no longer producing learning. Why does this happen?

The most widely accepted explanation is that, as conditioning proceeds, presentations of the US lose their effectiveness. We know from a number of research studies that, during learning, the formation of a CS (light)–US (food) association allows presentation of the CS to evoke activity in the US representation before the US occurs. To adopt the terms used by the influential theorist Wagner (e.g. 1981), the CS induces a state of secondary activation in the US representation (as opposed to the primary

Pioneers

A.R. Wagner (1934–) and **R.A. Rescorla** (1940–) carried out research at Yale University in the late 1960s. Their experiments showed that simple contiguity of the CS and US is not sufficient to produce conditioning, and that it is also necessary for the CS to provide information about the likely occurrence of the US. (The phenomenon of blocking, described here, is an example.) The theoretical model they devised to explain this effect (published in 1972) was able to deal with a wide range of learning phenomena and set the agenda for almost all the research that has been done on associative learning mechanisms since then. Although the details of the Rescorla–Wagner model have been much debated, its central principles have been adopted by a wide range of associative (or 'connectionist') theorists attempting to explain not only simple learning processes, but human cognition in general.

activation produced by the US itself). Wagner proposes that secondary activation is not capable of supporting association formation; furthermore, it stops the US (food) from evoking the primary state of activation. The result is that the US becomes less effective as learning proceeds. As the CS–US link grows stronger, Wagner proposes that the CS (light) becomes more effective at producing the secondary state of activation and the US (food) becomes less able to produce the primary state necessary for further strengthening to occur.

So, while contiguity is important for learning, its nature needs precise specification. The events that must occur together are not so much the CS and US per se as the primary states of activation of their central representations.

Blocking – failure to learn

The phenomenon of *blocking* provides an interesting and much-studied instance of failure to learn, in spite of contiguous presentations of CS and US.

In a blocking experiment, animals receive training with what is termed a compound CS (Phase 2) – in this example (table 4.3) represented by the simultaneous presentation of a noise and a light followed by a shock reinforcer. However, the experimental group has first received a phase of training in which the noise alone is conditioned (Phase 1). The performance of the control group of participants shows that training (Phase 2) with a compound CS is normally sufficient to establish associations between individual CS elements (noise, light) and the US (shock). So in this control group the light, when subsequently presented on its own, will

blocking training an organism with one stimulus as a signal for an unconditioned stimulus to prevent the organism from learning about a second stimulus when both stimuli are subsequently presented together as signals for the same unconditioned stimulus

Table 4.3 Design and results of an experiment by Kamin (1969) on blocking.

Group	Phase 1	Phase 2	Test light
Experimental	Noise → shock	Light + Noise → shock	No CR
Control	–	Light + Noise → shock	CR

evoke a CR. But the experimental group shows no (or very little) evidence of learning about the light in Phase 2. Although they have received light–US pairings, just as the control participants have, in Phase 2, the formation of the light–US association appears to have been blocked by initial training with the noise in Phase 1.

A possible explanation of the blocking effect links directly to the asymptote phenomenon. Recall that a US representation in a secondary state of activation will not support association formation. In our blocking experiment, Phase 1 training for the experimental group establishes the noise as a CS, enabling it to activate the US representation in a secondary state of activation. So for these participants, during Phase 2, the presentation of the US will not be able to produce the state of primary activation, which means that the light introduced as part of the CS at this stage of testing will be unable to acquire associative strength.

Predictive power

Blocking has been of special interest not just because it provides an example of the failure of the contiguity principle, but also because it seems to demonstrate the operation of another principle. Animals in the experimental condition learn well about an event with predictive power (the noise in the first stage of training predicts that the US will shortly occur), but they do not learn about an uninformative event (the added light in Phase 2 supplies no added information). The principle here is that conditioning occurs only to a CS that gives information about the likely occurrence of a succeeding event – i.e. what we might term a predictive CS.

SELECTIVE ASSOCIATION FORMATION

A further challenge to the principle of contiguity came in the 1960s when psychologists began to realize that the principle might apply only to certain pairings of events. They had long suspected that some associations might form more readily than others, but they were usually able to find reasons to dismiss their worries. For example, when attempts to replicate Watson's demonstration of emotional conditioning in infants proved unsuccessful when an inanimate object, rather than a live animal, was used as the CS, researchers suggested that the CS was simply not salient enough to be noticed.

