

OPERANT BEHAVIOR

Some behavior makes sense in terms of the events that precede it; other behavior makes more sense in terms of the events that follow it. Reflexes are behavior of the first kind. The type and vigor of the reflex response are closely related to the type and intensity of the eliciting stimulus. Kineses are behavior of the second kind. The movement of the orienting bacterium from moment to moment is unrelated to any single stimulus, yet its behavior as a whole can be understood in relation to the prevailing chemical gradient: The organisms aggregate in high concentrations of an attractant and disperse away from high concentrations of a noxious substance.

Behavior of this type is *guided by its consequences*: Under normal conditions, the location of the bacterium is determined by the chemical gradients in its environment. Because the antecedent (proximal) causes of its behavior are many but the final cause is one, we can most simply describe the behavior in terms of its outcome: The bug finds the food.

Behavior guided by its consequences was called by Skinner *operant behavior*, and the term has become generally accepted. The word *operant* refers to an essential property of goal-directed behavior: that it have some effect on the environment. If the bacterium cannot move, or if movement has no effect on the organism's chemical environment, then its behavior will not appear to be guided by a goal.

Skinner was interested in learning and restricted his definition to operant behavior that reflects learning. It is useful to have a term for the wider concept, however, and so I use operant behavior to refer to any behavior that is guided by its consequences.¹

The consequence that guides operant behavior plays the same role as the set point of a feedback system. If the set point is external, as in a tracking system, then the position of the animal will appear to be controlled by the external reference: As the target moves, so does the animal. If tracking is very good, the illusion that the animal is physically attached to the target can be very strong. The term *control* is sometimes used to refer to operant behavior also, as behavior that is "controlled by" its consequences. The meaning is the same. In both cases, the behavior is guided in some way by the discrepancy between the present state of affairs and some ideal state.

In this chapter I provide the outline for a comparative psychology of learning. I discuss the kinds of mechanism an animal must possess for it to be capable of operant behavior and look at differences between the operant behavior of simple animals and the operant behavior of mammals and birds. This comparison suggests that a major difference between "simple" and "complex" animals is in the number of different "states of the world" they can distinguish. The middle part of the chapter is an overview of classical and operant conditioning and related effects, such as sensitization and pseudoconditioning. Conditioning is taken up again in more detail in the next chapter. The last section of the chapter deals with the concepts of *stimulus*, *response*, and *internal state*, and their mutual relations. This discussion leads to formal definitions for each of these terms. The chapter ends with a discussion of *learning* and *memory* and how they relate to behavioral mechanisms.

CAUSAL AND FUNCTIONAL ANALYSIS OF OPERANT BEHAVIOR

How does the operant behavior of mammals differ from the kinesic behavior of protozoa? There are, of course, trivial differences: Protozoans are small and mammals are big; protozoans can move about but can't do much else, whereas mammals can climb and fly and talk and press levers. These are the same sorts of difference that exist among different mammal species; they have nothing to do with the essential feature of operant behavior, which is its goal-directedness. Yet there are differences. One way to get at them is to look carefully at the operant behavior of a

simple animal and consider how it differs from the operant behavior of a mammal or a bird. This comparison can also tell us something about the mechanisms that must be involved in learning.

One of the most beautiful examples of intelligent behavior by protozoa is provided by H. S. Jennings' account of how *Stentor*, a single-celled pond animal (see Figure 4.1), copes with a mildly irritating substance introduced into the water in its vicinity. *Stentor* feeds by drawing water into the opening at the top of the "trumpet." The water currents are created by the beating of cilia that fringe the opening. Jennings describes the effect of the slightly noxious carmine particles on the animal's behavior as follows:

Let us now examine the behavior under conditions which are harmless when acting for a short time, but which, when continued, do interfere with the normal functions. Such conditions may be produced by bringing a large quantity of fine particles such as India ink or carmine, by means of a capillary pipette, into the water currents which are carried to the disk of *Stentor*.

Under these conditions the normal movements are at first not changed. The particles of carmine are taken into the pouch and into the mouth, whence they pass into the internal protoplasm. If the cloud of particles is very dense, or if it is accompanied by a slight chemical stimulus, as is usually the case with carmine grains, this behavior lasts but a short time; then a definite reaction supervenes. The animal bends to one side...It thus as a rule avoids the cloud of particles, unless the latter is very large. This simple method of reaction turns out to be more effective in getting rid of stimuli of all sorts than might be expected. If the first reaction is not successful, it is usually repeated one or more times.

If the repeated turning to one side does not relieve the animal, so that the particles of carmine continue to come in a dense cloud, another reaction is tried. The ciliary movement is suddenly reversed in direction, so that the particles against the disk and in the pouch are thrown off. The water current is driven away from the disk instead of toward it. This lasts but an instant, then the current is continued in the usual way. If the particles continue to come, the reversal is repeated two or three times in rapid succession. If this fails to relieve the organism, the next reaction — contraction — usually supervenes.

Sometimes the reversal of the current takes place before the turning away described first; but usually the two reactions are tried in the order we have given.

If the *Stentor* does not get rid of the stimulation in either of the ways just described, it contracts into its tube. In this way it, of course, escapes the stimulation completely, but at the expense of suspending its activity and losing all opportunity to obtain food. The animal usually remains in the tube about half a minute, then extends. When its body has reached about two-thirds original length, the ciliary disk begins to unfold and the cilia to act, causing currents of water to reach the disk, as before.

We have now reached a specially interesting point in the experiment. Suppose that the water currents again bring the carmine grains. The stimulus and all the external conditions are the same as they were at the beginning. Will the *Stentor* behave as it did at the beginning? Will it at first not react, then bend to one side, then reverse the current, then contract, passing anew through the whole series of reactions? Or shall we find that it has become changed by the experiences it has passed through [my italics], so that it will now contract again into its tube as soon as stimulated?

We find the latter to be the case. As soon as the carmine again reaches its disk, it at once contracts again. This may be repeated many times, as often as the particles come to the disk, for ten or fifteen minutes. Now the animal after each contraction stays a little longer in the tube than it did at first. Finally it ceases to extend, but contracts repeatedly and violently while still enclosed in its tube. In this way the attachment of its foot to the object on which it is situated is broken, and the animal is free. Now it leaves its tube and swims away. In leaving the tube it may swim forward out of the anterior end of the tube; but if this brings it into the region of the cloud of carmine, it often forces its way backward through the substance of the




Figure 4.1. *Stentor roeselii* attached to its tube and drawing a cloud of carmine particles into the ciliary disc as part of its normal feeding pattern. (From Jennings, 1906, Figure 111.)

tube, and thus gains the outside. Here it swims away, to form a new tube elsewhere (Jennings, 1906, pp. 174-5).

The behavior of *Stentor* as Jennings describes it is marvelously adaptive. How might we explain it, both mechanistically and from a functional point of view?

The animal has four levels of response to escape from the carmine, each more costly (but also more likely to be effective) than the preceding one: turning away (T) uses little energy and doesn't interfere with feeding (F); ciliary reversal (R) uses little energy, but is an interruption of feeding; contracting into the tube and waiting for a while (CW) is energetic and seriously interferes with feeding; breaking away (B) is most energetic of all, and means abandoning a known site. We don't know what causes the animal to shift from one mode of behavior to another. Jennings is typically cautious: "...shall we find that it has become changed by the experiences it has passed through...?" He avoids saying what aspect of the animal's past experience might be responsible for the change from CW to B, the most drastic change in the sequence.

There are two simple ways to explain the succession of avoidance reactions. One way is illustrated in Figure 4.2. The figure shows on the left the sequence of five activities, and on the right the sequence of decisions that allows each activity either to be repeated, or to be followed

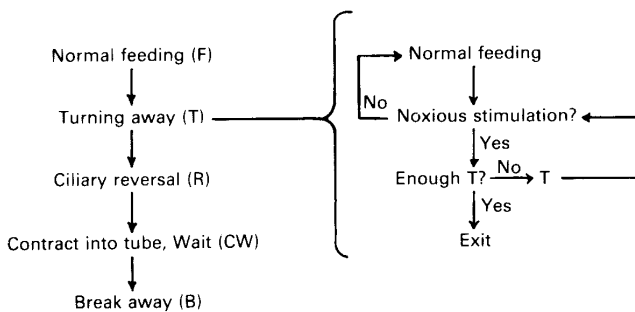


Figure 4.2. Chain-reflex model for *Stentor* avoidance behavior. Left: feeding and the four types of avoidance reaction, in the order they normally occur. Right: a program for each link in the chain.

by the next. The basic idea is that each activity is repeated for a while ("Enough?") and then, if the noxious stimulation persists, the next avoidance response is tried. "Enough" might correspond to number of repetitions, time spent in the activity, or some combination of time and number.

