

MEMORY AND TEMPORAL CONTROL

Memory is probably the most protean term in psychology: it has as many technical as non-technical meanings, and the numbers of both are large. Psychologists have at various times written of long- and short-term memory, of working and reference memory, of episodic and semantic memory, of primary and secondary memory. In common speech the term memory refers to a purely private event: “I remember...”, something not directly accessible to observation. Little wonder that there is still no consensus on what we mean by memory, or on its relations to learning.

In Chapter 4, I defined memory simply as a change of state caused by a stimulus: memory is involved if how the animal behaves at time t_2 depends on whether event A or event B occurred at previous time t_1 . Breaking a leg is a change of state in this sense, of course, so we need to restrict the definition to effects that are *specific* and to some extent *reversible*: the difference in behavior at t_2 should bear some sensible, informational relation to the difference between prior events A and B; and we should be able to change the effect by additional experience. Nevertheless, the advantage of the formal definition is that it commits us to no particular theoretical position — and it draws attention to the memory-like properties of habituation, dishabituation, spontaneous recovery and, particularly, temporal control — phenomena not traditionally considered memorial.

Much is known about temporal control. The first part of this chapter reviews the properties of temporal control and derives some general principles about the discrimination of recency. In the middle part of the chapter, I show that these principles also apply to more traditional situations used to study memory in animals, such as successive discrimination reversal and delayed matching to sample. The last part of the chapter brings together the idea of internal representation, described in Chapter 10, and the principles of memory described in this chapter, to explain behavior in the radial maze and related spatial situations.

TEMPORAL CONTROL

As we saw in earlier chapters, animals readily detect periodicities: If a pigeon is rewarded with food for the first key-peck T seconds after eating (i.e., a fixed-interval T -s schedule), he will usually not begin to peck until perhaps two-thirds of the time has elapsed, that is, his *postreinforcement pause* will stabilize at close to $.67T$.¹ The animal is able to do this by using food delivery as a *time marker*. Control of behavior by a past event is termed *temporal* control, to distinguish it from control of behavior by a present stimulus, which might be termed *synchronous* control (cf. Chapter 4). Temporal control is the instrumental equivalent of trace conditioning.

An attractive way to describe temporal control by food on an FI schedule is to say that food “resets” the animal’s “internal clock.” Pecking is initiated when the clock reaches a value which is an approximately constant proportion of T . Many features of temporal control are consistent with the clock idea. For example, as in a real clock, the *error* in timing is proportional to the interval to be timed: a clock that is one minute fast after an hour will be six minutes fast after six hours. If the clock is simply variable from day to day, then the variation in its error over an actual time T will be proportional to T . This is an example of Weber’s Law, which is characteristic of many sensory dimensions, such as brightness and loudness. In the context of timing, it is often termed the *scalar timing* property.²

The *reset* property of the time marker can be demonstrated by omitting it or replacing it with something that is not treated as a time marker. For example, Staddon and Innis (1969) trained pigeons and rats on a fixed-interval 2-min schedule and then shifted to a procedure in which food was delivered at the end of only 50% of intervals. The change this produces is shown in Figure 13.1. Intervals ending with “no food” (a brief blackout of the same duration as food delivery) are indicated by “N” in the figure. The animals show the usual pause after food, but if food is *omitted* at the end of a fixed interval, then the animals continue to respond until the next food delivery: in the absence of a “reset” (food), responding continues. With continued experience of reinforcement omission, some animals learn to pause a bit after the nonfood stimulus, but for most, the effect persists indefinitely in the form shown in the figure. This absence of pausing after a nonfood stimulus is known as the *reinforcement-omission effect*.




Figure 13.1. Sample cumulative records of stable performance in an experiment in which animals were trained on a fixed-interval 2-min schedule. At the end of intervals marked “N,” food delivery was omitted and a brief 3-s stimulus was presented in its stead. Record at the top is from a pigeon, at the bottom from a rat. The recorder pen was reset at the end of every interval. (From Staddon & Innis, 1969, Figure 3.)

Synchronous discriminative stimuli seem to tell the animal what to expect at a certain time, rather than affecting his estimate of time directly. For example, Church³ trained rats on a procedure in which 30-s and 60-s fixed intervals were intermixed, each with its own distinctive (synchronous) discriminative stimulus (i.e., a multiple FI-FI schedule). The animals soon developed pauses appropriate to the signaled intervals. Then, in test sessions, the stimulus in the 30-s interval was changed abruptly to the 60-s stimulus. The change could occur 6, 12, 16, 24 or 30 s after food. The rats behaved as if they had a single postfood clock and used the stimulus just to scale their rate of responding to the clock setting. Thus, an animal’s rate of responding t s after food in the presence of the 60-s stimulus was the same, whether the stimulus was present from the beginning of the interval (i.e., simple FI 60) or only appeared at some later time. The rats seem always to know what time it is; the synchronous stimulus just tells them whether to expect food at that time or not.

The clock idea is a convenient simplification, but I show later that the reset property is far from absolute: under many conditions, earlier events, preceding the resetting stimulus, can affect behavior. These interference effects, and the conditions under which they occur, show that temporal control reflects the same process studied in more conventional memory experiments: temporal control and memory seem to be different aspects of the same thing.

Excitatory and Inhibitory Temporal Control

In the preceding chapters we saw that synchronous stimuli can either enhance or suppress an instrumental response. Temporal control can also be excitatory or inhibitory, depending on circumstances. All the examples discussed so far are inhibitory, since the instrumental response is suppressed immediately after the time marker. As with inhibitory synchronous control, suppression of the instrumental response is usually associated with facilitation of competing responses. These are the *interim* activities discussed in earlier chapters: The postfood “pause” is a pause only in instrumental responding, other activities, such as pacing in pigeons, drinking or

wheel-running in rats, occur at their highest rates during the pause.⁴

Animals can learn to respond rapidly after a time marker just as easily as they can learn to pause: excitatory temporal control is as easy as inhibitory control. For example, in one experiment (Staddon, 1970, 1972a) pigeons were trained on a VI 60-s schedule where the response required for food was either pecking, or *refraining* from pecking for at least 10 sec. The signal telling the animals which response was required was postfood time: at postfood times less than 60 sec, pecking was required to obtain food, at longer postfood times, not-pecking was required. Cumulative records on the left in Figure 13.2 show stable performance on a “temporal go→no-go” schedule of this sort.




Figure 13.2. Cumulative records of four pigeons trained either on a procedure that reinforces a high rate after food (left, go? no-go schedule) or a low rate after food (right, no-go ? schedule). Diagonal “blips” indicate food deliveries. See text for details. (From Staddon, 1972a, Figure 2.)

The animals show a sustained period of high-rate responding for about 60 s after each food delivery (indicated by the blips on the record) followed by no responding until the next food delivery. The records on the right show the “FI-like” behavior produced by the reverse: a “no-go→go” schedule in which not-pecking is required at postfood times less than 60 sec, pecking at longer times.

Like synchronous control, temporal control shows generalization decrement: variation in an inhibitory temporal stimulus produces increases in subsequent responding, variation in an excitatory stimulus, produces decreases. Because the effect of a temporal stimulus is delayed, discrimination is not as fine as in the synchronous case: larger changes in the stimulus must be made to produce comparable percentage changes in response rate (I describe other examples of these memory limitations in a moment).

It is easy to show effects of variation in the stimulus complex associated with food delivery on fixed-interval schedules. For example, in a particularly clear-cut study with pigeons (Kello, 1974), fixed intervals ended unpredictably with one of three events: food paired with 3-s extinction of the key and house lights and 3-s illumination of the feeder light (F), all these events without feeder operation (NF), or extinction of the lights alone (N). The pigeons paused longest after F, least after N, and an intermediate amount after NF, and response rates over the interval following each kind of event were in the reverse order. Comparable experiments with the excitatory procedure (Figure 13.2, left) have shown the opposite result: slowest responding after N, highest after F. In both cases, the effect of the test time marker is of the same sort as the training time marker, and directly related to the similarity between the two.

These experiments are examples of control by stimulus elements (cf. Chapter 10). It is trickier to demonstrate temporal control by a stimulus dimension, simply because it is harder to establish temporal control by “neutral” stimuli such as colored lights and line tilts (I explain why in a moment). Nevertheless, when good control is established, gradients of the standard sort are obtained. If the time marker is inhibitory, then responding following it increases as the test stimulus varies from $S+$; if the time marker is excitatory, then responding following it decreases as the test stimulus varies from $S+$.⁵

Conditions for Temporal Control

Under what conditions will a stimulus such as food come to serve as a time marker? The general answer is the same for both temporal and synchronous stimuli: When it predicts some-

thing. A temporal stimulus acquires inhibitory control when it reliably signals a period free of food delivery — such as the early part of each interval on a fixed-interval schedule. It will acquire excitatory control when it signals a period when the conditions of reinforcement are better than other periods — as in the go→no-go schedule just described.

The signaled period need not be immediately after the time marker, however. Consider a modified fixed-interval T -s schedule in which food is sometimes available (in perhaps 50% of intervals, on a random basis) at postfood time $T/6$, as well as always available after time T — this is a simple (2-interval) variable-interval schedule, where the intervals are T and $T/6$ (simple VIs are sometimes termed *mixed-interval* schedules). Such a schedule is illustrated in Figure 13.3. The top panel shows the probability of food delivery at different postfood times: It shows just two vertical lines, one, of height 0.5 at postfood time $T/6$, the other, of height 1.0 at time T . The bottom panel shows the average response rate as a function of postfood time for a pigeon that has adapted to this procedure. It shows a high rate just after food (roughly corresponding to the first probability “spike”), followed by a period of low response rate, ending finally with a high response rate towards the end of the T -s interval.⁶ Thus a time marker can initiate a sequence of alternating “respond-not respond” periods.

If the availability of food is random in time, then at any instant the probability that food

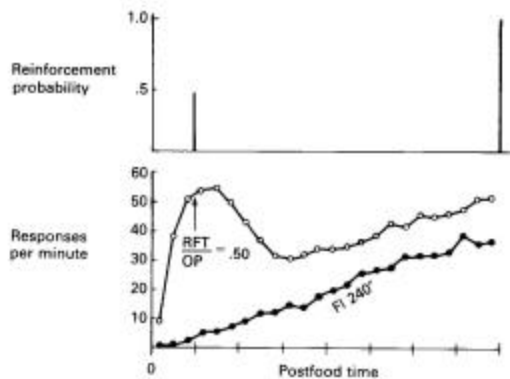


Figure 13.3. *Top panel:* food reinforcement probabilities as a function of postreinforcement time on a modified fixed-interval schedule. Reinforcement occurs with probability 0.5 30 s after food and with probability 1.0 240 s after food. *Bottom panel:* average rate of key pecking as a function of time since food for a pigeon trained on this schedule. (From Catania & Reynolds, 1968, Figure 19.)

will become available for a response is constant — food is no more likely at one postfood time than another. This is a *random-interval* schedule of the type already encountered in Chapters 5, 7 and 8. We would not expect, nor do we find, much patterning of responding as a function of postfood time: Average response rate is approximately constant. As we saw in Chapters 7 and 8, however, probability of food does increase as a function of post-response time on these schedules, since the longer the animal waits, the more likely that the random food-availability programmer has “set up.” We might expect, therefore, that probability of response will be low immediately after a response. This is true to some extent, although the effect is somewhat masked by the tendency of pigeons to respond in “bursts” of or two or three pecks at a time. Animals on concurrent variable-interval, variable-interval schedules show in their patterns of choice that they are quite sensitive to this property of vari-

able-interval schedules, as we saw in Chapter 8. *Spaced-responding* schedules make the temporal requirement explicit, only reinforcing responses longer than time T . If T is fairly short (less than a minute or so), pigeons and rats adapt by spacing their responses appropriately.