But an important experiment by Garcia and Koelling (1966) showed selectivity in association formation that could not be easily explained away. Their study demonstrates the phenomenon

of *preparedness*. The rats in this study appeared to be prepared to associate external cues with painful consequences and to associate illness with

taste cues. But taste did not become readily associated with shock, nor external cues with illness. The usefulness to the rat of having a learning system that operates in this way should be clear; after all, gastric illness is more likely to be caused by something the rat ate than something it heard or saw. But to the psychologist investigating general laws of learning, the preparedness effect constitutes a problem in need of explanation.

One possibility is that a *principle of similarity* operates in conditioning. By this principle, not only should the events to be associated occur together, but if learning is to take place they should also be similar to one another. Applying this principle to the Garcia and Koelling result, a taste and an illness might be readily associated because they are similar in that both are detected by receptors (called interoceptors) concerned with the animal's internal environment. External cues, on the other hand, have little in common with an internal state, making it difficult to associate auditory and visual events with illness.

Compared with the massive amount of experimental effort that has been expended on establishing the finer points of the contiguity principle, investigation of the similarity principle has been almost totally neglected. Perhaps we will see more studies in this area before too long.

preparedness tendency of certain combinations of events to form associations more readily than others

principle of similarity suggestion that association formation occurs particularly readily when the events are similar to one another

NON-ASSOCIATIVE LEARNING

Laboratory studies of learning have concentrated on conditioning procedures in which the participants experience two events (two stimuli, or a response and a stimulus) in close contiguity. It is hardly surprising, therefore, that association between events has proved so dominant in theories of learning. This approach has been justified by the assumption that the complex instances of learning shown in our everyday behaviour may well be governed by associative principles.

But this should not blind us to the fact that learning can also result from procedures in which there is no intentional pairing of two events.

RESPONDING TO A SINGLE STIMULUS

Repeated presentation of a stimulus that elicits a particular UR will result in *habituation* – a gradual reduction in the magnitude of the response. A good instance in vertebrates is the startle response produced

habituation waning of the unconditioned response with repeated presentation of the eliciting stimulus

Research close-up 2

Selectivity in aversive conditioning

The research issue

In laboratory studies of conditioning, participants and procedures have generally been chosen on the basis of convenience. The laboratory rat, for instance, is easy to obtain, cheap to keep and, when properly treated, easy to handle. Pairings of tone and shock or flavour with nausea will readily establish conditioned responses that are easy to observe. Although these kinds of studies are rather artificial, researchers have assumed that the results obtained from them reveal general principles about the nature of association formation, which apply to other species and other stimuli.

The experiment by Garcia and Koelling (1966) presented an important challenge to this assumption by showing, for laboratory rats, that animals appear to be especially 'prepared' to associate some combinations of events and to have difficulty in forming associations between other combinations. Similar principles of preparedness may well apply to humans.

Design and procedure

Rats were allowed to drink a saccharin-flavoured solution while a light and noise were being presented: each lick at the drinking tube closed a circuit that produced a flash of light and a click. So they experienced a compound CS comprising a taste element and an auditory-visual element. Some rats then received a nausea-inducing injection of lithium chloride (LiCl) as the US; other rats received an electric shock to the feet as an aversive US. Both groups of rats showed a reduction in willingness to drink over the course of several training sessions. In the final test, animals received either access to the saccharin solution in the absence of the auditory-visual cue, or access to unflavoured water but with the auditory-visual cue still being presented.

Results and implications

The results of the test phase are presented in figure 4.9. Animals given LiCl as the US showed an aversion to saccharin but were willing to drink plain water even when it was accompanied by the light and the click. Animals given shock as the US drank saccharin readily but shunned the 'bright noisy' water. Evidently, the events used as CSs and USs in this experiment are capable of entering into associations but show a certain 'choosiness' about which other events they are prepared to become linked to: the LiCl as an aversive US became linked to the taste of the saccharin water, and the foot shock as an aversive US became linked to the light and click, but not vice versa.

This result, and others like it, have led some researchers to suggest that we might be misguided in our attempt to establish general laws of learning. How can we retain contiguity as a general principle when the pairing of two events sometimes results in association formation and sometimes does not? An alternative view is that the preparedness effect reflects the operation of a further, previously unsuspected, general principle that must be added to contiguity in order to produce a comprehensive theoretical account.

Garcia, J., & Koelling, R.A., 1966, 'The relation of cue to consequence in avoidance learning', *Psychonomic Science*, 5, 123-4.