The model in Figure 4.2 is an example of a reflex chain: Each activity is the cause of the next one in the sequence. Chains are the first explanation to come to mind when one ponders how a simple animal might put together a sequence of

acts. They were very popular with early stimulus-response psychologists. But chains are a rather rigid form of behavioral organization. If one activity fails to occur or is blocked in some way, the next activity cannot occur, for example. Perhaps for this reason, chain reflexes are rarely found, even in something like walking, where it might seem reasonable that the movement of one limb should precipitate the movement of the next.² Moreover, a problem with the usual chain account is that the order of occurrence of activities is rigidly fixed whereas, as Jennings points, out the animal "has been changed by the experiences it has passed through" so that even though all external conditions are the same, the behavior may differ.




Figure 4.3. A temporal decision rule for *Stentor* avoidance.

behavior may differ.

The account in Figure 4.2 might also be termed *cognitive*, or even *symbolic* (cf. Gallistel & Gibbon, 2002), because it uses if-then logic. But there is an alternative way

to organize avoidance behavior that is not symbolic at all, and also allows for variability in the temporal sequence — and is even simpler than the chain described in Figure 4.2. The model is not logical, but *temporal*. The rationale is as follows: Risk to *Stentor* is directly related to the *time* when the carmine is present. The four avoidance reactions are progressively more costly. The animal needs a rule that will match the cost of the avoidance response to the risk. One way

to determine which behavior should occur, therefore, is simply to identify three cutoff times, t_T , t_R , and t_{CW} , such that $t_T < t_R < t_{CW}$. Then activity R (say) occurs if $t_T < t < t_R$, and so on, where t is the time since the beginning of the episode. The temporal rule describing the animal's behavior is shown graphically in Figure 4.3. Thus, an internal clock that starts running with the onset of the carmine, with settings corresponding to changes in behavior. $T \rightarrow R, R \rightarrow CW$, and so forth, is an alternative method of generating the observed sequence of avoidance reactions¹.

Noxious substances come in different concentrations, and last for different lengths of time. Risk is directly related to concentration as well as time. The scheme in Figure 4.3 can easily be generalized to two or more risk factors. For example, suppose that risk (H) is proportional to the product of concentration (C) multiplied by time: $H = tC$. Then the animal simply has to identify three values of H - H_T, H_R , etc. — that define transitions from one behavior to the next, as before. The whole thing can be represented in a two-dimensional space (rather than the one-dimensional space of Figure 4.3), as shown in Figure 4.4. Obviously, the scheme is not restricted to a product decision rule or to two dimensions: Any function, $H = f(t, C)$, can be used to partition the space into regions associated with different activities, and the space can have as many dimensions as there are risk factors.




Figure 4.4. A time-concentration decision rule for *Stentor* avoidance.

tion the space into regions associated with different activities, and the space can have as many dimensions as there are risk factors.

The one-dimensional, temporal decision rule suggests an obvious mechanism for the *Stentor* behavior — a clock. But the two-dimensional case makes it clear that an explanation by decision rules is functional, not mechanistic. It is not at all obvious what mechanism the animal uses to compute the relative importance of time and concentration in deciding on an avoidance reaction. Yet in this relatively simple case it is easy to imagine possible mechanisms. For example, suppose that the threshold concentrations at which

each of the four reaction occur are ranked according to their cost $C_T < C_R < C_{CW} < C_B$, and that the reactions are mutually exclusive (see the discussion of thresholds and reflex competition in Chapter 2). Then reaction T , with the lowest threshold, will occur first, blocking the others. Suppose further that the threshold of any reaction *increases* as the reaction continues to occur (habituation). This mechanism will ensure that eventually it will be supplanted by R , as the threshold for T increases owing to habituation. The same process repeated leads to the transition from R to CW , and so on. Given appropriate choices for time constants and inhibition relationships, a plausible mechanism for *Stentor* avoidance might be arrived at (see note 2).

This mechanism is not the only possible one. Even if we restrict ourselves just to processes involving time-varying thresholds, several mechanisms for the *Stentor* avoidance reaction might be invented. To discriminate among them would require careful behavioral and perhaps even physiological experiments. To eliminate the chain-reflex possibility, some way would have to be found to prevent one of the reactions. If reaction R can occur without preceding reaction T , then a chain is ruled out. Since this apparently occurs once in a while without any intervention, the chain reflex idea is unlikely from the outset. To discriminate among other possibilities, we need to know the dynamic properties of the system: How quickly do the activities succeed one another in a constant carmine concentration? What is the response to a step change and to alternations between different concentrations? and so on (see the Moorhouse et al. experiment dis-

¹ A dynamic model consisting of competing leaky integrators can simulate the sequence of avoidance activities and also allow for changes in sequence of the sort that Jennings observed (Staddon, 2001b, Chapter 4).

cussed in Chapter 2). Whatever the answer, it will have to be compatible with the very simple, static functional account.

The functional account is more general than the particular mechanism the underlies it, because it will apply not just to *Stentor*, but to any small animal similarly situated — even though different species (and perhaps even different individual *Stentors*) might well do the job in different ways. If we can think more than one way to solve a behavioral problem, there will often be comparable diversity in nature's solutions.

An animal's ability to behave according to the prescriptions of a functional account depends upon particular mechanisms — here of the reflex type, but in other cases, more complicated. For some purposes the functional account is more useful, for others we might like to know more of the details of how the animal does it. For example, if we are interested in how an animal's behavior relates to its niche, the functional account is better; but if we want to get at the physiological basis for behavior, the mechanistic account is more useful.

Obviously, the more intelligent the animal, the more possible mechanisms may underlie a given example of adaptive behavior. Consequently, when we are dealing with mammals and birds, functional or descriptive theories generally do much better than mechanistic ones.

Operant behavior and learning

How does the avoidance behavior of *Stentor* stack up against the operant behavior of higher animals — with the behavior of a cat escaping from a puzzle box, for example, Thorndike's original learning situation?³ At one level, the similarity is great. In both cases, a change for the worse in the animal's environment causes a change in its behavior. Further changes continue to occur in a systematic way until the irritating condition is relieved (variation and selection). The increasing rapidity of contraction (CW and B) in response to the carmine resembles the increasing efficiency with which the cat escapes on successive trials. The cat's improvement would probably be called "learning": Is *Stentor* also "learning"?

The same external situation — carmine, plus normal feeding behavior by the animal — leads initially just to turning away, but later to breaking away. This means that there has been a change in *Stentor's* internal state. Such changes must have been occurring right along, otherwise the initial avoidance reaction, turning away, would have persisted indefinitely. Thus, the animal shows one essential feature of learning, a change in behavior potential as a result of experience. So far our *Stentor* is doing about as well as Thorndike's cat. Is there any difference?

Yes. The big difference is in the effect of either lapse of time or change in situation on the animal's behavior. The same *Stentor* exposed to carmine after 24 hours will almost certainly behave as it did when first exposed, not as it did at the end of the first exposure. A small lapse of time will have lesser effects on *Stentor's* behavior, but some new experience during that time — a new chemical stimulus, for example — is likely to abolish the effect of the earlier experience. But the cat that escapes efficiently from the box after a few successive trials will not be the same even after a week as it was when first put in the box. It may not escape quite as rapidly as it did on its last trial, but it will certainly do better than it did on its very first trial. The ability of the cat to respond correctly after lapse of time is not immune to the effects of intervening experience, but is likely to be much less affected than the behavior of *Stentor*. What does this difference between cat and protozoan mean?

The obvious answer is that the protozoan cannot remember things for as long as the cat. Leaving aside for a moment the problem of exactly what we mean by "remember," I think that this begs the question. The problem is: *Why* doesn't *Stentor* remember things as long as the cat? The difficulty is unlikely to be some general limitation on the persistence of behavioral changes in very small animals. It is in theory and in fact a simple matter to set a molecular "switch" that affects behavior after long delays. Even precisely timed delays are not rare in nature. Many very simple animals show periodicities that extend over very long times — circannual rhythms, the

periodicity of the 17-year locust, and so on. Nature has no problem in building clocks or retaining changes over long periods.