What determines exactly when an animal will begin to respond after a time marker? The functional answer suggested by the discussion of optimal behavior in Chapter 7 is that it will depend on what other activities are available to the animal, and the constraints to which he is subject. An important but hard-to-estimate constraint is set by the psychophysical limits to the animal’s ability to tell time. For example, suppose that in addition to pecking the key for food, the animal has at least one other thing he likes to do. On a fixed-interval schedule, the amount of

time available for the “other” activity is critically determined by the accuracy with which the animal can estimate the interfood interval. If he is very accurate, then he can defer key pecking until just before the end of the interval (i.e., to a late setting of its internal clock), with the assurance that he will only rarely respond late, and so receive food after a longer time than necessary. On the other hand, if the animal is very inaccurate, he cannot safely defer pecking until almost the whole interval has elapsed, because by so doing he will often wait much too long. This problem was discussed earlier in Chapter 5 (Figure 5.18). The animal’s choice of pause “setting” is necessarily a compromise: If he sets it long, so as to waste few responses and have a maximum amount of time for interim activities, he gets food less frequently than he might if he set it shorter — but then he would have less interim time and waste more terminal responses. Moreover, as the timer is set longer, the variance of the pause distribution increases, according to the Weber-law property, which worsens the terms of the tradeoff.

The costs and benefits associated with any particular setting for mean pause can be derived from the minimum-distance model discussed in Chapter 7. This could be done analytically, but a qualitative analysis is quite sufficient to show that typical behavior on fixed-interval schedules is just about what we would expect. We already know from the earlier discussion that the *cost* of deviations from the preferred rate of eating is typically high relative to the cost of deviations from the preferred rate of key pecking or lever pressing. Hence, we can expect that animals will set their pause distributions so as to keep the area to the right of the line in Figure 5.18 quite small (so that eating rate is as close as possible to the maximum allowed by the interval schedule). On the other hand, we might expect that if offered a relatively attractive interim activity, the pause setting might shift to longer values. This usually happens: rats on a fixed-interval schedule in the usual bare Skinner box will show shorter pauses than rats also offered the opportunity to run in a wheel. Pigeons trained in a small box, or restrained in a body-cuff, show shorter pauses than animals responding in large enclosures (Frank & Staddon, 1974).⁷

Characteristics of the time marker

Food is not the only stimulus that is effective as a time marker, but it is more effective than “neutral” stimuli such as tones and lights — or even the animal’s own response (recall that pigeons and rats can only learn to space their responses if the delay times are quite short). For example, consider again the procedure illustrated in Figure 13.1. Food was omitted at the end of half the fixed intervals in that experiment, but something happened even at the end of no-food intervals: The light on the response key went out (for the pigeon) and the “house” lights went out (for both rat and pigeon) for about three seconds — a period equal to the duration of access to food at the end of food intervals. This brief “timeout” period tells the animal exactly as much about the time until the next opportunity to eat as does food; in both cases, the next food opportunity is after two minutes. Yet both rat and pigeon paused after food but not after the timeout. Why?

There are obviously two possibilities: Either the original hypothesis — that the pause is determined by the predictive properties of the time marker — is wrong, or there is something special about food (and electric shock and other “hedonic” stimuli) that makes it more effective than a neutral stimulus. There is too much other support for the predictiveness idea to give it up, and there is much evidence that there is indeed something special about hedonic stimuli such as food.

Consider some other experimental results. In the first experiment (Staddon, 1970b; see also Lowe, Davey, & Harzem, 1974) pigeons were trained on a fixed-interval 1-min schedule in which each interval ended with food reinforcement, with the special proviso that the *duration* of access to food varied unpredictably from interval to interval. There were five different feeder

durations, ranging from 1.3 to 9.0 seconds. In this experiment the pigeons had food available as a time marker in every interval; there is no obvious reason why they should not have paused in much the same way as on the usual FI schedule in which all food deliveries are of the same duration. But they did not. The results, for three pigeons, are shown in Figure 13.4. The bottom panel shows the average pause after each of the five food durations; all the pigeons show a positive relation: the longer the food duration, the longer the pause. The upper panel shows similar data on rate of responding over the whole interfood interval. Response rate decreases as food duration increases.




Figure 13.4. *Top panel:* Response rate in intervals following the access-to-food durations shown on the abscissa for three pigeons trained on a fixed-interval 60-s schedule in which intervals terminated unpredictably with one of five different food durations. *Bottom panel:* time-to-first-response (postfood pause) following different food durations. (From Staddon, 1972b, Figure 1.)

There is an uninteresting explanation for this result: Perhaps the pigeons simply take longer to swallow their food after eating for nine seconds as compared to eating for one or two? This plausible explanation is wrong for at least two reasons. First, it implies that pigeons should pause after food on any schedule, not just fixed-interval, but as we know, they do not — they show minimal pauses on variable-interval schedules, for example. As we have seen with the go? no-go schedule it is also relatively easy to train animals to respond *especially fast* after food, rather than pausing. Second, and more directly relevant, other experiments⁸ have shown that the differential-pause effect depends upon the animals experiencing different food durations *within the same experimental session*, or at least in an intermixed fashion. If, instead of daily experiencing five different food durations, the animals are given several days at one duration, then several more days at another and so on, then pausing after the short durations increases from one session to the next (or decreases, if the food duration is long) so that soon all differences disappear and the animals pause about the same amount after any food duration. The differential-pause effect depends on *intercalation* of different food durations.

Evidently, the effectiveness of a stimulus as a time marker depends on its freedom from interference from other remembered events. The long feeder durations were evidently less susceptible to interference than the short, when long and short were intercalated, so that postfood pause was longest after the long. When only short intervals occurred, however, they did not interfere with each other, so pause lengthened.

The destructive effects of interference between intercalated stimuli can be shown directly. In the following experiment (Staddon, 1975a) pigeons' ability to use a brief stimulus as a time marker was impaired by intercalating it with another stimulus with no predictive significance. The birds were first trained to respond for food on a variable-interval 1-min schedule. After a little experience, the birds showed characteristic steady responding, with no postfood pausing. In the second phase, every two minutes a brief (3-s) stimulus (three vertical lines) was projected on the response key. This stimulus signaled that the *next* reinforcement would be programmed on a *fixed*-interval 2-min schedule. Thus, after food, or at any other time, the animal could expect food after some unpredictable time averaging one minute; but after the 3-s vertical-line stimulus, the animal knew that food would be available only after exactly two minutes.

The pigeons adapted to this new time marker by developing a post-stimulus pause appropriate to the two-minute fixed-interval duration. This pattern is illustrated for one animal by the cumulative record on the left in Figure 13.5. The record reset at the end of each 4-min cycle (i.e., after each FI 2-min food delivery), and the recording pen was depressed during the fixed interval. The pause after the brief stimulus is clear in four of the five intervals in the figure, a typical proportion. This result shows that when there are no interfering events, pigeons can learn to use a brief, neutral stimulus as a temporal cue.

The right-hand record in Figure 13.5 shows the effect of an apparently trivial modification of this procedure. Instead of scheduling the 2-min fixed interval exactly once during each 4-min cycle, it was scheduled on only half the cycles. By itself, it is not likely that this change would have had any significant effect, providing the brief vertical-line stimulus continued to provide a reliable temporal cue. However, during those cycles when no fixed interval was




Figure 13.5. Left record: cumulative records of stable performance on a variable-interval 1-min schedule in which a brief vertical-line stimulus (indicated by depression of the response record) occurred every 4 min; the stimulus signaled a 2-min fixed-interval schedule. The record reset at the end of each FI 2 reinforcement. Right record: performance on this schedule when some 4-min periods were initiated by a brief horizontal-line stimulus (arrows) signaling “no change” (i.e., the VI 1 schedule continued in effect). (From Staddon, 1975, Figure 3.)

scheduled (i.e., when the VI 1-min schedule remained in effect), a brief *horizontal-line* stimulus was projected on the response key. The first effect of this change was that the pigeons paused indiscriminately after both horizontal and vertical stimuli. This necessarily led to their sometimes waiting longer than necessary after the horizontal stimulus. Eventually, the animals ceased to pause after either stimulus, as shown in the right-hand record in Figure 13.5.

Pigeons have no difficulty in telling vertical from horizontal lines when they are required to do so in standard simultaneous or successive discrimination procedures. So their failure to pause differentially after the two stimuli in this experiment doesn't reflect some kind of perceptual limitation: They can

recognize a vertical-line stimulus when they see it, and learn to respond differently in the presence of vertical and horizontal lines. The problem seems to be that in this experiment the animals were not required to respond *in the presence* of the stimuli. Instead they had to behave differently *after* the stimuli had come and gone — pausing after the vertical lines, not after the horizontal lines. In the intercalated (horizontal and vertical) condition, the vertical lines produced only brief pauses: the effect of the stimulus was restricted to a relatively brief post-stimulus period. Why?

A commonsense explanation is just that in the intercalated condition, the animals *could not remember* for more than a few seconds which stimulus had just occurred. A more descriptive way to put the same thing is to say that the effect of the informative vertical-line stimulus was abolished by prior presentation of the (uninformative) horizontal-line stimulus, an effect termed *proactive interference* — impairment of recall of a stimulus by occurrence of a prior stimulus.

There is a complementary effect, well known in human memory studies — *retroactive interference*: impairment of control by an earlier stimulus by the interpolation of a later one. Is there a parallel effect in the study of temporal control? The evidence is less clear than for proactive interference, but there is a common phenomenon that may be related: If a novel stimulus is

presented during the pause on a fixed-interval schedule, animals will often begin responding at once. This is a sort of *disinhibition*, first seen by Pavlov in delayed-conditioning experiments. This effect may be owing to retroactive memory interference, the later event (the novel stimulus) impairing recall of the earlier one (food, the time marker)¹. The converse effect, insulation from proactive interference by means of differential stimuli during the fixed interval, has been demonstrated, as I show in a moment.

If the interference demonstrated in these experiments has something to do with memory, we can make some additional predictions. All theories of memory agree that under normal circumstances an event is better remembered after short times than after long. If, therefore, the effects of reinforcement omission reflect failure of memory they should be reduced when the time intervals involved are short. This seems to be the case. For example, if we repeat the fixed-interval reinforcement-omission experiment with 15 or 30-s intervals rather than 60 or 120-s intervals, the pigeons pause as long after brief stimuli presented in lieu of food as they do after food (Neuringer & Chung, 1967; Starr & Staddon, 1974).

Interference between to-be-remembered events can also be reduced if they are associated with different situations or contexts. For example, if a person is required to learn two lists of similar items, words or nonsense syllables, one after the other in the same room, and then asked to say whether a particular item comes from a particular list, he will often make mistakes, identifying an item from list B as coming from list A and vice versa. If the same experiment is done with the two lists either separated in time, or learned in different rooms or from different experimenters, confusions of this sort are much reduced.