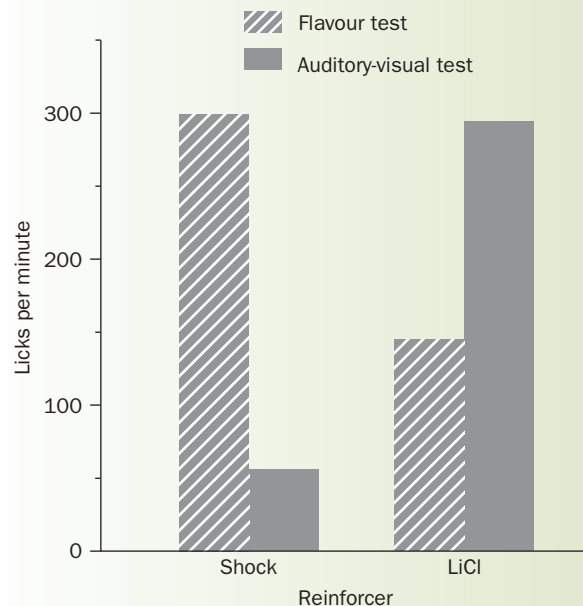


Figure 4.9

Drinking by rats of flavoured water or water associated with auditory/visual cues. In previous training the rats had experienced the flavour and auditory-visual cues together followed, for some animals, by a shock reinforcer, and for other animals by a nausea-inducing injection of lithium chloride (LiCl). Rats conditioned with illness as the reinforcer show a stronger aversion to the flavour cue; rats conditioned with shock show a stronger aversion to the auditory-visual cues. Source: Adapted from Garcia and Koelling (1966).

by a sudden loud noise, a response that reliably declines if the noise is regularly repeated. See figure 4.10, which also shows the

dishabituation restoration of a habituated response by presentation of a strong extraneous stimulus

phenomenon of *dishabituation*, whereby the response returns when a salient extraneous stimulus (e.g. a flashing light) is presented just before a trial with the habituated noise.

The observation that the response can be easily restored in this way shows that habituation is not solely a matter of sensory or motor fatigue – it is a genuine case of learning. And since habituation occurs as a consequence of the presentation of a single event, it is difficult to interpret this form of learning in terms of association formation. The most likely explanation, at least for simple instances of the phenomenon, is that changes occur in the neuronal pathway connecting the S and R that make transmission of nervous impulses less likely to occur.

A series of elegant neurophysiological studies by Kandel and colleagues (e.g. Kandel, 1979) using the marine mollusc *Aplysia* (see figure 4.11) has gone some way towards establishing which synaptic connection loses effectiveness during habituation, and the biochemical basis of this loss. (For this work Kandel was awarded the Nobel prize for medicine.)

imprinting the development of filial responses by newly hatched birds to an object (usually the mother) experienced early in life, or more generally the early formation of social attachments in animals

Loss of the UR is not the only effect produced by stimulus exposure. Consider the phenomenon of *imprinting*,

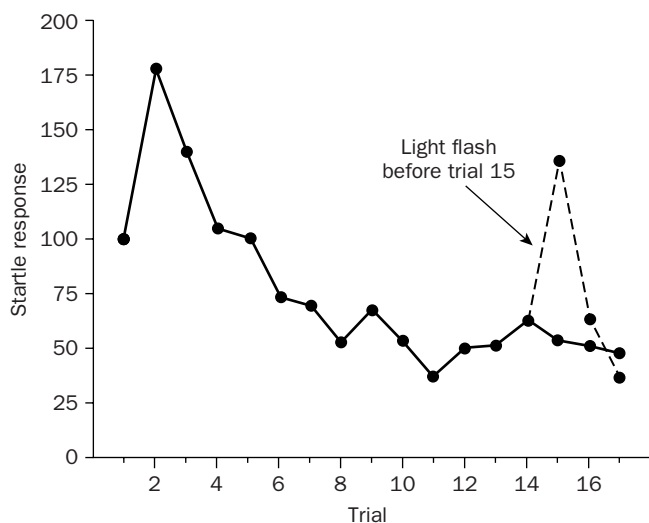


Figure 4.10

Habituation of the startle response of rats to an auditory stimulus. The response magnitude is expressed with respect to the initial level, which is given a score of 100. For half the animals a light flash was presented before trial 15 resulting in a temporary recovery of the startle response (dishabituation). Source: Based on Groves and Thompson (1970).