We can get an idea of the real difference between cat and *Stentor* by looking at what each creature *ought* to do in its respective problem situation — given what we can guess of their different discriminative abilities. The cat faces a much more clear-cut problem. The cat can discriminate puzzle boxes from the many other things with which she is familiar. When she senses the box again, she is unlikely to confuse it with anything else and has no reason to suppose that her previous solution is not appropriate. The *Stentor*, however, is unlikely to be able to discriminate carmine from many other chemical mixtures. Consequently, when it gets the carmine on a later occasion, it may simply be unable to identify it as the same event: Even if the carmine has exactly the same sensory effect on both occasions, the effect may be sufficiently similar to the effect of other chemicals, experienced before and after, that the *Stentor* may still be literally unable to identify its second carmine experience as a repeat of its first. Even a highly intelligent *Stentor* might be well advised to treat a second exposure (especially if much delayed after the first) as a completely new situation.

Stentor learning about carmine may be like a man with very poor vision meeting again someone he has met before: If he has to rely entirely on vision, he will necessarily be unsure whether he has really met the person before or not. No matter how smart he is; no matter how good his memory for faces; the blind man must necessarily be unsure because his vision gives him inadequate information. The poorly sighted person (like the hearing-impaired person) may appear dumb not because he cannot remember and understand, but because he cannot discriminate. In other words, *Stentor* probably hasn't sufficient information to justify repetition on the second carmine exposure of the behavior it had developed in response to the first exposure. *Stentor* may fail to learn, therefore, not because it cannot remember, but because it cannot be sufficiently sure when the same situation has recurred.

The abilities of animals are always in harmony with one another; an animal will not develop one ability to a high level if lack of another makes the first useless. An animal that can only move slowly will not usually have highly developed distance vision; a short-lived animal such as a butterfly may have little ability to repair injuries; a fish that lives in caves will often be blind. In similar fashion the persistence of memory is functionally related to the *number of*

		<i>Hedonic category</i>	
		Bad	Good
<i>Physical category</i>	A		
	B		

Table 4.1. Contingency table for a simple animal that divides all stimuli into two categories, A and B.

things that can be discriminated. This relation can be illustrated as follows. Imagine a very simple animal that can classify stimuli in just four ways: good-bad, and two other, "neutral" categories unrelated to good-bad (two categories of physical size, or intensity, for example). Every situation that the animal encounters must then fit into one of the four cells defined by this 2 x 2 table (Table 4.1). In order for it to be worthwhile for an animal to form an association between a given state of the world and its hedonic consequences, there must obviously be a real association. For example, if our animal can only partition neutral events into the two categories of "intense" and "weak," then there must really be some general correlation between the intensity of a physical event

and its hedonic consequences; loud noises must be generally associated with bad things and soft noises with good or neutral things, for example.

It is obvious that the more possible categories that animal has available, the higher the real correlation between events in a given category and their hedonic consequences is likely to be. For example, many loud noises are quite harmless, but loud, staccato noises coming from long objects are always potentially dangerous. The more qualifiers (additional categories) the

animal has available, the more accurately it can identify the real concomitants of good and bad. The more accurately the animal can identify signals for good and bad, the more worthwhile it becomes to remember them.

I conclude that the poor long-term memory of *Stentor* is much more likely to reflect a limitation on its ability to classify physical events into a number of categories than some limitation on storage or the persistence of physical change. The problem is an *information-processing* one. In order to react to carmine as it did on the first occasion *Stentor* must be able to identify the relevant features of its environment accurately enough to minimize the chance of the two kinds of possible error: reacting to the wrong stimulus and failing to react to the correct one. If its categories are so few⁴ that carmine does not uniquely occupy one of them, then it lacks the ability to tell whether a given chemical stimulus is carmine or something else, and so cannot detect the special properties of carmine presentations. Hence, if we want to show learning in simple organisms, we need to pick stimulus categories so separated that the animal is unlikely to confuse the experimental stimuli with others encountered naturally. The major difference between animals that can, and cannot, learn — between protists and pigeons — is in their ability to differentiate among states of nature: The difference is in *what they know*, rather than in how well they can remember.

Operant and classical conditioning

The basic operation in learning is using something as a signal for something else. The signal can be used in different ways, corresponding to habituation, sensitization, pseudoconditioning, and classical (Pavlovian) and operant (instrumental) conditioning. Sometimes only the latter two are classified as *learning*; this is the narrow sense of the term, but I will refer to all five types as *learning*, in the broad sense. I define these five types of learning now in a general way. The next chapter goes into more detail about particular operant and classical conditioning procedures. The mechanisms of classical and operant conditioning are discussed at length in the last two chapters.

Every stimulus has two aspects: good-bad (hedonic quality), and other, nonhedonic properties. For example, an electric shock is bad, but it also has a certain “tickling” quality, a certain duration, a limited spatial extent, and so on, and these things have no particular good-bad aspect to them. Similarly, food is good but it also has a certain taste, texture, temperature, and color. Sometimes the hedonic and nonhedonic aspects are distinct, but sometimes they are not. Spatial location rarely has any hedonic quality, for example, but the taste of a food and its value may be inseparable. In the simplest kinds of learning, the animal uses the nonhedonic characteristics of a stimulus as a cue for its hedonic qualities.

Habituation. As we saw in Chapter 2, habituation is the selective waning in strength of a response to repeated stimulation. For example, anything novel is potentially dangerous. Hence any sudden novel sound tends to be treated as potentially harmful: *Stentor* retreats into its tube; a rabbit pricks up its ears, turns toward the sound, and gets ready to bolt for its burrow. But as the sound is repeated and is not accompanied by any bad consequences, its nonhedonic properties, which at first were taken as potentially alarming, are eventually treated as a “safety” signal and the sound alarms no more. Habituation is, therefore, a kind of learning in which the nonhedonic aspects of a stimulus are treated as a signal for safety or neutrality. The stimulus is reclassified as “harmless.”

In higher animals, the effects of habituation may be long lasting or not, depending upon the stimulus. Wild birds may take alarm at a new feeding station, for example, but their alarm will habituate and will not reappear even after a few days away. On the other hand, the effects of habituation to a series of pistol shots will not persist unless they recur many times with no evil consequence. The more intense the stimulus, the more transient the habituation to it. In lower animals, the effects of habituation rarely persist.

Habituation can often be abolished by some new experience; this is *dishabituation*. For example, if an animal has been habituated to a series of loud noises and is now placed in a new situation, or presented with a bright light, another animal or some other striking stimulus, presentation of another loud noise will often alarm again.

Sensitization. This and the other three types of learning I will discuss all involve two kinds of stimulus. First, a stimulus with hedonic value and a stimulus with less or no hedonic value: a *reinforcer* or *unconditioned stimulus* (US) and a *conditioned stimulus* (CS). These terms derive from the use of hedonic stimuli such as food and electric shock in Pavlovian and instrumental conditioning situations. A noxious stimulus, such as an electric shock or a loud sound, will usually elicit a startle reaction from a human or an animal. A loud sound will startle, a weak sound will not; but a weak sound (CS) presented after one or two shocks (US) may startle again. This is not because the animal has learned anything about the relation of the shock and the sound: The sound is not a signal for the shock. Sensitization is a sort of *confusion* effect: The animal is anticipating the shock; sudden shock and sudden sound share nonhedonic qualities — so sound after shock is reacted to like shock.

The term *anticipation* here says rather more than necessary. All that is involved is that the shock changes the animal's state (lowers its reaction threshold) in such a way that another shock, or stimulus like shock, will be reacted to more strongly than if the first shock had not occurred. In the same way, *Stentor* after a second or two exposure to carmine reacted differently to further carmine: The initial few seconds of carmine exposure sensitized the animal so that further exposure elicited more and more extreme reactions. Only local memory is involved; a weak sound long delayed after shock will elicit no reaction, just as a second dose of carmine delayed after the first produces the same effect as the first rather than a stronger effect.

Any frightening situation sensitizes the animal; it is not necessary actually to present a noxious stimulus. For example, the following situation provides a good classroom demonstration of the human startle response: A student is invited to participate in front of the class in an experiment to measure skin resistance. He (or perhaps more usually, she) is asked to place two fingers on two metallic contacts protruding from an unfamiliar piece of obviously electrical apparatus.

At the instant that the finger touches the (completely harmless) contacts, the demonstrator sounds a loud siren. The siren produces no, or at most a weak, startle when presented without the buildup. When it coincides with what the student fears may be a painful, or at least a novel, experience, however, the reaction is vigorous.

Pseudoconditioning. This is very similar to sensitization. In both cases, prior presentation of a US causes a CS to elicit the same reaction. The effect is called *sensitization* if the CS at a stronger intensity can elicit the reaction (shock → startle; loud sound → startle; weak sound ≠ startle unless preceded by shock or loud sound). The effect is called *pseudoconditioning* if the CS never elicits the reaction on its own.