A similar experiment has been done with temporal control.⁹ Two pigeons were trained on a fixed-interval 2-min schedule in which intervals ended either with food (F) or a brief blackout (N) with probabilities 1/3 and 2/3. The response key was either red or green during each fixed interval. Both stimuli gave the same information about the *outcome* of an interval: in either case, the probability the interval would end with food was 1/3. But when the stimulus during the interval was green, that interval had *begun* with N (i.e., blackout) — the green stimulus was a consistent context for remembering N (green *retrodicted* N). The red stimulus was an ambiguous context, because red intervals began indiscriminately with N or F.

The critical question, obviously, is whether the animals were better able to use the neutral time marker beginning green intervals than the same time marker when it began red intervals: would they pause longer following N in green than following N in red? The answer is “yes;” these two pigeons, and two others similarly trained with a shorter fixed interval, all paused almost as long after N as after F in green, but paused much less after N in red. Evidently, the distinctive context was able to mitigate the usual interference between N and F in fixed-interval reinforcement-omission procedures where intervals beginning with N and F are not otherwise distinguished.

Some “neutral” stimuli are more memorable than others. For example, the extraordinary human memory for faces has often been noted. The reason why people are able to identify hundreds, or even thousands of faces, but only a few (for example) telephone numbers is still being studied, but a popular suggestion is that it has something to do with the *multidimensional* property of “natural” stimuli like faces and scenes (see Chapter 10). There is some evidence that animals’ ability to use a stimulus as a time marker in fixed-interval schedules is similarly af-

¹ Some of the effects of the so-called *gap procedure*, in which a stimulus (usually the same as the stimulus signaling the inter-trial interval) is briefly presented during a trial on the *peak-interval procedure*, may represent a kind of retroactive interference, but stimulus generalization may also be involved. See Staddon & Cerutti (2003) and references therein for additional discussion of these procedures.

ected by stimulus complexity. For example, in an unpublished experiment in my own laboratory we have found that if instead of the usual simple color or blackout stimulus we present a color slide of a pigeon as the neutral omission stimulus, the birds show essentially normal fixed-interval pauses.

Events of longer duration are usually easier to remember than events of shorter duration. Suppose, in the standard reinforcement-omission paradigm, we present neutral events of variable duration at the end of half the fixed intervals — being sure to time the intervals from the end of the event in every case. Will animals pause longer after longer stimuli, as they did when the stimulus was food, and as the memory argument implies? The answer is “yes” for pigeons, but (except for a transient effect) “no” for rats (Staddon & Innis, 1969; see also Roberts & Grant, 1974).

Conclusion: the discrimination of recency

The reinforcement-omission effect — shorter pausing after a neutral stimulus presented in lieu of food than after food on fixed-interval schedules — seems to reflect a competition for control of the animal’s behavior between two past events: food, which is the earlier event, and the neutral stimulus. The animal must attend to the most recent event and ignore the earlier one. Both events have the same temporal significance, but food is more valued. Evidently, a few seconds after the neutral stimulus the animal attends to food rather than the neutral stimulus. Since the last food delivery is relatively remote in time, the animal responds (long postfood times signaling further food), resulting in a too-short pause after the neutral stimulus (the reinforcement-omission effect). The same process accounts for diminished pausing after short FI feeder durations when long and short are intercalated.

The general conclusion is that trace (temporal) stimulus control is vulnerable to the kinds of proactive and retroactive interference studied in memory experiments. Things that give a stimulus value, such as reinforcing properties, and stimulus complexity or “meaningfulness,” facilitate temporal control. Separation, in time, or by context, minimizes interference between events. Conversely, the occurrence of similar interfering events (the horizontal-vertical experiment) or more memorable events with similar significance (the reinforcement-omission effect) impairs temporal control. When the interfering event is similar in properties, but different in temporal significance, to the event of interest, the resulting impairment of temporal control may be termed a *recency confusion* effect, since the animal is evidently uncertain about which stimulus just occurred. When the interfering event is highly salient or valued, the resulting impairment is better termed a *recency overshadowing* effect, since the more salient, older, event exerts control at the expense of the more recent, less salient, event.¹⁰

Other methods for measuring temporal control and memory

Fixed-interval schedules might be termed a *production* method for studying temporal discrimination in animals, in the sense that the animal determines how long he waits. Animals and people can also be asked to *estimate* time intervals. For example, in the popular discrete-trial *bisection procedure* the animal is provided with two response alternatives (e.g., two pecking keys, for a pigeon), one signifying “too long”, the other “too short.” Each cycle of the procedure has two parts: In the first part the

Table 13.1: *Payoff matrix for signal-detection study of time estimation*

	Response	
	“Short”	“Long”
Time interval $>T$	C_1	V_2
Time interval $<T$	V_1	C_2

keys are dark and ineffective; after a variable period of time, t , the key-lights come on and the

animal must respond. If time t is less than some target time, T , a response on the left key (say) is reinforced with food and a response on the right key is either unreinforced or mildly punished by a timeout. If $t > T$, a response on the right-hand key is reinforced, and one on the left key punished. The advantage of this procedure is that the costs and benefits to the animal associated with different kinds of errors and correct responses can be explicitly manipulated, rather than being an accidental consequence of the interim activities that happen to be available (as in the fixed-interval situation). In this way, we can get some idea of the limitations on the timing process itself, apart from biases to respond or not respond associated with competition from activities other than the measured response.

This experimental arrangement is obviously well suited to signal-detection (ROC) analysis (cf. Chapter 9). Bias (the animal's criterion) can be manipulated either by varying payoffs, or by varying the relative frequency of too-long vs. too-short time intervals. The payoff matrix is illustrated in Table 13.1: V_1 and V_2 represent the benefits (reinforcement probabilities or magnitudes) associated with the two kinds of correct responses (correct identifications of "long" and "short" intervals); C_1 and C_2 represent the costs (timeout probabilities or durations) associated with the two different kinds of errors (long→short confusions and short→long confusions).




Figure 13.6. ROC plots for three pigeons trained on a procedure in which the duration of a time interval was judged; responses on a green key signaled "too long," to a red key, "too short." Abscissas show the probability of a green-key response given a sample interval shorter than the target duration; ordinates show the probability of a green-key response given a sample interval longer than the target duration. (From Stubbs, 1976, Figure 2.)

Figure 13.6 shows results from a pigeon experiment of this sort (Stubbs, 1976; see also the chapters by Stubbs & Platt in the volume edited by Zeiler & Harzem, 1979) in which pecking a green key was the correct response for a "long" stimulus and pecking a red key the correct response for a "short" stimulus. All three pigeons show ROC curves of the expected type, indicating that there is a stable limit on temporal discrimination that can be separated from the bias induced by payoffs —

although other experiments already discussed suggest that this situation is not as pure a measure of the limits of time perception as it might appear, because the intercalated long and short test intervals undoubtedly interfere somewhat with one another.

A very similar procedure, *delayed matching to sample* (DMTS), can also be used to study memory interference. In this procedure, a pigeon (for example) is confronted by three response keys. At the beginning of each cycle, only the center key is illuminated, with one of two stimuli, S_1 or S_2 . One or a few pecks on this *sample* key turns it off. After a delay of a few seconds, the two side keys come on, one showing S_1 the other S_2 . The animal's task is to peck the key showing the sample stimulus. A correct response yields food, an incorrect a timeout. After either event, the cycle resumes. The location of S_1 and S_2 in the choice phase varies unpredictably from trial to trial, so that the animal must recall the most recently presented sample to make a correct choice.

This procedure offers the same possibilities for confusion as the reinforcement-mission temporal-control experiments just described: On each choice trial the animal must be able to discriminate between the most recent sample (or time-marker) and earlier samples. This suggests that there should be fewer errors if the delay value is short or if sample stimulus duration is long, and both effects are generally found. One also suspects that performance is likely to be better if there is a substantial interval between trials, because each sample is then more widely

separated in time from preceding samples, which should reduce interference from them. This intertrial effect has not been shown reliably with pigeons with nontemporal sample stimuli, but experiments with rats have shown it (Riley & Roitblat, 1978; Lett, 1975). On the other hand, no experiment appears to have been done in which the intertrial interval varies from trial to trial. The earlier results with variable food and timeout duration on fixed-interval schedules suggest that intertrial-interval variation within each session may be necessary to get a reliable effect. Choice performance is more accurate the longer the preceding intertrial interval when the stimulus is itself a time interval, however (Riley & Roitblat, 1978).

An older version of the DMTS procedure nicely demonstrates that it is the memory for the *sample* that is important, not the subsequent choice arrangement. In an experiment by Harrison and Nissen¹¹ with chimpanzees the procedure was as follows: the animals were presented with two buckets, one covering food the other not. The animal saw the trainer baiting one or other bucket. After a delay, out of sight of the buckets, the animal was released and allowed to select one. In between baiting and choice, the buckets were moved closer together or farther apart. Nissen concluded that it is the separation of the buckets *at the time of baiting* that is important to accurate choice, not their separation at the time of choice. Evidently it is the way the sample is represented or *coded* by the animal that determines how well it can be responded to after a delay. If the baited and unbaited buckets are coded with very different spatial coordinates, they are not confused, even after substantial delays. But if their spatial coordinates (or other properties) are similar, then with lapse of time the animal cannot reliably distinguish between them.

As might be expected from the importance to recall of stimulus value, performance in this task depends upon the magnitude of the bait. If large and small baits are shown on different trials, accuracy is better on trials with the large baits even if the animal is always rewarded with bait of the same, small size.

Proaction and retroaction

The laws of memory define the limits on control of present action by past events. These limits are of two kinds: *proaction* effects, in which an earlier event (S_1) interferes with control of behavior by a later event (S_2), and *retroaction* effects, in which a later event interferes with control of behavior by an earlier event.




Figure 13.7. Perspective metaphor for the temporal resolution of memory. Top panel: the “mind’s eye” viewing past events B (long duration, remote past) and A (short duration, recent past). Bottom panel: Perspective projections of A and B and A' and B' , from the viewpoint of the present.

The degree to which one event interferes with control by another depends on two properties: the *similarity* of the two events (including similar time coordinates), and the difference between the behavior controlled by each. For example, suppose that S_1 normally elicits R_1 and S_2 normally elicits R_2 . In the delayed-match-to-sample (DMTS) situation S_1 might be a red sample key and S_2 a green key, R_1 would then be pecking the red choice key and R_2 pecking the green choice key. In DMTS R_i is always the next response required after S_i , so that only proaction effects are possible. Since the responses required by S_1 and S_2 are very different (S_1 and S_2 are not confused when both are present), interference in this situation depends upon the similarity of the

stimuli: the more similar are S_1 and S_2 , the worse the performance. If instead of red and green we used red and pink (say) as sample stimuli, we could expect choice accuracy to decrease: As we just saw, animals make more accurate choices when the two baited buckets are far apart than when they are close together. Response similarity makes much less difference than stimulus similarity, because the responses are usually available all the time and are not subject to interference or decay: the animal always knows what the response alternatives are on each trial, he is less sure about which alternative is correct (indeed, this asymmetry is a logical necessity, since the animal cannot even learn the task unless he can perfectly discriminate the choices). Knowing the choices requires only that the animal have learned a discrimination; knowing which choice to make requires in addition that he be able to remember which of the discriminated events *occurred last*.