Figure 4.11

Perhaps surprisingly, studies of molluscs have helped researchers find out more about how the brain works.

in which a chick becomes attached to a conspicuous object experienced early in life. This behaviour pattern is found only in some species, but other features of the imprinting process appear to be more general. Most animals exposed to complex objects are able to learn the characteristics of the object, and subsequently to distinguish more easily the object from other similar things. This phenomenon is known as *perceptual learning*. The nature of the mechanism responsible for it is not fully known, but it seems likely that associative processes are involved, in that learning the characteristics of a complex object involves learning that its various features go together. This is achieved by the formation of associative links among its component parts.

perceptual learning exposure to events, increasing subsequent ability to discriminate between them

The perceptual learning process, which enables the animal to build up an accurate representation of the stimulus, probably plays a role in some instances of habituation. When animals are habituated to a complex event, the response can be restored if some element of that complex is omitted or changed. Such dishabituation occurs, it has been suggested (Sokolov, 1963), because animals are sensitive to any mismatch between incoming stimulation and the central representations of events they have already experienced.

SPATIAL LEARNING

Figure 4.12 shows apparatus used to assess rats' ability to learn about the layout of their environment. A rat is set to swim in a pool of water from which it naturally wants to escape. It can do this by reaching a small platform, which is just below the surface and not visible to the animal (because of the 'milky' of the water). Finding the platform on the first trial is a matter of




Figure 4.12

Swimming pool apparatus introduced by Morris (1981). Although it cannot see, smell or hear the target, the rat learns to swim straight to the small, just-submerged, platform from wherever it is put into the pool.

chance, but with further training the rat rapidly learns to swim straight to the platform. Since the rat cannot see the platform to home in on it, how can it be performing this feat?

One obvious possibility is that the rat learns to swim in the general direction of some feature of the room outside the pool, which lies on a continuation of the line between its starting point and the platform. But this cannot be the whole story, because in other trials, rats were put back in the pool at a different starting position from that used in training. Figure 4.13 shows the paths that the rats followed. Clearly, in these trials, following a line to an extra-pool cue would not work. However, as the results show, even under these conditions the rats were still very good at finding the platform.

To explain this in terms of standard conditioning processes, we must assume that the rat learns to approach not single cues, but

complex configurations of cues. We know from other training procedures that rats can learn about combined (often referred to as configural) cues. But such learning tends to occur painfully slowly, whereas spatial tasks are mastered much more easily by rats. This suggests that spatial learning operates according to principles quite different from those that underlie classical and instrumental conditioning procedures. It is possible that exposure to an environment allows the animal to form a *cognitive map* of that environment – some sort of internal representation of the spatial relationships among the cues it has experienced. The animal is then able to navigate because it knows its own position with respect to this internal representation. But no one has yet supplied a full account of the process by which the map is constructed, how the animal knows its own position, and so on.

cognitive map postulated internalized representation of the layout of the environment in which information about the relative spatial relationships of various features is preserved

DISCRIMINATION LEARNING

In a discrimination learning task the animal is presented with two stimuli (sometimes more) that are associated with different outcomes. For example, a pigeon might be presented with a choice between two discs, one coloured red and the other green; pecking at the green disc will produce food, but pecking at the red disc will not. The pigeon will solve this problem, coming reliably to choose the green disc after a few dozen trials. Its ability to do this task is no puzzle and can be fully explained in terms of standard conditioning processes.

More intriguing is the fact that training on such a task will transfer to other similar tasks. If the pigeon is now asked to solve a similar discrimination problem, in which the choice is between blue and yellow discs, learning can occur very rapidly: we call this *positive transfer*. The original associations involving red and green