An experiment by Wickens and Wickens (1942) is a neat demonstration of the role of confusion or *generalization* in these effects. They trained two groups of rats in a box with two chambers to run to the second chamber when shocked in the first. For one group the shock came on suddenly; for the other, its onset was gradual. Half the animals in each group were then tested with a light that came on suddenly; the other half with a light that came on slowly. As you can see from Table 4.2, the light nearly always elicited running when its speed of onset matched the training condition, but rarely did so when it did not match.⁵

Sometimes animals habituate to a stimulus; sometimes they are sensitized by it, so that their reaction increases with successive stimulus presentations. We seem to have things both ways here: Since a reaction can only decrease or increase, we can explain everything. Fortunately, the effect to be expected from a given stimulus does seem to depend on measurable stimulus properties, most notably stimulus intensity. Animals seem to habituate to stimuli of low

		Shock Onset (training)	
		Gradual	Sudden
Light Onset (testing)	Sudden	2/7	7/9
	Gradual	8/10	1/8

Table 4.2. Summary results from a pseudoconditioning experiment by Wickens and Wickens (1942). Cells show numbers of animals per group showing the four possible results.

or moderate intensity; but become sensitized by intense stimuli. For example, a sudden, loud tone will elicit a startle response from rats. But after repeated presentations, the response habituates. However, if the same loud tone is presented against a background of loud white noise, the reaction not only fails to habituate, it increases across tone presentations. The critical factor seems to be the damage-potential of the situation.

Classical conditioning. If a hedonic stimulus (US) is reliably preceded (signaled) by a neutral stimulus (CS), many animals can learn to use the CS as a signal for the US. The process differs from sensitization and pseudoconditioning in that the CS must really be a good predictor of the US; it is not sufficient that they occur more or less close together in time. In the

standard procedure studied so extensively with dogs by Pavlov, a US such as food is repeatedly preceded by a neutral CS such as a tone. After a few such pairings, the salivation produced originally only by food is now produced by the tone as well. The reaction to the US is called the *unconditioned response* (UR); the reaction to the CS is called the *conditioned response* (CR). A comparable experiment with *Stentor* would involve pairing brief carmine presentations (the US) with some other stimulus, such as a change in illumination or temperature (the potential CS). The carmine elicits turning away; the light-change initially elicits nothing. If after a few pairings the light produces turning, then we *may* have classical conditioning. Additional control experiments in which the order of CS and US is varied (to rule out sensitization and pseudoconditioning) are necessary to be sure — we want to be sure that it is the predictive relation between CS and US that is important. And additional tests with the CS alone, after a delay, are necessary to see if the change is a relatively permanent one — as it usually is in higher animals. In practice, of course, protozoans rarely pass all these tests; most of their learning is habituation, sensitization, or pseudoconditioning.

Classical (Pavlovian) conditioning is the prototype for all signal learning. The CS is a signal for the US, and the animal reacts to the CS as if it were the US — although careful scrutiny usually shows that the reaction to the CS is anticipatory, rather than just a copy of the reaction to the US.

Operant conditioning. Suppose we pair a tone with food a few times, in Pavlovian fashion, but then present the tone alone: What will the animal do? Pavlov knew perfectly well that his dogs would not sit quiet under such conditions (that's one reason he restrained them in a harness). Given the opportunity, the dog does things that might produce food: If the experimenter is in the room, the dog will beg from him. If not, the dog will paw at the food bowl and try to get out of the room — whatever the environment permits by way of exploration. Suppose one of these explorations is, in fact, effective, as it might be if the experimenter had merely hidden the full food bowl, for example. The dog then eats and is removed from the room. If the experiment is repeated on the following day, the tone, the room, and the various other stimuli will not now produce the random searching we saw the day before. Instead, the dog is likely to go rather directly to the place where he previously found the food.

This is a two-phase process: (a) The first phase is unsystematic behavior that eventually leads to something good (or avoids something bad). (b) The second phase is the recurrence of efficient behavior when the animal is later returned to the same situation. The two phases together are called *operant conditioning* (an essentially equivalent term is *instrumental learning*). The first phase fits the control-by-consequences definition of operant behavior. The second

phase is referred to as the control of operant behavior by a *discriminative stimulus* — more on this in a moment.

Operant and classical conditioning are closely related. The parallel between the two is easy to see in avoidance or escape conditioning. A situation that signals shock or some other aversive stimulus elicits a range of avoidance and escape reactions of three types: fight (attack any plausible target, such as a conspecific), flight (escape from the situation), or “freeze” (immobility).

These reactions tend to occur in a systematic sequence (as in the *Stentor* example), depending upon the severity and duration of the threat and the opportunities offered by the environment. If no reaction is successful in escaping from the threat, then the animal settles on the final reaction in the series, which, in the case of shock, is generally immobility. This would usually be described as classical conditioning: An aversive CS produces conditioned immobility or “helplessness.”

But if one of the reactions *is* effective in eliminating the threat, then a recurrence of the situation is likely after a few repetitions to lead to reliable repetition of this effective avoidance response. This is operant conditioning, and the response is an *operant response*.

The relation between classical and operant conditioning is clear: Classical conditioning is the process that permits an animal to detect that the CS predicts the US. In this way the animal is able to identify the hedonic qualities of the situation: Is it good or bad, and what kind of good or bad is it? Given this information the animal, be it rat or *Stentor*, has available a repertoire of potentially useful reactions that nature (in the case of *Stentor*) or nature and past experience (in the case of the rat) has given it to try. If some reaction is effective, it is selected (in ways I discuss later) and recurs when the situation recurs: This control of operant behavior by the situation in which it has developed is termed *discriminative control*. The environmental features that are effective in controlling behavior are collectively termed the *discriminative stimulus* for the response. Classical conditioning obviously sets the stage for discriminative control.⁶

Generalization and discrimination. Adaptive behavior demands as a minimum that animals respond differently in different situations. The different kinds of learning just discussed define “situation” in different ways. In habituation, the period just following a stimulus is treated by the animal as a different situation from other times. Sensitization and pseudoconditioning work in the same way, the time just after a US is different from other times. In operant and classical conditioning, the animal further differentiates the world according to stimuli in addition to the US. The CS and its associated context (discriminative stimulus) define not only a particular hedonic quality, but also, in the operant case, a pattern, perhaps a unique pattern, of behavior.

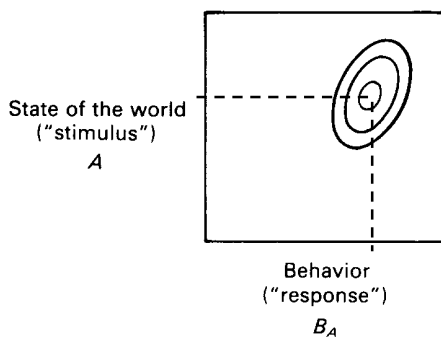


Figure 4.5. Definition of an *operant* as a region in a stimulus-response space. Closed contours are lines of equal probability.

Just what the term *situation* signifies in the operant and classical conditioning of birds and mammals is the topic of later chapters. For now, just consider the idealized view illustrated in Figure 4.5. The figure shows the various states of the world (“stimuli”) that can be discriminated by an animal, arranged along the vertical axis; the various modes of behavior of which the beast is capable are ranged similarly along the horizontal axis. Suppose that in the situation defined “A,” behavior B_A develops as a consequence of operant conditioning. Physical situations never recur identically — one day is sunny, the next is overcast or rainy, and so on. Consequently, animals are designed to accept minor variations stimulus situation. Animals never repeat a response identically either — nor should they, since some variants are more effective than others, and variability itself is sometimes advantageous in

operant and classical conditioning of birds and mammals is the topic of later chapters. For now, just consider the idealized view illustrated in Figure 4.5. The figure shows the various states of the world (“stimuli”) that can be discriminated by an animal, arranged along the vertical axis; the various modes of behavior of which the beast is capable are ranged similarly along the horizontal axis. Suppose that in the situation defined “A,” behavior B_A develops as a consequence of operant conditioning. Physical situations never recur identically — one day is sunny, the next is overcast or rainy, and so on. Consequently, animals are designed to accept minor variations stimulus situation. Animals never repeat a response identically either — nor should they, since some variants are more effective than others, and variability itself is sometimes advantageous in

avoiding traps (as I noted in Chapter 3). These two sources of variation mean that the pattern of behavior that develops in situation A is better represented as a region or *probability distribution* in the stimulus-response space in Figure 4.5 — rather than as a point where a unique stimulus produces a perfectly stereotyped response. One can think of this relation between situation and action, termed an *operant* (*habit* is the older term), as a location in the space centered on the point (A, B_A), where the third dimension, out of the page, corresponds to response strength or probability.