In temporal-control experiments, however, the responses controlled by the interfering events can be either the same or different: in the FI reinforcement-omission experiments, F and N control the same behavior; in the first condition of the horizontal-vertical experiment (Figure 13.5, left), F and N (vertical lines) controlled different patterns. Interference was less when the two events controlled different patterns. Under some conditions, therefore, interference seems to depend on response as well as stimulus aspects of the task. The ways in which stimulus and response factors interact are yet to be fully worked out.

Proaction and retroaction effects in these experiments measure the *temporal resolution* of memory.¹² A useful analogy here is shown in Figure 13.7. This model assumes that past events are separately represented by the animal, and implies that the (temporal) similarity relations among events change continuously with the passage of time. The upper part of the figure shows the “mind’s eye” looking back over the record of past events, arranged in a time line, where the distance of the event from the eye (the present) is proportional to how long ago the event occurred. The bottom part of the figure shows how the events seen by the eye will appear on the “retina”. Let us suppose that this eye (unlike a normal visual system) perceives the size of events solely in terms of their retinal size. Then the “salience” of the various events is given by their projected size, as shown by the vertical line on the right in the bottom half of the figure. Thus, a long-lasting but remote event may appear only as large as a shorter but more recent event: B' is the same size as A' , even though B is an event of longer duration than A . Moreover, the relative sizes (salience) of events will change with lapse of time, i.e., as the vantage point moves to the right.

The various effects I have described are generally consistent with this metaphor. For example, events of longer duration are obviously easier to “see” than events of short duration. This fits in with the results of reinforcement-omission experiments in which longer events produce longer post-event pausing. Events widely separated in time are easier to tell apart (i.e., interfere less) than events close together in time; this is consistent with the effect of long intertrial intervals in improving DMTS performance. Moreover, the interference (proximity in the projected “memory image”) between adjacent events should increase with time. No matter how brief an event, at short delays it will appear more salient than any earlier event; but as time elapses, longer, long-past events will gain relative to shorter, more recent events — just as a mountain will loom over a house when both are viewed from a distance, but the house will blot out the mountain when the viewpoint is close to the house. As time passes, A' , the representation of event A , will therefore lose in size relative to B' , the representation of earlier, but longer, event B . This is a venerable principle in the study of memory, *Jost’s Law*: Given two associations of the same strength, but of different ages, the older falls off less rapidly in a given length of time.¹³ This principle accounts for the ability of even a not-very-salient event to control behavior over

brief time intervals.

Figure 13.7 shows that any limit on the animal's "visual acuity" means that in a given context, only a limited number of past events can be distinguished: This model therefore implies a limit to the capacity of *event memory*.

Hedonic value, as an important factor in memorability, can be represented in this analogy by the dimension of *height*: just as a distant high structure looms over a lower, closer one as the viewpoint recedes, so a preceding reinforcement seems to overshadow a neutral stimulus in the fixed-interval reinforcement-omission paradigm. Thus both confusion and overshadowing effects fit easily into the analogy. I show in a moment that this visual analogy corresponds to a decaying-trace view of memory.

Discrimination reversal and learning set

I end this section on methods of studying animal memory with an account of two procedures much used to study species differences in intelligence. These two procedures exemplify the memory processes already discussed, and also illustrate the fallacy of comparing species in terms of their performance on some task, rather than in terms of the processes that underlie performance differences — a variation on the theme first played in the discussion of probability learning in Chapter 8.

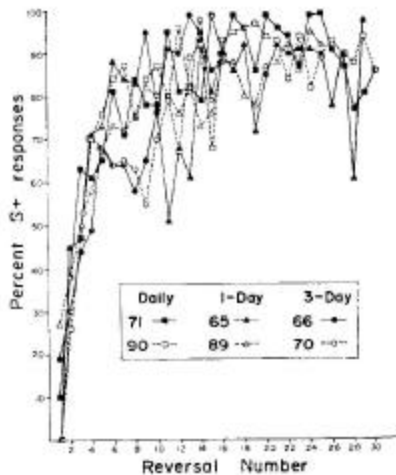


Figure 13.8. Performance of six individual pigeons exposed to daily reversal of a red-green successive discrimination, where $S+$ was reinforced according to a VI 60-s schedule, no reinforcement occurred in $S-$, and the change from $S-$ to $S+$ occurred only if no $S-$ response preceded the change by less than 30 sec. The animals had learned the simple red-green discrimination perfectly before being exposed to reversal training; this experience, plus the correction procedure, accounts for the close to 0% $S+$ responses on the first reversal. (From Staddon & Frank, 1974, Figure 1.)

Memory limitations enter in an interesting way into two tasks originally devised to study "higher mental processes" in animals: *discrimination reversal*, and *learning set*. Both tasks were intended to assess animals' flexibility by requiring them frequently to learn a new task, either a discrimination opposite to the one already learned (discrimination reversal) or a completely new discrimination (learning set). There are several versions of each procedure. One that has been used with pigeons is as follows.¹⁴ The animals are trained on a multiple schedule, familiar from Chapters 11 and 12. Two one-minute components occur in strict alternation. In one component, key pecks produce food according to a VI 60-s schedule; in the other, pecks are ineffective (extinction). The extinction stimulus ($S-$) changes to the VI stimulus ($S+$) after 60 s only if no peck has occurred in the preceding 30 sec; thus by pecking on $S-$ the animal can prolong its duration indefinitely. This "correction" procedure imposes a cost for responding to $S-$ — and also provides an additional cue to the identity of $S+$ each day (if a stimulus changes within 30 s of a peck, it must be $S+$). The stimuli are red and green key lights. After an initial period during which the animals learn a specific discrimination

(e.g., GREEN: VI, RED: EXT), the significance of the two stimuli is changed daily, GREEN signifying VI reinforcement on odd numbered days, EXT on even-numbered days.

The first question is, Do the pigeons improve in their reversal performance from one discrimination reversal to another? Figure 13.8 shows the percentage of “correct” (i.e., VI-stimulus) responses on the first day of each reversal for animals reversed every day, or less frequently. The results for all six pigeons are similar: a steady improvement in performance, settling down to perhaps 90% correct responses after several reversals. What does this result tell us about the flexibility of these animals’ learning processes? Two other results from this experiment — the effects of a shift to a new pair of stimuli and of days off — shed some light on this question.

After good performance had been achieved on the red-green reversal problem, the two stimuli were changed — to blue and yellow. The pigeons were given a total of 11 daily reversals with this new pair of stimuli. Then the animals were simply not run for a period of four days, then run for a single day, then not run for a further eight days. The effect on discrimination performance is shown in Figure 13.9: The animals performed quite well on the first day with the new stimuli, but discrimination was very poor the next day, i.e., on the *first reversal* after the change, and took several further reversals to recover almost to its previous level. In a similar way, the pigeons performed well after the four days off and after the next eight days off, but on the first *reversal* after the eight days off performance was poor and remained so for several subsequent reversals.




Figure 13.9. Average % $S+$ responses for six pigeons following a shift to blue-yellow reversal after extended training on red-green. (From Staddon & Frank, 1974, Figure 8.)

day, then performance each day will not be very good: the animal will always begin by responding a lot to $S-$ and not much to $S+$, thus ensuring a mediocre discrimination score. Conversely, if the animal treats the stimuli afresh each day but is poor at detecting $S+$ and $S-$, performance will also be inferior. The properties of memory are involved in the first prerequisite — control of behavior on day N by the significances established on day $N-1$.

We can get an idea of the relative importance of these two factors, memory and speed of learning, from the way that performance changes within and across experimental sessions, and from the results in Figure 13.9. For example, consider the animal that treats the stimuli afresh each day (this would be a beast capable of only local memory, in the sense of Chapter 4). If such an animal learns fast, then the absolute level of performance will be quite good, but there will be no improvement across successive reversals. In the case of this animal neither the days-off manipulation, nor the shift to a new problem (i.e., new pair of stimuli), should have any effect. Certainly, there is no reason at all to expect any special effect on the *first reversal*.

How about the animal that learns fast, but remembers the stimulus significances established on the previous day? The critical thing here, obviously, is *how well* the animal remembers. The previous discussion of the effects of stimulus intercalation makes some suggestions on this point. Early on, one might expect animals to remember quite well. In particular, on the sec-

ond day with any new pair of stimuli there will be no sources of interference and animals should remember well the significances established on the first day. But since the correct response is now the opposite one, we might expect to see especially poor performance on the second day, the *first reversal*, of any new problem — exactly what is shown in Figure 13.9. In short, anything that *improves* the animal's memory for what happened in the previous experimental session should *impair* discrimination-reversal performance.

Days off is such a factor. We might expect that the longer the time elapsed since discrimination N-1 — especially if discriminations N-2, N-3 etc. precede N-1 and provide sources of proactive interference — the smaller the effect the significances established then should have at the outset of discrimination N. So it proved in the experiment shown in Figure 13.9. After a days-off period, performance is slightly better than before the days-off period (reversals 12 and 13 in the figure): If the pigeons are run on Saturday (say), but then not run again until the following Thursday, performance on Thursday is good. But the very factors that minimize interference on Thursday from what was learned on Saturday act to *maximize* interference on the next day, Friday, from the significances established on Thursday. Thursday is temporally isolated from the discrimination sessions preceding it, so that its effect on Friday is unimpaired by the contrary stimulus significances established on the preceding Saturday. The result is poor performance on Friday.

This same line of argument leads us to expect that performance at the beginning of each experimental session should change systematically with successive reversals. Early on, the animal should respond incorrectly at the beginning of each session, responding most to S- (i.e., the previous day's S+). But with continued training, recall of the previous day's S+ and S- should be progressively impaired, so that at the beginning of each experimental session the animal should respond more or less equally, at an intermediate level, to both stimuli. This is more or less what happens with pigeons: at first, errors are high chiefly because the animal consistently picks the wrong stimulus at the beginning of each session. With continued training, this initial bias disappears and the animal appears more hesitant, responding at a slower rate, but more or less equally to both stimuli (presumably the hesitancy reflects the ambiguous status of both stimuli: the animal cannot, at this stage, recall which stimulus was S+ yesterday, but it has no difficulty recalling that both stimuli have served as both S+ and S-).

The discrimination-reversal task is not ideal as a test of “intelligence” in animals, because good performance can be achieved in several ways, not all of which correspond to superior ability. For example, poor temporal resolution of memory, i.e., a relative inability to distinguish yesterday's S+ from S+ the day before that, can aid performance on the task. It is possible to imagine three type of performance on the task, depending on the temporal resolution of memory: (a) At the lowest level, temporal resolution is exceedingly poor (this amounts to just local memory). Hence each day is treated as a separate experience, and discrimination-reversal performance is little different from simple discrimination performance. There should be no improvement across successive reversals. (b) At an intermediate level, temporal resolution is intermediate, hence discrimination-reversal performance is initially poor, but improves as proactive interference accumulates and weakens the effect of day N-1 training on day N performance. (c) At the highest level, temporal resolution is sufficiently good that the animal can show *spontaneous reversal*, using the S+ on day N-1 as a cue to S+ on day N. Spontaneous reversal is not possible at the two earlier stages, because late in training the day-N-1 S+ cannot be recalled, so cannot be used as a cue, on day N. (Spontaneous reversal may fail to occur even if memory permits because of built-in performance constraints: the animal may in some sense know that today's S+ is opposite to

yesterday's yet be constrained to respond to the most recent $S+$ — yesterday's — anyway, as in negative automaintenance.)