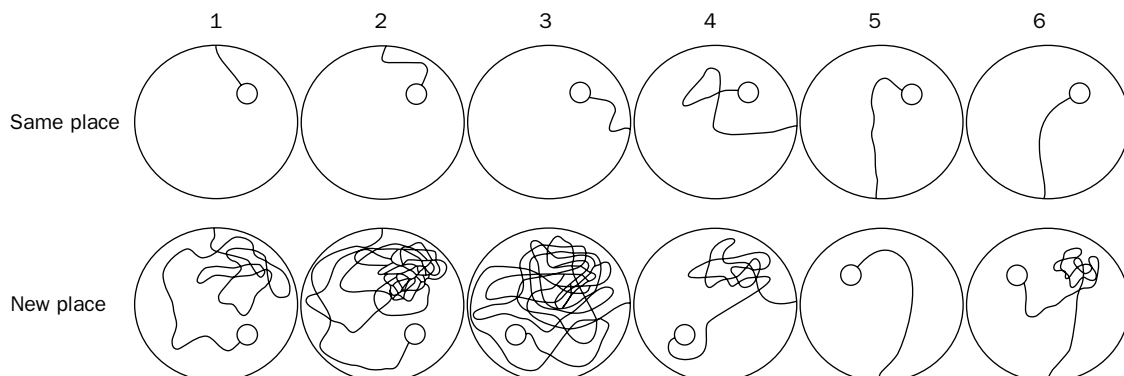


Figure 4.13

Plan view of the paths taken by rats on test trials in the swimming pool of Figure 4.12. The top row shows the performance of six rats swimming from a novel starting point to a platform remaining in the same place as was used in training. The second row shows performance for rats required to swim to the platform in a new place in the pool.

are clearly irrelevant to this new discrimination task, so the transfer must have some other source. The pigeon appears to have acquired a fairly abstract concept in the course of acquiring the first discrimination – something along the lines of ‘differences in colour are important and should be attended to’.

Studies involving primates have produced more dramatic examples of abstraction. In the *learning-set* procedure (first introduced by Harlow, 1949), a rhesus monkey is presented with two objects and given a small amount of food for picking up one of them. After six trials the original objects are replaced with two new ones and, again, responding to only one of the objects was rewarded. After six trials on this new problem, the objects were again changed, and so on for many, many pairs of objects.

Early in training performance is unremarkable, six trials being insufficient for the monkey to solve the problem. But as training proceeds, performance begins to improve, until finally it is as near perfect as it can be (see figure 4.14). After training on hundreds of these problems, the monkey is able to solve a new problem with no more than a single error, switching its choice to the other object if its first choice is wrong, but staying with its original choice if this proves correct. By experiencing many problems of a similar type, the animal appears to abstract some general rule about how to behave in this situation – a rule that allows the near-instantaneous solution of a problem that it had, in fact, never faced before.

The rule that operates in this case is the *win-stay, lose-shift* strategy: in other words, the animal learns to persist with a choice that yields food, but shift to the other object if it does not. Associative theory can go some way towards explaining this. The occurrence of reward (or non-reward) can be regarded as a stimulus that, like any other, can enter into associations or acquire discriminative control over an instrumental action. The special feature of the learning-set procedure is that these stimuli and associations come to dominate the animal’s behaviour to the exclusion of all others.

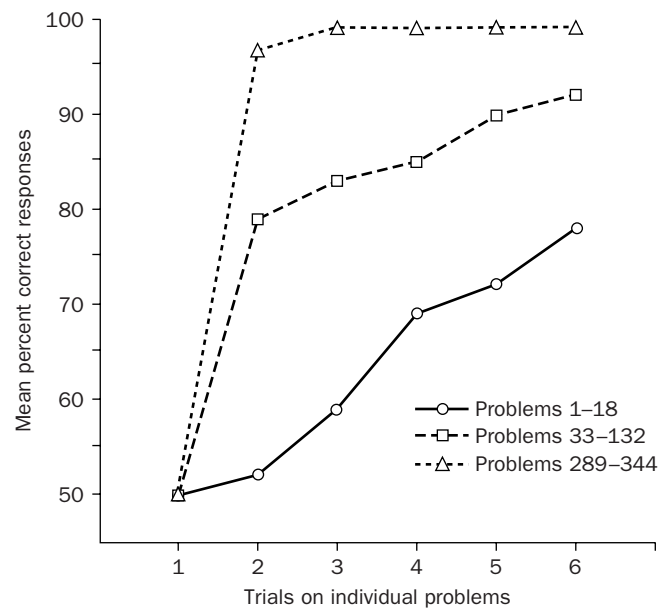


Figure 4.14

Performance of monkeys over a series of 344 six-trial discriminations. On each new problem, performance starts at chance, but the rate at which learning occurs is more rapid for problems encountered late in the series than for those encountered earlier. Source: Adapted from Harlow (1949).

So the animal learns to focus on classes of cues that are accurate predictors of reward and to ignore others that are not. Intensive research is currently going into the nature of such higher-level learning processes that might modulate the mechanisms of simpler associative processes.

FINAL THOUGHTS

At the start of this chapter we said that the aim of the psychologist whose focus of interest is learning was to discover general principles that govern the way in which organisms interact with their environment and become changed by the experience.