This picture is an enormous simplification, but it enables me to define generalization and discrimination in a simple way. *Generalization* refers to the effect of systematic variation in the physical stimulus situation on the operant response controlled by that situation. Usually the effect is to reduce the probability of the measured response, as suggested by the hillock pictured in Figure 4.5. Thus a pigeon trained to peck a green disk for food reward will peck less frequently on a yellow or red disk. Generalization leads to two kinds of “confusion”: A given operant response can occur in situations other than (but similar to) the training situation; and responses other than the given operant can show up in the training situation.

The extent of generalization obviously depends upon the animal’s individual and species experience of variability. Aspects of the environment that never vary can be ignored: Most lizards can move up or down a temperature gradient to find the temperature that suits them best. But some tropical rain-forest species have lost this ability: because of the constant temperature of their environment they have lost their sensitivity to temperature variation. This is an extreme deficit in discrimination and reflects the history of the species. An animal reared in an experimental environment where large changes have little significance is likely to generalize more broadly than an animal reared in an environment where large significance routinely attaches to small changes. This is true generalization, an effect of individual history. Measured generalization reflects both individual and species history.

Discrimination refers to the limits on an animal’s ability to treat similar situations differently. Generalization and discrimination are loosely, and asymmetrically, related. Generalization refers to the stimulus variations the animal willing to accept and still define the situation as essentially unchanged. Discrimination refers to the minimal variations in situation (stimulus) that the animal *capable* of treating as different. A pigeon may be able to detect the difference between wavelengths of 500 and 600 nm (they look green and yellow to us), but be willing to treat them alike because in the past they have been associated with the same consequences. On the other hand, wavelengths of 500 and 501 nm can be associated with consequences as different as we like, yet the pigeon will not treat them differently because it cannot tell them apart.

An animal may be able to detect small differences, but be inclined to accept large ones as of no significance. Hence, broad generalization need not imply poor discrimination. But an animal cannot treat similar situations as different if it cannot tell them apart. Hence poor discrimination implies broad generalization

Learning in the broad sense requires a change in the animal’s internal state caused by stimulation. Any change of state is a sort of memory. Habituation, sensitization, and pseudoconditioning require only local memory: a change in state initiated by US presentation that decays with post-US time. Learning in the narrow sense (often termed *associative learning*) requires a change of state of a more subtle kind, one that can be reinstated by only part of the original stimulus complex — the CS alone, without the US. This change is termed *long-term memory*.

Students reading with yellow marker in hand will notice that I have used the terms *memory*, *stimulus*, *response*, and *internal state* without defining them precisely. The last section of the chapter plugs this gap.

THE LOGIC OF HISTORICAL SYSTEMS

The eccentric and sometimes profound philosopher Ludwig Wittgenstein was discussing the topic of *time* with a colleague. As philosophers often will, the colleague soon asked Wittgenstein to define what he meant by time. Wittgenstein replied along the following lines: “If you know what I mean by *time*, let us continue our discussion. If not, let us discuss something else.” I have followed a similar strategy. The terms *stimulus*, *response*, *memory*, and so on are understood by everyone at a level sufficient to use them in discussing simple learning phenomena. Nevertheless, unlike “time,” the terms can be defined precisely. Definition is helpful because it shows that the role of discriminative and conditioned stimuli must go well beyond their obvious effect in eliciting a response. The analysis also clarifies limitations on our ability to understand *historical systems*, that is, systems whose future behavior is not predictable from what we can discern of their present state, but depends on past events: A computer is a historical system because we can’t tell from looking at it what it will do when we type something at the console; but a mechanical clock is not, because we know everything of interest about it from the position of its hands.

The first thing to remember is that when we look at the behavior of an animal and try to understand it, we necessarily deal in *models*. We cannot hope to comprehend the real complexities of any animal, nor would such detailed knowledge be very useful. We are always talking about a simplified “ideal animal” that behaves like the real animal only in respect of the things we happen to be interested in. When we use terms like “stimulus,” “response,” and “internal state,” when we assume that the animal is the same when we repeat earlier treatments, it is understood that we are thinking of our model animal, not the real one. Nothing is ever the same from day to day, least of all a living organism. Every stimulus is physically different from every other, as is every response. Any real animal has an infinity of physically different internal states. Nevertheless, the first step in understanding behavior is to put similar things into classes and look for regularities. The infinity of things that are different about the animal from occasion to occasion is of less interest than the much smaller number of things that are the same, that repeat themselves in a lawful manner.

Let us begin with the concept of *stimulus*. There are two problems: First, how are we to classify stimuli that have no overt effect on the animal? Second, are we to define stimuli physically or functionally?

Imagine a passenger in a bus on a busy city street. The bus halts briefly at a stop and at the same time a blue car pulls up on the other side of the street. Within a minute the bus pulls away and the car is lost to view. The blue car has elicited no reaction, perhaps not even a glance, from the bus passenger. Yet the next day, when questioned by police interested in a robbery, he may be able to recall details about the car such as its color and make. Clearly, the car was a stimulus in some sense. Hence, stimuli need not elicit immediate reactions.

This example suggests an answer to the second question also. The passenger describes his experience not in terms of the physical properties of the car, but in terms of the knowledge categories he shares with the interrogator: The car was a blue Ford sedan. Someone who had never seen a car and did not, therefore, possess the appropriate categories, would have had to resort to a much more cumbersome and detailed, and probably less accurate, description. Hence, stimuli are most usefully defined in terms of what the subject knows, that is, functionally rather than physically.

*Finite-state systems*⁷

These conclusions can be formalized as three definitions:

1. A *stimulus* (input) is something that either elicits a response, or changes the internal state of the animal.

2. A *response* (output) is something that is jointly determined by a stimulus and an internal state.

3. An *internal state* determines which stimulus will elicit which response and how a stimulus changes the internal state.

Stimulus and *response* correspond approximately to their commonsense meanings. I argue in a moment that *internal state* corresponds to what I have been calling a *situation*. These three definitions formalize something long accepted as a working rule by most psychologists, namely that stimulus, response, and internal state are all defined in terms of one another.

If we simplify things, the relations among these three terms can be illustrated by means of two tables. Suppose that changes inside our model animal are controlled by a clock, and occur

	<i>S - R</i> State		<i>S - O</i> State	
	X	S	X	S
<i>Stimulus</i>	s? n	y	c S	S
x	x	x	X	X
<i>Stimuli</i>	<i>Responses</i>		<i>States</i>	
s?: Did you see a car	y: Yes, I saw it		S: having seen the car	
c: car	n: no		X: unknown	
x: unknown	x: unknown			

only with each clock “tick,” rather than continuously as they probably do in real animals. (This is termed a *discrete-time* model.) In between clock ticks, the model animal stays the same. This simplification does not limit the generality of my conclusions, because any continuous system can be mimicked as closely as we pleased by a discrete system with a large number of states.

According to our definitions, there are just two ways that an internal state can affect behavior: (a) It can affect the relation between stimuli and current responses. For example, in state 1, stimulus S_1 might elicit response R_1 ,

Table 4.3. Finite-state tables for a two-state model. Left: Stimulus → Response relations. Right: Stimulus → State-change relations.

whereas in state 2, S_1 might produce R_3 . (b) A state can affect the relation between stimuli and future states. For example, in state 1, stimulus 1 might cause a transition to state 2; but in state 2, stimulus 1 might cause a transition to state 3. In other words, if we are going to describe a historical system in terms of stimuli, responses, and internal states, we can describe everything we want to know about it by means of two tables. One table shows the relations between stimuli and current responses, with state as the parameter. Given the state, this *S-R* table shows which stimulus elicits which response. A second table shows the relations between stimuli occurring at one time (tick of the clock) n and the subsequent state at time $n + 1$, with state at time n as the parameter: given the state at time n , table *S-O* (*O* for organism) shows which stimulus causes which state *transition*.

These two tables can be used to illustrate the logic of the earlier example. To make the tables easier to read I adopt the convention that STATES of the model are represented by capital letters, stimuli by lowercase letters, and *responses* by lowercase *italics*. The model obviously does not define all the behavior of the organism: unknown states, stimuli and responses are denoted by x 's.