All three of these cases permit good steady-state reversal performance. They differ in the means used to achieve it — but these differences can be revealed only by appropriate tests. Available results suggest that goldfish correspond more or less to case (a), pigeons and rats to (b), and some higher primates to (c).

The learning-set task, in which a new pair of stimuli must be discriminated each day, seems like a better test of “learning ability”, whatever that might be, because the role of temporal resolution is minimized. Performance in this task does depend on something that might be termed “memory resolution”, however: the ability to keep separate (not confuse) a number of different pairs of stimuli, since if new stimuli are confused with old ones, then on some days the animals will pick the wrong one of the two new stimuli as $S+$ and this will retard acquisition of the discrimination. Thus, learning-set is subject to the same dual process as discrimination reversal: Any improvement across problems (i.e., pairs of stimuli) can reflect improvement in attending to the procedural features — differential reinforcement, correction — that signal $S+$ and $S-$. But improvement may also reflect increasing confusion among past stimuli with concomitant reduction in their ability to affect preferences for new stimuli.

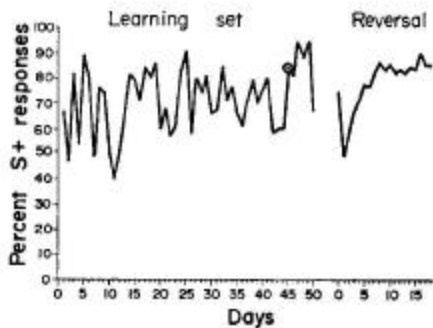


Figure 13.10. Performance of six pigeons on a series of daily learning-set problems after extended reversal training. Circled problem (no. 45) was also used for the reversal series on the right-hand side of the figure. (From Staddon & Frank, 1974, Figure 10.)

If the pairs of stimuli used each day are very different (i.e., the animals have good “memory resolution”), then animals will treat each day as a fresh discrimination. If the animals are pre-trained on discrimination reversal, they should transfer perfectly to such a task, having already learned how to identify $S+$ each day. Figure 13.10 shows learning-set performance of a group of six pigeons shifted to learning set after extensive experience with discrimination reversal. The performance of the group changed little across a series of 50 daily problems. Moreover, resumption of discrimination reversal again showed the first-reversal performance decrement and slow improvement required by the memory analysis.¹⁶

Final performance on the reversal problem was at about the same level seen when that problem was one of the learning-set series. All these characteristics — little or no improvement in learning-set performance after reversal training, first-reversal decrement after learning-set training, and similar performance within a learning-set series and at asymptote in a reversal series — are what we would expect from the properties of event memory just discussed.

MEMORY AND SPATIAL LEARNING

This discussion of reversal learning and learning set emphasizes a major difficulty in studying learning and memory: We see and can measure performance, but the same *performance* can usually come about in several ways, and even simple tasks call on more than one ability. Performance can never be taken at face value; we must always ask about the component abilities that make it up. An analysis into components can never be satisfactorily proved by a single “crucial” experiment. The best we can do is to take our hypothesized basic abilities or processes and show how by putting them together in various combinations we can bring together a wide range of

facts. The smaller the number of processes, and the larger the number of facts explained, the more reasonable the set of processes will appear.

In this last section, I take the principles of memory derived earlier and put them together with the notion of spatial representation discussed in Chapter 10. I use the resulting set of temporal and spatial principles to explain a wide range of experimental results with the radial-arm maze and other spatial apparatus.

The radial-arm maze

The¹⁷ multi-arm radial maze, first discussed in Chapter 10, has turned out to be an ideal arena for demonstrating the properties of memory in rats and other small animals. The standard radial-maze experiment involves an open 8-arm maze (Figure 13.11, left) with a baited food cup at the end of each arm. Hungry rats are allowed to choose freely among these arms, until either all eight have been visited or a fixed time (amply sufficient for eight choices) has elapsed. One or a few trials of this type are given each day until performance ceases to improve further. The usual finding is that rats soon learn to visit each arm almost without repetition, so that all eight arms are visited within the first eight or nine choices of a trial.




Figure 13.11. *Left:* eight-arm radial maze. *Right:* a maze with eight parallel arms used to assess the role of spatial encoding in a radial-maze performance.

Several experiments have shown that under normal conditions this efficient performance depends primarily upon the animal's memory for arms it has previously visited, rather than upon any kind of response strategy such as "choose the next clockwise arm".

I first describe the factors involved in radial-maze performance, then develop a model based on representation of events in the maze by a two-part code. The two parts are an identification code for places visited, and a temporal code reset by each visit (the reset assumption is not essential, however, and there is some evidence, discussed later, against it). The identification code is usually spatial (cf. Chapter 10); the temporal code has the properties described in the first part of the chapter. I then show how the major experimental findings with the radial maze and related T-maze tasks fit into this model. The last part of the chapter deals with apparent exceptions posed by findings from delay experiments with fewer than eight choices and more than a few trials per day. These exceptions can be accommodated by the assumption that a limited number of events are distinguishable in event memory, something implied by the kind of model illustrated earlier in Figure 13.7.

An overview of radial-maze performance

Performance in the radial maze seems to be determined by three factors: the coded values of past choices plus a *response rule*, response strategies, and "error", i.e., unaccounted-for factors.

These three factors represent three ways of solving the maze problem. The first method is the most satisfactory: Given some way to identify each arm, choices can be made in such a way as to avoid previously entered arms (appropriate for the radial-maze task), or select a particular arm (appropriate in most traditional maze experiments). Arms might be identified either by cues outside the maze (extramaze cues: room features) or by cues within the maze (intra-maze cues: arm color and texture, etc.). Intra-maze arm identification is more difficult, since the arms are usually similar -- and may also be less useful, since a code based solely on intra-maze cues

lacks spatial information. A code based on extramaze cues can be map-like, with information about adjacency and other spatial properties.

The second method, response strategies, is less efficient than either of the arm-identification methods, for two reasons: A given arm can be reached only after entering all the preceding arms in the sequence; and if the sequence is interrupted for some reason, succeeding arm choices will be inappropriate unless the sequence is resumed exactly where it was left off — which requires accurate memory. Since the main advantage of the response-sequence solution is that it makes minimal demands on memory, rats should depend on response strategies only when memory is impaired or arms cannot be accurately identified.

The third method, unsystematic choice, is not a solution, but serves to guide choice when other factors provide no guidance — early in training, for example.

The balance among these three methods is determined by factors such as stage of training, availability of extramaze cues, maze configuration, and the animal's information-processing capabilities.

Under usual conditions, behavior is guided by extramaze rather than intra-maze cues. This permits animals to develop a map-like representation of the maze, rather than responding in push-pull fashion to specific stimuli (Suzuki, Augerinos, & Black, 1980).¹⁸ Absence of extramaze cues (a maze with high walls on each arm, or with tube arms, for example) favors a more primitive representation, tied to intra-maze cues — or a response pattern. Rats seem to behave as this reasoning suggests, adjusting their relative reliance on extramaze, intra-maze and response-pattern factors so as to do as well as possible: Rats trained with few extramaze cues are more likely to show response patterning than rats trained under normal conditions, for example. The use of extra- or intra-maze cues depends upon memory: no matter how each arm is encoded, the animal must be able to distinguish entered from un-entered arms. Hence animals with impaired memories should show more reliance on response patterning, the only strategy left open to them. Young rats have poorer memories than adults and also show more response patterning.

Even when extramaze cues are available, and spatial coding is therefore possible, rats often enter arms in a systematic sequence. Experiments in which choices early in a trial are determined by the experimenter (forced-choice trials) show that these patterns are not necessary for them to learn the maze, however, and many successful animals show no obvious response pattern as they learn. Well-trained rats will abandon their response patterns, without loss of accuracy, if patterning conflicts with correct choice.

The radial-maze task constrains the order in which animals can learn different things. Since animals do not know the requirements of the task at the start of training, they must learn that only the first visit to each location (arm) is rewarded. To do so without relying on a response strategy, they must be able to identify each arm. They must also know whether an arm has been previously visited or not, which implies some form of temporal code. Only then can the animal apply the appropriate response rule based on this arm-identification and temporal knowledge. The response rule can neither be learned nor used effectively until the identification and temporal codes are relatively unambiguous.

A two-part code

Several studies show that rats can reliably determine which arms of the radial maze they have just visited, and that arms visited within the last few minutes are not confused with those arms visited on previous days. These findings suggest that animals encode two properties of each maze arm visited: the identity of the arm, and the time at which the arm was last visited (temporal location). The temporal code corresponds roughly to what Olton (1978), Honig (1978)

and others have termed *working* memory. The spatial code is one aspect of what Honig and Olson have termed *reference* memory.

Under normal conditions (ample extramaze cues, spatially separated arms) arms seem to be identified by a map-like spatial code. Temporal location can be represented by a decaying trace or *tag* (a temporal code). As each arm is visited, a temporal tag is attached to the spatial code for that arm; the tags for all arms decay at the same rate so that the *recency* with which any arm has been visited is given by the current value of its tag.

This scheme implies that rats represent the maze as a *list* of pairs, of the form $P_i T_i$, where P_i corresponds to the spatial (“place”) coordinate for the i th arm and T_i to the temporal tag for the i th arm. The reinforcement contingencies (e.g., all arms baited or only one arm baited) then determine how the rat responds with respect to his internal representation, his *response rule*. Let’s look at the properties of the spatial and temporal codes and the response rule.

Spatial code. Classic studies of human memory show the power of mnemonics prescribing that the to-be-remembered items be “stored” in spatially separated locations.¹⁹ The more widely separated the locations the better: items stored in different bureau drawers are more likely to be confused than items stored in different rooms in one’s house, for example — recall the sample-separation experiment by Harrison and Nissen, discussed earlier. These results strongly suggest that the spatial code incorporates information about spatial proximity: neighboring locations should have similar codes and be confused more easily than disparate locations.

The spatial coordinate may be bivariate, to reflect the two-dimensional structure of the maze, so that the animal’s experience with the maze is represented as a set of triples. It is likely that the form of the code depends both on the form of the apparatus — a hierarchical apparatus lends itself to a hierarchical code, for example — and on memory constraints such as the rodent equivalent of George Miller’s²⁰ “magical number seven”: humans cannot remember more than about seven unrelated items; no doubt rats are similarly limited.

Temporal Code. The temporal code carries information about *when* an arm has been visited. Performance in the radial maze and related situations can be accommodated by two assumptions: (a) when the rat visits an arm, a single trace variable is reset to a maximum value (or boosted by a fixed amount), and (b) that the trace decays with negative acceleration thereafter. The second assumption is not controversial: trace decay is old hat in theories of memory. The assumption that trace values are completely reset after each arm visit is less easily accepted because it implies that rats cannot learn to behave differentially depending on whether they have visited a place once or more than once: yet under appropriate conditions, rats can perfectly well discriminate the number of occurrences of a repeated stimulus. I point out later that in fact each arm visit is probably represented separately — but other memory limitations, built-in to the trace model, mean that we can treat the 8-arm radial maze as if memory traces are reset (i.e., as if each arm visit is represented only once). But the alternative assumption, that visiting an arm simply adds an increment to the trace value works about as well.

The form of the temporal code (trace) is constrained by two times: the time over which animals can distinguish entered from un-entered (i.e., never-entered, or last-entered-a-long-time-ago) arms, and the time over which they can distinguish the least-recently entered arm from the next-to-least recently entered arm — *absolute* and *differential* recency, respectively. Experiments show that the time of absolute recency is on the order of hours, of differential recency, minutes.