What has emerged from this brief survey is the central notion of association. The principle of association provides a useful theoretical underpinning for the outcome of many learning studies. Building on this foundation, researchers in the future will need to:

1. continue the detailed study of conditioning procedures to refine our understanding of the laws of association (e.g. contiguity) and, perhaps, to detail the need for other principles (e.g. similarity); and
2. extend the study of more complex forms of learning in order to assess how far these can be understood in terms of underlying associative processes, and, where they cannot, to specify what further principles should be introduced to supplement the principle of association.

Ethical issues arising from research into learning include the implications of emotional states being subject to classical conditioning, and the use of punishment to eliminate unwanted behaviour. Is it really appropriate to call behaviour ‘voluntary’ if it is in fact controlled by a response–outcome association?

Summary

- Learning is defined as the process whereby an organism interacts with its environment and becomes changed by the experience so that its subsequent behaviour is modified. Note that we infer that learning has occurred through our observations of changes in behaviour.
- The basic principles of learning have been established through laboratory studies of animals but are also applicable to humans. Indeed, these basic principles have been applied to the analysis of human conditions such as obsessive–compulsive disorder.
- Classical conditioning reflects the formation of stimulus–stimulus associations. Such associations constitute the main way in which an organism represents information about the relations between environmental events; and they can endow previously neutral events with emotional significance.
- In instrumental learning (the other major form of conditioning), an association is formed between a response and its consequences. When the consequences of the response are pleasant, the likelihood of the response will increase; when the consequences are unpleasant, the likelihood will decrease (the law of effect).
- Behaviour controlled by such associations may be described as voluntary or goal-directed. For example, when in our car we may learn that the response of pressing our foot on the car accelerator when the traffic lights are green results in a positive outcome (increasing the likelihood of this behaviour). By contrast, performing this same response when the traffic lights are amber is likely to result in an unpleasant outcome, thereby decreasing the likelihood of this response in the future.
- Conditioning procedures are used to investigate the laws of association. They have shown that the co-occurrence of the events to be associated is important (principle of contiguity) but also that associations may fail to form unless one event supplies information about the occurrence of the other (principle of predictiveness).
- The associative principle has been thoroughly tested and shown to have wide relevance. It has difficulty, however, in explaining some examples of complex (e.g. the learning of abstract concepts or rules) or very simple (e.g. habituation) forms of learning.

REVISION QUESTIONS

1. What is the defining feature of the procedure used to establish classical conditioning?
2. What is the nature of the association formed as a result of the classical conditioning procedure?
3. What role does classical conditioning play in psychological phenomena observed outside the conditioning laboratory?
4. What feature is common to the various procedures used in the study of instrumental learning?
5. What is the nature of the association formed as a result of instrumental training procedures?
6. What factors determine when instrumental learning will be expressed in behaviour?
7. How has the laboratory study of instrumental learning helped enhance our understanding of normal and abnormal human behaviour?
8. What determines whether or not the pairing of two events will result in association formation?
9. What forms of animal learning cannot be explained in terms of association formation?

FURTHER READING

Dickinson, A. (1980). *Contemporary Animal Learning Theory*. Cambridge: Cambridge University Press. Focuses on the importance of association formation in learning.

Domjan, M. (2003). *Principles of Learning and Behaviour*. Belmont, CA: Thomson/Wadsworth. A wide-ranging textbook, which covers all the main areas in the study of learning.

Mackintosh, N.J. (ed.) (1994). *Handbook of Perception and Cognition: Animal Learning and Cognition*. San Diego, CA: Academic Press. An advanced and demanding text, which nevertheless provides state-of-the-art expositions by specialist contributors on most of the topics covered in this chapter.

O'Donohue, W. (1998). *Learning and Behavior Therapy*. Boston, MA: Allyn and Bacon.

An up-to-date exposition of how studies of animal learning can help us to understand human fears and anxiety (and also help in devising treatments for anxiety disorders).

Pearce, J.M. (1997). *Animal Learning and Cognition: An Introduction*. Hove: Psychology Press.

An introductory textbook covering in more detail and with great clarity all the phenomena dealt with in this chapter. It also includes material on other cognitive processes that have been studied in animals, such as navigation, communication and social learning.

Shanks, D.R. (1995). *The Psychology of Associative Learning*. Cambridge: Cambridge University Press.

Describes in detail how the simple associative notions outlined in this chapter have been developed to explain a wide range of human cognitive processes.

Shettleworth, S.J. (1998). *Cognition, Evolution, and Behavior*. New York: Oxford University Press.

The evolutionary perspective adopted by this book makes an interesting contrast to the approach adopted in this chapter.

Contributing author:
Geoffrey Hall