Table 4.3 illustrates the car story: The *S-R* table is on the left and the *S-O* table on the right. The *S-R* table shows the response, n , y , or x , elicited by the stimulus (the question) or by some other stimulus, x , when the individual is in state X (“other”) or S . The *S-O* table shows the state produced by stimuli c (the car) or x (“other”). The tables have a lot of x 's in them, because all we know about the situation is that if the individual saw the car (i.e., his state was changed by it), then he responds “yes” to the question, but otherwise he does not. The critical entries are in the *S-R* table where state S yields a positive response to the question ($s? \rightarrow y$), and in the *S-O* table, where the sight of the car changes any state (X) to S ($c \rightarrow S$). The other entries are essen-

tially “don’t know” — we don’t know the effects of stimuli other than the ones whose effects we are describing.

The essential information in Table 4.3 is shown in another way in the upper part of Figure 4.6. Here states are represented as nodes in a directed graph. Each arrow in the graph begins with the stimulus received and ends with the response made, and the arrow connects the initial and final states. The lower part of Figure 4.6 shows a slightly more realistic representation: Instead of just two states, “having seen the car” (S) and “not having seen the car” (X), it shows three, the third one being “recalling the car” (S^*). The idea here is that the question “Did you see it?” changes the individual’s state into one where he is actually recalling the car. The evidence for this additional state is that the first question about the car is usually answered more slowly than later ones, suggesting that the first question somehow reinstates or makes more readily available other information about the original experience. Obviously still more complex representations could be devised to accommodate additional experimental facts.⁸




Figure 4.6. Two finite-state models for recall.

State S^* corresponds to the *recall situation*; it is the “state of mind” created by the interrogator’s question. Discriminative and conditioned stimuli seem to have a similar effect, in the sense that they reinstate a certain set of behavior potentialities. The term *expectancy* is sometimes used but, like the term *anticipation* I used earlier, it may say more than necessary. I return to this question in later chapters.

Equivalent histories. The common feature of all these representations is that “seeing the car” produces a persistent change that can be detected by the different responses made to a later question. All the models are oversimplified in obvious way. The schemes in Figure 4.6 make no provision for our passenger remember anything other than the car. Clearly state S must be further differentiated into states S_1 , S_2 , and so forth, to take account of the effects of later experiences. In the same way, state X must also be differentiated into numerous states to accommodate the varied possibilities for prior experience. Nevertheless if we ask our passenger only the one question, everything we know about him can be represented by the upper diagram in Figure 4.6. The states of our model must be inferred from the questions we ask of the subject, and they are limited by those questions.

A series of questions, or stimuli, is a *history*. The *states* in our models (any models) are *equivalent histories*. This idea can be explained in the following way. Imagine that you have an infinite number of *replicas* of the historical system (animal, human, machine) that you wish to study. Imagine further that you know in advance what things are stimuli (inputs) for the system and what things constitute its responses (outputs). Then take replica number one, and subject it to a sequence of inputs of infinite length. Then take replica number two and do the same thing, but with a different, infinite, sequence of inputs; and so on for all the replicas and all possible sequences (this is a long experiment!). Now, sort all the stimulus-response sequences you have obtained. Suppose you find, for example, that if the first three inputs are the sequences abc, acb, abb, or acc then the response to the fourth input, whatever it may be, is the same; and similarly for the fifth, sixth, and so on. Apparently these four histories are equivalent in the sense that the system’s behavior afterward is the same for all. It is natural to summarize this list of experimental results by saying that these histories all put the system into the same *state*. Thus from the ex-

perimeter's point of view, postulation of a state is just a way of summarizing the effect of various past histories that are equivalent in terms of the system's future behavior².

The experimenter's hope is that by classifying the various histories he has observed, he can reduce the infinity of input-output relations to a finite number of states. Since a real experimenter has less than infinite time and few or no real replica organisms, he must proceed by shortcuts, making guesses about possible states based on a preliminary classification, and then choosing "test histories" that will check out his guesses. The classifying process is called *induction* or the *inductive method*; the guess-and-test is called the *hypothetico-deductive* method. All science involves both, although the proportions are different for different sciences: Paleontology and sociology, for example, are largely inductive; physics is largely hypothetico-deductive.

It should be clear that understanding a system with very many states is a formidable task. Moreover, the logic of historical systems imposes impossible conditions on the conscientious experimenter — an infinity of replicas, and infinite time with each. These conditions are never satisfied. In the matter of replicas, there are only two options: either take different animals and assume they are essentially the same, or take the same animal at different times and assume that the intervening history has had negligible effect. The first option corresponds to the *between-group* method of experiment. Different histories, the experimental treatment and one or more control treatments, are given to each group and the average results are compared. We can never be certain here whether the animals are really equivalent to each other, or whether it is legitimate to average the results.⁹ The second option is the *within-animal* method. Treatments to be compared are presented successively to the same beast. To check that no systematic changes are taking place, "control" treatments are often repeated in between different "experimental" treatments; but we can't be certain here whether similar behavior on successive occasions means that the animal returns to essentially the same state, or not.

Despite these logical difficulties, both within- and between-animal methods are widely and successfully used, albeit for rather different kinds of problems. We have no real choice. When the causes of present behavior can lie in the remote past, the problem of experimental analysis becomes essentially insoluble, in the sense that there is no practicable method that will *guarantee* a solution. Hence, the experimenter has no choice but to set up situations where he hopes he can rule out the effects of remote past history. He tries to ensure that the animal is in the same state when he makes the manipulations whose effects he is comparing, which means either assuming that different animals are equivalent, or that certain kinds of intervening experience have negligible effect, so that "behavior" is *reversible*¹⁰ — i.e., that the animal can be restored to the same state.

Memory. Memory in the broad sense is simply an effect of a past event on future behavior. Any system with more than one state has memory of a sort. As we have seen, there are two main types of memory: habituation, sensitization, and pseudoconditioning require only local memory; classical and operant conditioning require long-term memory. The adjective "long-term" is unfortunate in its emphasis on the persistence of memory. The distinctive thing about long-term memory is not so much its persistence as its *context sensitivity*, that is, the power of a stimulus situation to recreate a given behavior potential.

Long-term memory can be further subdivided into two types, synchronous and temporal. *Synchronous* refers to control of operant behavior by the current physical stimulus situation: the rat pressing a lever in a particular Skinner box, the porpoise doing its tricks on the trainer's signal. *Temporal* refers to control of operant behavior by an event in the past. For example, if fed at fixed time periods, most animals will learn to anticipate food delivery and will nearly always

² Radical behaviorists object to the idea of state, but it is hard to see why. The generic definition of stimuli and responses (see, for example, Skinner, 1935; Staddon, 1967) defines stimuli and responses as classes (see Endnote 1). The idea of state simply does the same thing for histories.

be close to the feeder as feeding time approaches. Each food delivery is a temporal signal for the next and controls future food anticipation. Animals can also learn to choose one stimulus rather than another on the basis of the most-recent, or least-recent, stimulus they have seen. Temporal control differs from local memory in two main ways: (a) It is context-sensitive — an animal can be controlled by one past event in one context, by another in another. Apart from dishabituation, local-memory effects are largely independent of context. (b) It is modifiable by training: The animal can learn to do different things under the control of different stimuli after different times, in more or less arbitrary combinations.

Temporal and synchronous control are not mutually exclusive; all temporal control is context-dependent, and the effects of current stimuli are almost always modified by events in the animal's immediate past. For example, an animal once fed for pressing a lever and now not will eventually cease to press; but a single food delivery in the familiar context will at once reinstate the old pattern.

Endogenous rhythms. These rhythms, such as *circadian* (approximately 24-hour) rhythms shown by almost every organism, are intermediate between synchronous and temporal control. Rhythms have no single temporal stimulus; external stimuli are involved in the form of *zeitgeber* or “Timekeepers” such as the light-dark cycle that advance or retard the natural rhythm so as to synchronize it with an environmental cycle. A number of attempts to find a link between circadian rhythms, which are more or less fixed and not dependent on learning, and interval timing of the sort shown by fixed-interval schedules of reinforcement, have failed. They seem to be different and to depend on different neural systems.

SUMMARY

The most important difference between kinesic behavior and the operant behavior of mammals and birds is in its sensitivity to context. The internal state of the kinesic animal depends only on its current environment and the environment in the immediate past (local memory). The bacterium sees the world through a brief, moving time-window, with no memory of the past and no prevision of the future. Mammals, birds, and other intelligent animals can behave in a variety of ways, however; and the ways that they behave depend on both their current environment and their past history (long-term memory).

This chapter has dealt in a rather philosophical way with three related topics: (a) the relation between mechanistic and functional explanations of behavior; that is, between explanations in terms of antecedent events (stimuli past and present) and in terms of consequences (goals, rewards and punishments); (b) the differences between the operant behavior of simple and higher animals; and (c) the logic of explaining behavior in terms of antecedent events.