In one experiment, for example, delays ranging from 5 s to 24 hours were imposed between rats’ first and last four radial-maze choices. Choice accuracy over the last four was very high with delays as long as four hours, and above chance even after 12 hours. Evidently abso-

lute-recency discrimination is very good. Another experiment ran three successive trials each day, with only a minute or so between trials. Performance on the second and third trials was worse than on the first, showing that differential-recency discrimination is relatively poor (Roberts & Dale, 1981; Beatty & Shavalia, 1980).




Figure 13.12. Decay functions for the temporal codes for two arms, entered at times $-I$ and 0 . The absolute recency (AB) and differential recency (BC) change at different rates.

Absolute recency is determined by the animal's ability to discriminate whether the trace associated with a given arm is different from zero, differential recency by its ability to discriminate which of two traces is higher or lower in value. The relation between absolute and differential recency is illustrated in Figure 13.12, which shows the decaying traces associated with two arms, entered at times $-I$ and 0 . The levels of these traces at time t (point A) are denoted by heights AB and AC. Thus, height AB represents the animal's ability to discriminate an arm that has been entered from one that has never been entered (absolute recency) — height AC is the absolute-recency value for the newer trace. Height BC ($=AC-AB$) represents the animal's ability to tell which of two arms has been entered most (or least) recently (differential recency).

The simplest model for a memory trace is exponential decay, but this cannot account for Jost's law, one of the best-established memory principles. A simple function that is satisfactory (it corresponds closely to the perspective metaphor in Figure 13.7) is the hyperbolic function,

$$T(t) = 1/(a+bt), \quad (13.1)$$

where $T(t)$ is the trace value at time t after the event and a and b are constants².

Using the hyperbolic function, the absolute value of one trace and the differential recency value for two traces can be derived from Equation 13.1 (see Figure 13.12) as follows:

$$\text{Absolute recency: } AB = 1/(a + b(t + I)), \quad (13.2)$$

$$\begin{aligned} \text{Differential recency: } BC &= AC - AB = 1/(a+b(t+I)) - 1/(a+bt), \\ &= bI/(b^2t^2 + (2ab + Ib^2)t + a^2 + abI) \\ &= bI/(At^2 + Bt + C) \end{aligned} \quad (13.3)$$

where I is the time between arm choices, t the time since the most recent choice (OA in Figure 13.12) and A , B and C are lumped constants. An important property of the hyperbolic function is that differential recency decreases faster than absolute recency since the first is inversely related to t , the second to t^2 : in other words, the difference between two traces declines faster than the absolute value of a trace, as Figure 13.12 shows.

Response rule. The task in the radial maze is to find a particular place (spatial code) or find the place visited least- or most-recently (temporal code), depending on the reinforcement contingency. In terms of the processing the animal must carry out on his representation of the maze, these possibilities correspond to a different *sorts* of the list of time-place pairs that represent the maze: the animal can sort by either spatial or temporal code, depending on which is appropriate.

² The multiple time scale (MTS) memory model incorporates a trace that approximates the hyperbolic function used here and is consistent with Jost's Law (Staddon & Higa, 1999; Staddon, 2001b).

An animal's ability to learn any discrimination task involving delay will be directly related to the *discriminability* of the traces for the stimuli at the time when a choice is made. In delayed matching to sample, for example, the animal has to identify the stimulus seen *most recently*. We can be no more certain of the details of trace discrimination than of discrimination among sets of more familiar unidimensional objects such as wavelengths or loudnesses (see Chapter 10); but whatever the details, discriminability is obviously related to differential recency as defined by Equation 13.3.

In the radial maze, the probability an arm will be chosen is directly related to the discriminability of its trace value from the set of trace values for already-chosen arms. Under usual conditions (widely spaced trials) this means that the subject chooses the arm with trace value closest to zero (i.e., the *least-recently visited* arm). In terms of the list representation, the animal is always selecting the place-time pair with the lowest value of T . Consequently, pairs will always be confused because of the similarity of their temporal, rather than their spatial, codes. *Experimental results*

This model has three parts: spatial and temporal codes, and a response rule. Let's look at some experiments relating to each:

Response rule. Rats require little training to learn not to revisit arms in the eight-arm maze, but if the four arms entered in the first half of a trial are also the ones baited on the second half, they can learn (albeit more slowly) to repeat their first four choices rather than choosing the other four, unbaited arms. It's conventional to assume that these different performances reflect a difference in response rules, rather than a difference in the process by which arms are encoded, but it's important to realize that this is an assumption, not something that can be taken for granted. What's the evidence?

Under usual conditions (all eight arms baited) the response rule is that the spatial code with the oldest trace is selected over others (least-recent choice). There are three types of evidence for this rule: (a) Evidence that in spatial situations, rats behave spontaneously in accordance with the least-recent rule. (b) Results showing that tasks for which this rule is appropriate are learned rapidly, whereas tasks for which the rule is inappropriate are learned more slowly. (c) Results showing that even if the rule is appropriate, learning is rapid only if choices are spatially encoded, i.e., if extramaze cues are available and goal directions are varied.

Rats and many other animals have a spontaneous tendency to avoid places recently visited. This tendency was first noticed as *spontaneous alternation* in T-mazes, but the same tendency is exhibited as "patrolling" in residential mazes, in mazes with more than two alternative goal-routes — and in the radial maze. This least-recent tendency makes adaptive sense from two points of view: The least-recently visited place is the one where things are most likely to have changed. Consequently, if there is value in keeping up to date about the state of the world, the least-recent rule is the one to follow. For an opportunistic forager like the rat, many food sources will correspond to a random-interval schedule: depleted by each visit, and replenishing unpredictably with time. The least-recent strategy is optimal for exploiting such sources (see Chapter 9, Note 16).

Rapid learning in the radial maze is consistent with the least-recent rule: spatial encoding is ensured because the mazes are typically large and open, with ample extramaze visual cues, because the arms of the radial maze differ in two ways, direction and location, and because the goal boxes are widely separated. The usual reward contingency (bait in every goal) makes the least-recent response rule appropriate. This rule is also appropriate for the parallel maze (Figure 13.11, right), but here the arms differ only in location, and goal boxes are adjacent, making spa-

tial encoding difficult — and rats find this maze much harder than the radial maze. Spontaneous alternation does not occur in a maze with two parallel arms — further support for the idea that rats have a natural tendency to select the least-recently visited *place*.

Spatial effects. If information about spatial proximity is included in the spatial code, then anything that reduces the distance between maze arms should impair learning. Conversely, the learning even of very similar mazes should not be subject to interference, providing the mazes are in separate locations. There are experiments of both types.

Distance between maze arms can be varied either by making them adjacent, as in the parallel maze (Figure 13.11, right) or by reducing the size of the maze. Learning the parallel maze is much more difficult for rats than learning a radial maze. It is not yet clear whether the difficulty of the parallel maze reflects the physical proximity of the arms, their similar direction, or both.

Rats can learn radial mazes in different places without interference, up to a limit set by memory capacity. For example, one experiment showed that exposure to a second radial maze halfway through a trial on a maze produces subsequent performance decrements on the first maze only when the second maze is in *exactly* the same position as the first (not next to it or even two feet above it), and when there are several intervening trials (on identical mazes in similar rooms) before the interrupted trial is completed. The first result is expected from the spatial-coding hypothesis: mazes in different places don't interfere because their arms have different spatial codes.

The interfering effect of a large number of intervening mazes may reflect memory-capacity limitations: each room provides a different context, and as we have already seen, separating potentially interfering events by means of different contexts can reduce interference. But when the number of contexts gets very large, memory limitations are a factor. For example, if an animal is trained with the same maze in three separate rooms, he might encode each arm in the form R_iA_j , $i = 1-3$, $j = 1-8$, where R denotes the room code and A the arm code. Since the number of rooms is small, we may expect few confusions among the R -codes, hence no interference among mazes. If the number of rooms is large, however, the number of room codes may exceed the rat's memory span for unrelated items and confusion will occur. These results suggest both a context-specific code, in which each arm is related to a particular maze (whose center may be coded in terms of absolute spatial location), and that the rats can restrict their trace sort to one context.

A puzzling feature of radial-maze experiments is the complete failure to find evidence for spatial, as opposed to temporal, generalization. For example, several experiments have found that when rats make mistakes — that is, reenter already-chosen arms — they do not selectively repeat arms near to correct (i.e., unchosen) arms. If locations are identified according to a spatial code, why don't spatial confusions occur?

The present model gives a straightforward answer to this question: If the maze is represented by a bivariate code, and arm choices are made by sorting place-time pairs on the basis of the *time* coordinate, errors *should* have nothing to do with the place coordinate. The only way that spatial generalization can occur is in the mapping of the animal's spatial representation onto the physical maze. For example, if the animal's identification of "true north" deviates by more than 22.5° from the actual north, then in selecting the coordinate for the north arm he may actually enter the NE arm. Granted the multiplicity of extramaze cues typically available, it is unlikely that rats make errors as large as this in relating their spatial representation to the actual maze. Consequently, spatial generalization in the radial maze is not to be expected.

But what of the parallel maze? We know that rats find this much more difficult to learn than the radial maze, so we might expect measurable errors in their mapping of their representation onto the actual maze. Yet here also there is no evidence for spatial generalization. One answer is that behavior in the parallel maze is not guided by a spatial representation. The rats seem to be guided by intra-maze cues and a nonspatial representation in preference to a spatial-map method made unreliable by the spatial proximity and identical direction of the arms.

Evidence for a nonspatial representation in the parallel maze comes from experiments on spontaneous alternation and with blind and sighted rats. I have already discussed evidence suggesting that spontaneous alternation is conditional on the formation of a spatial representation, and noted that rats do not show spontaneous alternation in a two-arm parallel maze. Sighted rats learn the radial maze much faster than blind rats, but show no such superiority in the parallel maze. Under normal circumstances (i.e., when they have not been trained before being blinded)

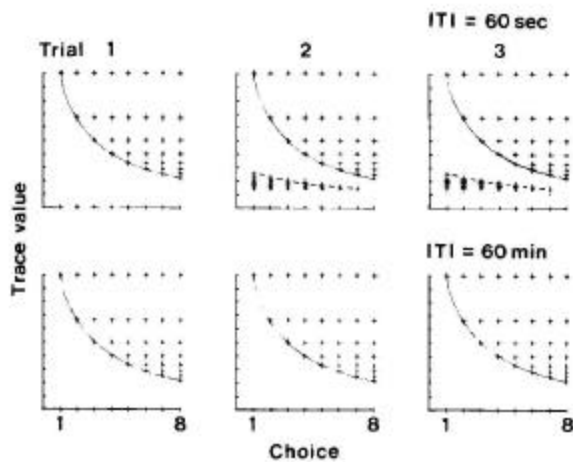


Figure 13.13. Trace values for each of the eight arms of a radial maze immediately after each of eight choices on three successive trials. The trials are either 60 s apart (upper) or 60 min apart (lower). Choices within a trial were always 10 s apart. The solid lines indicate the lowest trace value among the arms in the “previously chosen” set; the dashed line shows the highest trace value among those arms yet to be chosen on a trial.

blind rats do not seem to use a spatial representation in either type of maze. The similar performance of blind and sighted rats on the parallel maze implies that even sighted rats do not encode the maze spatially. If the parallel-maze arms are not encoded spatially, spatial generalization is not to be expected even in sighted animals.

But it is not even necessary to invoke differences in the way that radial and parallel mazes are represented. The lack of spatial-generalization errors follows from the task constraint that gives priority to arm identification: The animal cannot apply the temporal tag and relevant response rule *until he can recognize an arm without error*. Consequently, if the maze problem can be solved at all, and response strategies are precluded, arm identification *must* be close to 100% accurate. If extramaze cues are degraded so that a spatial code becomes inaccurate enough to show spatial-generalization errors, performance based upon such a code cannot approach typical levels of accuracy.

Hence accurate performance under these conditions implies a *nonspatial* code. In either case, if the animal can learn the task at all, he will not show spatial-generalization errors.

Temporal effects. The model assumes that the temporal tag associated with each maze arm is initially zero, is reset to a maximum on each visit, and then decays hyperbolically with time thereafter (Figure 13.12, Equation 13.1). The pattern of trace values to be expected on the first and two subsequent trials within a day is illustrated in Figure 13.13. The curves in the figure are based on the assumptions that the rat chooses an arm every 10 s during a trial, and that the inter-trial interval is either 60 s (top three panels) or 60 min (bottom three panels). Parameters a and b in Equation 13.1 were determined by Equations 13.2 and 13.3 so that absolute recency and differential recency are in approximately the right ratio (roughly 60:1). These numbers are not critical; the figure shows a pattern to be expected under a wide range of conditions.

The +s and curves in Figure 13.13 were derived as follows. Look at the upper-left panel, which is typical of all. The leftmost upper + represents the maximal (reset) trace value associated with the first arm entered on this first choice of the first trial of a day. The descending curve beginning from that point shows the hyperbolically decaying trace value associated with that arm; this line is the value of the *oldest* trace among all the arms visited on that trial. The second + from the left in the upper left panel represents the reset value for the trace of the second arm entered; this trace value also decays with time, as indicated by the +s just above the solid line showing the oldest trace. The other +s similarly represent the trace values for the other arms entered.

The dashed line below the solid line in upper panels two and three represents the decaying trace of the arm entered last on the preceding trial; the +s below it represent traces for arms entered earlier. The difference between the solid and dashed lines provides a measure of the recency discriminability of chosen vs. unchosen arms: the closer the two lines, the harder it should be for the animal to distinguish arms chosen on this trial from unchosen arms (i.e., arms chosen on the previous trial).

Discriminability (the separation of the solid and dashed envelopes) decreases during each trial, and is much lower on the second and third trials when trials are 60 s apart, but not when they are 60 min apart. The first prediction is consistent with the universal experimental finding that the probability of repeating an arm increases during a trial, even when the increased opportunities for repetition are corrected for. The second prediction is consistent with the finding that choice accuracy decreases between Trials 1 and 2 when the trials occur 1-2 min apart, but not when they are separated by an hour. Figure 13.13 also predicts that choice accuracy will not decrease any further after Trial 2, since the trace values are the same at the start of all trials after the first. This surprising prediction is also confirmed by data.

Figure 13.13 shows that discriminability is always high after the first choice of a trial, even if the trial follows quite closely on a preceding trial. Figure 13.13 also suggests, and the data confirm, that at the start of a trial, rats should prefer arms chosen early in the preceding trial. Both these results argue against “resetting” and in favor of some sort of boost given to the trace on each arm visit.

Figure 13.13 also shows that if trials are well spaced, discriminability depends only on choices already made (i.e., nonzero trace values) not on choices yet to be made, which will all have zero trace values. Hence, rats should show the same accuracy on the first N choices of any radial maze, no matter how many arms it may have (and so long as it can be learned at all). In confirmation, experimental results show that accuracy over the first eight choices on a 17-arm radial maze is the same as choice accuracy over all eight choices on the 8-arm maze.

After eight correct choices have been made, the set of trace values appears as shown in the rightmost column of +s each panel in Figure 13.13. The least-recent response rule implies that the animal should have difficulty deciding on his ninth choice because the oldest traces are all close together in value: the set of traces does not divide easily into a set of chosen and a set of unchosen arms, so that discriminability is low. This suggests that animals should hesitate after visiting all eight arms in a trial. After the ninth choice has been made, however (first column of +s in each panel), the traces once again divide up into two relatively discriminable sets because the trace for the arm just chosen is at a high value and the other seven are relatively low; hence, choice time should decrease. Experimental results fit in with this picture: Rats typically spend only a second or so in the center between arm choices, until all arms have been chosen; then they wait in the center of the maze 30 seconds or more before making the ninth choice. But after the ninth choice, arms are again chosen at a high rate.

Other spatial situations

The notion of an identification code, a resetting or visit-incremented, hyperbolic trace and the least-recent response rule accounts for most radial-maze experiments, and experiments on spontaneous alternation, that have used one or just a few trials per day. These assumptions also account for many results from delayed-alternation experiments, and, if the response rule is suitably changed, for results from delayed-reaction and delayed-matching-to-sample experiments as well. The theory runs into difficulties when the number of choices is small, as in the T-maze, and when there are many trials per day, so that the ratio of intertrial interval to retention interval is relatively small. The critical problem here is the resetting assumption. I first show how the theory works, when it does; then show where it fails to work — then reconcile the two.

Animals have difficulty learning a discrimination when reward is long delayed after the correct choice. Lett (1975; reviewed in Revusky, 1977) has done a series of experiments in which rats learned to either return to, or avoid, one arm of a T-maze, even though they were rewarded only after delays of from one minute up to an hour or more. For example, in one experiment, Lett gave rats one trial per day on a simple T-maze discrimination, removing the animals to their home cages as soon as they entered the correct arm, and then rewarding them in the start box after a 60-min delay. Over a period of ten days or so, the rats learned to perform correctly.

Two features of the experiment seem important to its success: removal of the rat from the goal box as soon as it entered, and the long intertrial interval in relation to the delay interval. Removal from the goal box seems to be important to minimize context-specific retroactive interference from activities that may occur in the goal box after it is entered. The long intertrial interval is necessary for the animal to be able to discriminate the most recent goal entry from previous goal entries: the task in this experiment is to discriminate between the trace of the correct goal box, visited 60 min ago (the retention interval), and the trace of the wrong box, visited perhaps as recently as the previous trial (i.e., with a time separation equal to the intertrial interval plus twice the retention interval): the animal must discriminate between two traces originating t and $2t + I$ s ago. Equation 13.3 shows that discriminability of one trace from another is a positive function of intertrial interval (I), with asymptote inversely related to the retention interval (t). Hence, the model suggests that animals should have no difficulty in delayed-reward experiments so long as the intertrial interval is large in relation to the retention interval, as it was in the Lett experiments.

In delayed-alternation experiments, animals are forced to visit one arm of a T-maze then, after a delay, are rewarded for visiting the other arm. The required rule is exactly the same as in the radial maze, but the number of arms is smaller. Despite the similarity of the tasks, delayed-alternation performance typically falls to chance after delays much shorter than the several hours required to produce impairment of radial-maze performance. An important variable here seems to be the intertrial interval, which is typically less than the 24 hours usual in radial-maze studies. The shorter time would lead one to expect more rapid impairment in the delayed-alternation task. Yet this cannot be the whole story, because (as we have already seen) radial-maze performance is quite good even when trials are only 60 s apart. Why should rats be able to discriminate arms visited 30-s ago from arms visited 90-s ago in the radial maze, but not the T-maze?

There is a straightforward explanation for this apparent paradox. The perspective metaphor in Figure 13.7 implies that each visit to a maze arm is separately represented by the rat, but that (for a given context) there will be some limit to the number of visits that can be separately

represented — simply because as time passes, past events are bunched closer and closer, so that at some point they must cease to be discriminable. This limitation is part and parcel of the trace model: Equation 13.3 shows that the differential recency of any pair of events decreases very rapidly with the passage of time. We might also expect that the larger the number of different traces (events), the more difficult it will be to discriminate one from the others. Thus, both *time* and *number of intervening events* should affect animals' ability to identify a particular trace. The data seem to agree: both time and number of intervening choices have been shown to affect radial-maze performance, for example.

Suppose that this limit on event memory means that after perhaps 6-8 visits to the same arm or different arms, information begins to be lost. In the 8-arm radial maze, this will occur at about the time when every arm has been visited. On a succeeding trial, therefore, as arms are revisited, information about prior visits will be lost, precluding any interfering effects from such visits. The effect of this capacity limitation on event memory is that only the most recent visit to an arm will be represented, so that traces will appear to be completely reset after every tour of the radial maze.

Not so in T-maze delayed alternation: Here several visits to each arm will be separately represented, so that the animal cannot just sort the set of place-time pairs by the time coordinate. For example, suppose we denote the two arms by *L* and *R* (T-maze) or *A-H* (radial maze) and the trace values by numbers. Then the radial-maze task, after seven choices in a day, is to sort a list of pairs A99, B95, C90, D87, E83, F79, G76, H2, say, by their trace values alone: In this example, 99 (close to 100%) is the value of the most-recent trace, and 2 the value of the least-recent (the last arm chosen on the previous day): 2 would be selected, and H would be the correct choice. In T-maze delayed alternation, after the same number of choices (i.e., three and a half trials, each a forced choice followed by a free choice, and assuming no errors) the set of pairs might be L99, L95, R90, L87, R83, L79, R76, L2 (L2 is again the last arm chosen on the previous day). The animal cannot solve this problem merely by picking the oldest trace and choosing the associated arm. He must identify the most-recently visited arm, and then choose the other — a much more difficult task, for two reasons: Rats find the least-recent-choice rule much easier than most-recent choice. And the trace-value difference between the most-recently visited arm and the arm visited next-most-recently (4 in the illustration) is necessarily much smaller than the difference between the oldest trace and the next oldest (74).

Thus, the poorer performance in T-maze tasks vs. radial-maze tasks may reflect response-rule and recency-discrimination limitations that combine to favor the radial maze: The animal's limited event memory, together with the large number of arms in the radial maze and the typically long ITI, protects performance from interference from remote trials; whereas the smaller number of arms in the T-maze, and the typically short ITI, ensure proactive interference.

Delays on the order of seconds (not minutes or hours) between sample and choice stimuli severely impair the performance of pigeons on delayed matching to sample. DMTS experiments can be explained in the same way as delayed-response, with two added difficulties: spatial stimuli, such as two maze arms, are almost certainly easier to remember than the colored lights and patterns typically employed (with pigeon subjects) in DMTS studies; and the animals have no built-in response rule appropriate for the reinforcement contingencies, as rats do in delayed-alternation or the radial maze. Taken together, these factors seem ample to account for the poor performance of most animals in the DMTS procedure with even quite moderate delays.

SUMMARY

Memory in the most general sense is implied by any difference in behavior that can be traced to a

difference in an organism's past history. Control of behavior by a time marker (temporal control) is perhaps the simplest example of a memory effect. I showed in the first part of the chapter that control by a time marker can be impaired by both prior and subsequent events — interactions corresponding to the proactive and retroactive interference much studied in human and animal memory experiments. These interactions limit animals' ability to discriminate the absolute and differential recency of events. This memory for events is sometimes known as *working* memory.

Recency discrimination depends on the particular properties of time markers. Hedonic stimuli, such as food and electric shock, are particularly effective as time markers. Hedonic events preempt the animal's attention so that a remote hedonic event may control behavior even though a more recent, but less salient, event is a better predictor.