There are three main conclusions: (a) Mechanistic and functional accounts answer different questions: The mechanistic account is specific and leads naturally to physiological questions. If we are curious about how *this animal* works, this is the kind of explanation we want. The functional account is more general and tells us something about the animal in relation to its niche — not as an individual, but as a member of a species. (b) The difference between the operant behavior of simple and complex animals seems to be in the richness with which the animal's environment is represented internally. The world of the bacterium has few categories; the world of the rat or chicken has many. This difference allows the “higher” animal to detect predictive relations between neutral and valued events — and makes it worth remembering them. (c) Stimulus, response, and internal state are all defined in terms of one another. From the observer's point of view, an internal state is a set of equivalent histories. The term *context* is functionally equivalent to *internal state*. The special feature of operant and classical conditioning (as opposed to other forms of learning) is the ability of different contexts to reinstate previously acquired patterns of behavior.

The operant behavior described in this and the preceding two chapters can be arranged in three levels, defined by the type of behavioral variation: (a) innate unsystematic variation, as in the kinesic behavior of bacteria; (b) innate systematic variation, as in the different reactions of the snail *Littorina* depending upon whether it was in or out of water (this is context-sensitive behavior, but of a built-in sort); and (c) learned systematic variation, as in operant and classical conditioning, in which varied reactions (which may be built in, or may themselves have been built up by past selection) are selected and come to occur in varied contexts. The remainder of the book is concerned with learned systematic variation. I first discuss the *steady state*, that is, the ways that animals adapt to particular fixed learning situations. The last four chapters look at the processes by which these steady states are arrived at.

NOTES

1. *Skinner's definition of operant behavior.* B. F. Skinner in his various writings has used the term *operant behavior* in three independent senses: (a) By exclusion, as behavior that is not tied to an eliciting stimulus. It is, if not spontaneous, at least behavior for which "... no correlated stimulus can be detected upon occasions when it is observed to occur" (Skinner, 1938, p. 21). (b) As behavior involving units, operants, defined by the correlation between stimulus and response. Skinner argued that both stimulus and response are classes of events, and the definition of each class is an empirical matter: Each should be defined in terms of the other so that the elements show "orderly changes" as a function of relevant variables (see Staddon, 1967, and Schick, 1971, for discussions of the philosophical basis for this position). A stimulus becomes "that which causes a response," and a response "that which is caused by a stimulus." This kind of bootstrap definition may seem peculiar but is, in fact, quite characteristic of scientific theory generally (see Staddon, 1973). I provide a formal definition along these lines for stimulus and response later in the chapter. (c) As behavior that is "controlled" by its consequences. The third definition is the one that has achieved widest currency and carries the fewest theoretical overtones. This is the sense in which I use the term.

There is some contradiction between Skinner's first two characteristics of operant behavior: How can a behavior occur in the absence of any stimulus and yet be defined by the closeness of association of stimulus and response? Skinner was aware of this problem and has always insisted on using the term "occasioned by," rather than "elicited by" when referring to the stimuli that cause operant behavior. The paradox arises from his considering only *learned* behavior. In higher animals, behavior that is effective in securing a reward or escaping a punishment becomes associated with the circumstances under which the animal learned to make the effective response. Features of the environment become *discriminative stimuli*, having the power to reinstate this behavior or later occasions: The dog learns to avoid shock by pressing a lever when the buzzer comes on so that in due time the buzzer becomes a discriminative stimulus for lever pressing. These are the stimuli to which Skinner refers. There is nothing in the notion of behavior guided by consequences that requires either learning or (which is implied) the existence of discriminative stimuli, however: The behavior of bacteria aggregating in a nutrient medium, or the escape behavior of *Stentor*, are both examples of operant behavior in the third sense, although neither involves either learning (in this sense) or discriminative stimuli. There is no compelling reason to reserve the term operant behavior just to the more complex case, and I use it here in the widest sense.

The first and third of Skinner's defining properties, spontaneity and control by consequences, are not contradictory but complementary. If behavior is to be determined by its consequences, it must first occur in sufficient variety that effective variants are available to be selected. Unfortunately, Skinner took this process of generation for granted and concentrated almost entirely on the fact of selection.

The phrase “control by consequences” is, of course, an oblique reference to feedback control, as I note in the text.

2. *Parallel models.* Chain reflexes are historically the oldest kind of mechanism proposed for adaptive behavior. The most implausible thing about chain reflexes is that they are *serial* processes: Each activity depends on the preceding one. Serial processes are vulnerable to disruption and imply a degree of stereotypy rarely found in behavior.

Obviously many other mechanisms could be invented for *Stentor* avoidance. Parallel models explain behavior as the outcome of competitive interactions among a set of possible activities, only one of which can occur at a time. Parallel models adapt readily if one activity is blocked and derive behavioral variability from variations in thresholds, a well-known phenomenon. The behavior of Jennings’ *Stentor* is most plausibly explained in this way.

The reasoning is as follows. (a) Since the animal engages in one or other of the five possible activities shown in Figure 4.2, never in more than one at a time, we assume that they *compete* for access to the “behavioral final common path.” (b) Since the activities occur in a fairly rigid order, we assume that their *thresholds* differ, “turning away” occurring at low values of an irritant stimulus, “breaking away” at high concentrations. (c) Under constant stimulation, all activities habituate after a while. Hence, “turning away” will tend to weaken after some time. As it weakens, its inhibitory effect on the other four possible activities must also diminish. The activity with the next-highest threshold is presumably “ciliary reversal,” which, accordingly, supplants “turning away” after a while. Soon, it also habituates, to be supplanted by “contraction,” and so on. In this way the entire avoidance sequence can be generated.

The key features of this kind of model are (a) that the tendency for any activity to occur depends both on the current stimulus and the state of other activities; and (b) that activities tend to habituate (which can be thought of as self-inhibition), thus effectively reducing the threshold for those activities they inhibit.

Parallel models of this sort are simple in principle. In practice, they give rise to behavior that depends in complex ways on the quantitative properties of habituation and reciprocal inhibition. Nevertheless their simplicity, and consistency with what we know of the “wiring diagrams” of simple animals, suggests that they accurately represent an essential feature of much adaptive behavior (see Grossberg, 1982; Kennedy, 1967; Ludlow, 1976, 1980).

3. Edward L. Thorndike is the inventor of the *law of effect*, the basis of the modern principle of reinforcement. His original experiments were done with cats and other animals escaping from cages that could be opened from the inside by pressing a lever, pulling a string, or some other simple but relatively “unnatural” activity on the part of the animal. I return to Thorndike and the law of effect in Chapter 5. For a review of the work that has been done on learning and habituation in protozoa see Corning, Dyal, and Willows (1973).

4. *Information theory.* I write here and later as if the animals have a set of watertight compartments into which physical stimuli must fit. In reality, of course, things are not so clear-cut. In the case of an animal with two categories, for example, it is not that every stimulus that it ever encounters fits into category A or category B, with no overlap. Rather, some stimuli are always in category A, some always in B, some classified sometimes as A and sometimes as B, and perhaps others are sometimes C and sometimes B, and so on. The critical limitation is in the degree of *information* we have about the stimulus, knowing into what category the animal has classified it (or vice versa). If the information is one *bit*, then we can treat the animal as if it had two categories, with every stimulus falling uniquely into one, even though, in fact, it may have N categories, with considerable uncertainty about how a given stimulus is classified.

The term *bit* comes from the theory of communication (Shannon & Weaver, 1949; see Cherry, 1961, for an exposition for psychologists) and is defined as the information necessary to decide between two equiprobable alternatives. If you are curious about the sex of my cousin and I tell you she is female, I have transmitted, and you have received, one bit of information. On the other hand, if you know my cousin's name is Gail and I tell you she is female, I may have transmitted less than one bit, since the name Gail is more often a female than a male name (i.e., you already know that the alternatives are not equiprobable).

The general formula is,

$$I = -\sum_1^N p_i \log_2 p_i,$$

where p_i is the probability of the i th event of N total. Thus, if you know I live in the 489 telephone-exchange area and I tell you my phone number, I have transmitted one set of four decimal digits out of a possible 10,000, that is, 13.29 bits of information.

Protozoa seem to be limited as to the amount of information they can process about their environment. This corresponds to a limitation on the number of internal states (in the sense defined later in the chapter) that they have.