When interference is minimized, neutral events can serve as time markers, and they then behave in the same way as synchronous stimuli: excitatory temporal stimuli show decremental generalization gradients, inhibitory stimuli incremental gradients. This aspect of memory, encoding of the properties of particular stimuli, is sometimes known as *reference* memory.

Event recency can be represented by a trace variable that decays in negatively accelerated fashion with elapsed time. Such a model implies a context-specific limit on the number of events that can be separately represented in event memory. Animals learning tasks such as successive discrimination reversal and delayed matching to sample seem to behave as this model suggests: reversal performance improves with practice, is unaffected by a lapse of several days or by shifting to a new pair of stimuli, but is usually impaired on the first reversal after the shift.

The latter part of the chapter showed how trace discrimination, together with spatial encoding of maze arms (a bivariate spatiotemporal code) can explain a variety of experimental results with the radial maze and related apparatus. When every arm is baited (in the radial maze) or when alternation is rewarded (in T-maze delayed-alternation experiments) animals appear to "sort" the bivariate codes representing arm visits according to the trace value, selecting the spatial coordinate associated with the oldest trace.

This model explains why the learning of mazes in different spatial locations does not interfere and why there is no spatial generalization in either radial- or parallel-maze experiments. It also accounts for the dependence of maze performance on the time and number of intervening choices, for the similarity of performance on successive trials after the first within a day, and for pauses within a trial after each arm has been visited once.

Rats in the 8-arm radial maze behave as if each arm visited is represented only once. This "reset" property of the trace conflicts with other memory results showing that repetition affects performance in delayed-reaction and DMTS situations. A resetting-trace model also cannot explain why delayed-alternation performance in a T-maze is much more sensitive to retention-time delays than performance in the radial maze, since the two tasks are identical save for the number of choices involved. One solution to this contradiction is provided by the limit on event memory implied by a trace model. The number of arm visits within a trial in the radial maze exceeds this event limit so that a trial interferes only with the following trial and not with later trials — this is why the resetting-trace model works for the radial maze. The smaller number of different-arm choices in the T-maze means that unless the retention interval is short relative to the intertrial interval, multiple visits to the same arm are represented, complicating the animal's task and impairing choice accuracy. This analysis accounts for the effects of sample and non-sample repetition, and for the effects of different times between repetitions, in delayed-alternation experiments — as well as for the sensitivity to delay of delayed-response and DMTS performance. But a small change from resetting to the idea that visiting an arm simply gives the trace a boost

also leads to a model adequate to explain all these results, although I do not present all the details here.

The discussion of memory in this chapter begins to get at the general problem of complex learning. From this perspective, learning can be thought of as the change in an animal's performance brought about by the change in its representation of past events and their significance — where the change in representation depends both on the animal's behavior (acting through a feedback function to produce changes in the environment) and the passage of time. The next two chapters deal directly with learning in this sense.

NOTES

1. Schneider (1969). The book *Time in Animal Behaviour*, by Richelle and Lejeune (1980) provides a comprehensive review of work on temporal control as well as some discussion of circadian rhythms in animals. See also Staddon (2001b), Chapters 13-15.
2. The data support the assertion that error is proportional to the mean interval being timed, but this follows from the clock model only under certain conditions. For example, if the clock rate is constant from reset to reset, but varies randomly around the true mean from one reset to the next, then error will be proportional to the timed interval. On the other hand, if clock rate varies randomly from moment to moment, mean error will not bear such a simple relation to mean interval timed. This process can be analyzed as a *random walk*, for which the expected deviation will be proportional to the *square root* of the timed interval. Scalar timing, the evidence for it and its theoretical properties, is discussed at length by Gibbon (1977); see also Staddon & Higa (1999) for a critique and alternative.
3. This experiment, and other work related to the clock hypothesis, is described in Church (1978; see also 1980, Staddon & Higa (1999) and Staddon (2001b).
4. In Chapter 11 I noted that behavioral competition can improve discrimination. Effects that seem to reflect competition, such as local contrast, increase when stimuli are made more difficult to discriminate or stimulus control is weakened in other ways. It is possible, therefore, that the animal uses interim activities as a way of sharpening discrimination performance. Since time is a continuous dimension, recognizing the time-of-arrival of periodic events poses special problems. This difficulty may perhaps account for the special strength of interim activities (especially schedule-induced interim activities) on periodic schedules. Schedule-induced interim activities, such as drinking by rats, are *not* seen in the *S*- component of multiple schedules, for example; they seem to occur only in the temporally defined *S*-s on periodic schedules.
5. The difficulties encountered in obtaining intra-dimensional generalization gradients of temporal control, and some examples of such gradients, are described in Staddon (1975).
6. The Catania and Reynolds (1968) monograph from which this figure is taken describes a mass of data on the effects of postfood reinforcement probability on the pattern of postfood responding by pigeons.

7. The postfood pause on fixed-ratio schedules. It is tempting to apply this temporal-discrimination account to fixed-ratio schedules also: the postfood pause on FR has usually been explained by the roughly fixed time between food deliveries that emerges as a secondary consequence of the approximately constant rate of responding animals show on ratio schedules. Once the time between food deliveries stabilizes, the argument goes, a fixed-ratio schedule looks like a fixed interval, and the animals therefore begin to pause after food. The problem with this argument is that if the schedule looks like FI, the animals should also slow down (because typical response rates on FI are lower than rates on comparable FR, cf. Chapter 7). But if they slow down, the interfood interval should increase, so the pause should increase, and the result should be a very unstable pattern. This explanation also assumes that animals are in some sense unable to tell the difference between interval and ratio schedules, even though features of their behavior other than the pause are quite different on the two schedules. More recently, Wynne and Staddon (1988; see also Staddon & Cerutti, 2003) have suggested that both FI and FR pausing results from an obligatory timing process they term *linear waiting*.

The satiation-deprivation process discussed at the end of the last chapter provides yet another explanation for the postfood pause on ratio schedules. Satiation-deprivation implies changes in the competitiveness of interim activities early and late in the postfood interval. The instrumental (terminal) response necessarily occurs just before food. Hence, the tendency to engage in interim activities will always be highest *after* food, i.e., after a period when they have not been occurring.

Animals pause much less on variable-ratio schedules and this has usually been attributed to aperiodicity of food deliveries. But animals also respond somewhat more slowly on VR than during the “running rate” on FR, so there is time for interim activities after each response or group of responses, diminishing the tendency for them to occur after food (the same argument accounts for lack of postfood pausing on VI). The small pauses sometimes seen on VR may reflect the time taken up by eating (when interim activities cannot occur). Postfood pausing on VI is negligible because the tendency to make the instrumental response is weak just after a response (reflecting the lowered probability of food at that time), allowing time there for the occurrence of interim activities and correspondingly weakening their tendency to occur after food.

One prediction of this analysis is that the postfood pause on fixed-ratio schedules should be reduced by postfood timeouts longer than the typical pause, and this seems to be the case (Richards & Blackman, 1981), unless the timeouts are very long (which may produce counter-vailing increases in food motivation).

8. This point was confirmed experimentally by Hatten & Shull (1983).

9. Staddon (1974). Context effects in conventional animal-memory experiments have been shown by Roberts (1972) and Grant (1980).

10. Overshadowing. The term *overshadowing* is originally due to Pavlov (1927), who used it to describe the effect of an intense or salient stimulus element in gaining control at the expense of a less salient element. For example, if an animal is trained to discriminate between a positive stimulus consisting of a loud sound and a dim light and a negative stimulus consisting of darkness and silence, then in a test it is likely that the sound will be the most important controlling element: the animal will *attend* (in the sense the term was used in Chapter 10) primarily to the sound dimension. Later classical conditioning experiments have shown that a similar effect can be produced by pretraining: If an animal is first trained to discriminate between a tone and silence, then a light is added to the tone, little or no control will be acquired by the light. This is termed *blocking* and I discuss it later.

Overshadowing by hedonic time markers obviously conforms to Pavlov's usage, since it is the salience of these stimuli that leads to their prepotency, rather than any kind of pretraining. The temporal usage is different in that it refers to a phenomenon that can be repeatedly demonstrated in "steady-state" responding in individual animals, rather than to differences in the ease with which some behavior can be *acquired*. There may be no difference at the level of fundamental processes, however, since the ease with which a discrimination can be learned obviously depends upon the memorability of the stimuli to be discriminated.

11. See Fletcher (1965) for a review of a review of this and related experiments on animal memory. Roberts (1972) and Medin (1969) have reported similar findings.

12. D'Amato (1973) seems to have been the first person to propose that many memory effects can be explained in terms of the discrimination of recency.

13. This is actually Jost's second law; the first is "Given two associations of the same strength, but of different ages, the older one has greater value on a new repetition." Both refer to the relative gain in strength of old vs. new associations. This version is from Hovland (1951, p. 649), which also contains the original Jost (1897) reference. Jost's law implies that any decay component in memory cannot be simply exponential, at least if the same decay rate is assumed for all memories (Simon 1966).

14. Staddon & Frank (1974); see Mackintosh (1974) for a review of the numerous studies on discrimination reversal. The learning-set task was devised by Harlow (1949) as a test of intelligence in subhuman primates; it became a standard procedure in work with primates, but is now little used. Miles (1965) has reviewed much of this work.

15. An even better possibility is to *remember* yesterday's $S+$, and then respond to the other stimulus, so-called *spontaneous reversal*. This requires excellent memory on day N for the stimulus significances established on the previous day, and seems to occur in some higher primates. It also requires an ability *not* to respond to the best remembered rewarded stimulus. The best that most animals can do is to disregard prior stimulus significances. I discuss these differences later.

16. The first-reversal decrement shown in Figure 13.11 proves that the pigeons were able to keep the various stimulus pairs distinct, and that lack of interference across days was due to this rather than to complete confusion about past stimuli. If the animals had been completely confused each day about the identity of $S+$ on the previous day, then there would have been no decrement on the first reversal because the animals would not have been able to recall the identity of the previous-day's $S+$ (because of proactive interference from earlier, presumably similar, $S+s$). They did show a first-reversal decrement, hence must have been able to recall on reversal 1 $S+$ on day 0, even though that $S+$ was preceded by 50 days with novel pairs of stimuli each day.

17. The material in this section is based on an unpublished theoretical paper by Dale & Staddon (see also Staddon, 1985). The original experiments on the eight-arm radial maze were done by Olton (1978; Olton & Samuelson, 1976). See Staddon (2001b, Chapter 12) for a real-time theoretical approach to spatial search.

18. The original idea that animals get around with the aid of *cognitive maps* is due to the Berkeley psychologist E. C. Tolman (1886-1959). Tolman proposed what would now be called an in-

formation-processing theory of animal learning that made much use of concepts such as *expectancy*, *means-end-readiness* and the like, whose descendants (conscious or otherwise) are to be found in the writings of human psychologists such as Simon (e.g., 1979). Tolman's best known works are *Purposive Behavior in Animals and Men* (1932) and the contribution to the collection edited by Koch (1959).

19. These mnemonics are entertainingly reviewed in Crovitz (1970) and Yates (1966).

20. Miller (1956) wrote a paper, now a classic, entitled "The magical number seven, plus or minus two: Some limits on our capacity for processing information" in which he reviewed a range of experimental results pointing to this limitation on *immediate-memory span*.