The limited information-processing capacity of simple animals accounts for the importance of innate behavior to these animals. Even the most primitive animal must be able to identify a few things with high accuracy: a potential mate, common dangers, and food. Given limited computational capacity, it is more efficient to build in sensitivity to a few specific signals (usually chemical) for these predictable and necessary tasks, than to use limited general learning capability that may not be able to discriminate with the necessary precision.

5. The standard treatment of the properties of habituation is by Thompson and Spencer (1966). For a more recent, theoretical account and review, see Staddon & Higa (1996). Sensitization and pseudoconditioning are discussed at length in many of the older learning texts, such as Hilgard and Marquis (1940) and Kimble's update (1961).

6. The word *control* means quite different things in the two usages, *discriminative control* and *control by consequences*. Discriminative control is closer to the direct meaning — in many cases, discriminative control cannot be distinguished from simple elicitation of an operant response by its stimulus. I argue later that discriminative control is better thought of as a determination of the internal state of the animal by its present environment and past history, rather than as elicitation. Control by consequences means feedback control, as we saw earlier. It is unfortunate that the same word has these two senses, both in the rather restricted context of operant behavior.

7. The following discussion of finite-state systems follows the excellent account in M. Minsky's (1967) book: *Computation: Finite and Infinite Machines*. For an application of this approach to the causal analysis of operant behavior see Staddon (1973).

8. State X is the set of all states other than S ; equally obviously, S is also a set, the set of all states having the common property that the individual responds "yes" to the question. States in these representations are always sets of real states. This is just another way of saying that we necessarily always deal with models, simplifications of the enormously complex real organism.

Time-to-respond, *latency*, is a much-used dependent variable in numerous experiments on human cognition. These experiments show that questions about, or exposure to elements from previously experienced situations much facilitate responses to additional questions or elements. See, for example, Posner (1978) for a review of this work.

9. *Averaging of data.* The problem of the relation between averaged data and individual performance has been an enduring theme in methodological discussions over the years. For example, in the heyday of Hull and his students, numerous experiments studied the process of learning by looking at the improvement in performance across trials of groups of animals set to running mazes or solving discrimination problems. The resulting average curves are usually negatively accelerated, with each trial producing a substantial increment at first, and smaller and smaller ones as training progresses. The results can be fitted by a variety of negatively accelerated functions, of which the negative exponential is the most popular:

$$y = A(1 - \exp(-Bx)), \quad (\text{N4.1})$$

where y is a performance measure such as “percent correct,” x is the number of trials, and A and B are fitted constants.

How representative is Equation N4.1 of the performance of individual animals? This question usually cannot be answered in the obvious way, by looking at the learning curve of each animal, because these curves are typically very variable. If we take the variability at face value, then there is no point to averaging, since each animal is different. If we take the variability to be “noise” — random variation — then the legitimacy of averaging depends on what is assumed to be varying. What form of *individual* relation between y and x is consistent with the relation between averages given by Equation N4.1?

I’ll just look at two possibilities. One suggests that averaging is legitimate, the other, equally (some might say more) plausible, suggests that it is not.

If we take the “noise” assumption at face value, then each measured performance value, y , can be represented as the sum of the true value, plus a noise component, ϵ , which is a sample from a population with zero mean. We also assume that each individual is identical, so that the true relation between y and the independent variable x , $y = g(x)$, is the same for all

$$\begin{aligned} y_1 &= g(x) + \epsilon_1 \\ &\vdots \\ &\vdots \\ y_N &= g(x) + \epsilon_N. \end{aligned} \quad (\text{N4.2})$$

Averaging performance across animals therefore yields:

$$1/N \sum_{i=1}^N y_i = \bar{y} = g(x) + \bar{\epsilon} \quad (\text{N4.3})$$

By hypothesis, ϵ_i is drawn from a population with zero mean, so that $\bar{\epsilon}$ will tend to zero as N increases. To an approximation, therefore, the average result will be:

$$y = g(x),$$

which can be compared with the obtained relation,

$$y = F(x).$$

Hence the “noise” assumption allows us to equate $F(x)$, the average curve, with $g(x)$, the individual function.

This model has two rather implausible assumptions: that all animals are identical, and that all variability is “measurement error” of the same sort as the error in reading a meter in a physics experiment. All animals are almost certainly not identical; and there is nothing corresponding to measurement error in most psychological experiments. It is at least as reasonable to assume that even if each animal obeys the same general law, the parameters of the function will differ from individual to individual, and that this is the main source of individual variation. With this change, Equation 4.2 can be rewritten as:

$$\begin{aligned}
 y_1 &= g(x, a_{1,1}, \dots, a_{1,m}) + \varepsilon_1 \\
 &\vdots \\
 &\vdots \\
 y_N &= g(x, a_{N,1}, \dots, a_{N,m}) + \varepsilon_N,
 \end{aligned}
 \tag{N4.4}$$

where y depends on x according to a function with M parameters that may differ in value from individual to individual. By the same argument as before it follows that

$$F(x) = 1/N \sum_1^N g(x, a_{i,1}, \dots, a_{i,m}).$$

That is, the average function, $F(x)$, is equal to the average of each individual function of x . For what class of functions, $g(x)$, will the average function, $F(x)$, be of the same form as the functions averaged? The short answer is: not many. It is true if $g(x)$ is linear, as can be seen at once by averaging two

$$\begin{aligned}
 y_1 &= m_1x + c_1 \\
 y_2 &= m_2x + c_2,
 \end{aligned}$$

where m and c are constants. Hence,

$$y_1 + y_2 = x(m_1 + m_2) + c_1 + c_2,$$

or

$$y = x(m_1 + m_2)/2 + (c_1 + c_2)/2,$$

which is still a linear function.

It is not true for nonlinear relations such as Equation N4. 1. For example, suppose that $y_1 = m_1x^{a_1}$, and $y_2 = m_2x^{a_2}$, which are power functions, with m and a as constants. Then

$$y = m_1x^{a_1}/2 + m_2x^{a_2}/2,$$

which is not a power function.

Thus, rather little can be learned about the performance of individuals by looking at group averages. Despite this fact, most psychological research is done with group averages. For a more extensive discussion of this issue see Estes (1956), Stevens (1955), and Sidman (1960).

10. Unobservability. In addition to these difficulties, there are unavoidable uncertainties in the study of black-box systems. Because a stimulus can change the state of the system, some states may be literally unobservable by an experimenter. For example, even under ideal conditions, we can never be sure that a given event is *not* a stimulus. Failure to find any effect is always inconclusive. This limitation holds because the tests themselves can eliminate the evidence they are intended to elicit (see E. F. Moore, 1956). Stated formally, Moore's "uncertainty-principle" theorem is that "there exists a [finite-state] machine such that any pair of its states are distinguishable, but there is no simple experiment which can determine what state the machine was in at the beginning of the experiment" (1956, p. 138). It is easy to see the logic of this limitation in the case where there are three states to be distinguished call them A, B, and C. Let us suppose that A and B are distinguished by the fact that input 1 elicits output 1 from state A and output 2 from state B; B and C are distinguished by the fact that input 2 produces output 2 from B and 3 from C. Input 1 fails to distinguish B from C because it elicits 2 from both; similarly, input 2 fails to distinguish A from B because it elicits 2 from both. Finally, suppose that inputs 1 and 2 both send the system into state C.

To discover the initial state of the system we have only two choices, either input 1 or input 2. The first choice will distinguish between initial states A and B, but will fail to distinguish between states B and C. Moreover, additional inputs will not help, because the system will now be in state C no matter what its initial state; a similar problem results if the initial input is 2.

There is no theoretical solution to this problem, of course, other than having replicas known to be in the same state at the beginning of the experiment, so that different, noninterfering tests can be carried out on each replica. Unfortunately, real animals are replicas only in respect to restricted aspects of their behavior, and it is often far from clear what these aspects are. It is reasonable to assume equivalence in such matters as the properties of cells and organs, and possibly some psychological functions of sensory systems. It is much less clear that the processes involved in learning are essentially identical from individual to individual.

Everyday experience shows that Moore's theorem is relevant to psychology. For example, asking someone to recall an event may actually cause him to forget it, although some less direct test might have shown that he had, in some sense, remembered it. Many people have had the experience of being asked to recall verbally a familiar telephone number and being unable to do so; yet they could have dialed the number without difficulty. Sometimes the ability to dial correctly is retained even after failure of verbal recall, but more commonly failure to recall also abolishes (often only temporarily) the ability to dial correctly. Recent scandals involving "recovered" supposedly repressed memories show how easy it is to create false memories and how hard it is to distinguish them from real ones.

See, for example, <http://faculty.washington.edu/eloftus/Articles/lof93.htm>).