11

STIMULUS CONTROL AND PERFORMANCE

The last chapter concluded that animals develop an internal representation of their world that guides action. We are uncertain in most cases both about the properties of such an internal representation, and the effects of reward and punishment on it. The most parsimonious assumption is that the representation of simple objects is independent of reward and punishment: things look as they look, whether or not good or bad consequences are associated with them. Search images (*concepts,* in psychological jargon) may be an exception to this. Representations of very complex objects may perhaps be acquired only through a history of explicit reinforcement. Medieval teachers believed that Latin is learned only through the birch, and this general view of the motivation required for complex learning was almost universal until recently. Still, for recognition of simple stimuli, no special training seems to be required. The effect of reward and punishment is to give value to certain objects or places, as represented, rather than to create or modify the representations themselves.

Is performance then determined solely by the animal's external environment, as internally represented? In this chapter I argue that there is at least one other factor that must be taken into account: competition among activities for available time. These two factors, competition and external stimuli, taken together account for numerous experimental results on generalization and discrimination. The rest of the chapter explains how competition and stimulus control contribute to discrimination, behavioral contrast, generalization and peak shift.

INHIBITORY AND EXCITATORY CONTROL

Animals need to know both what to do and what not to do; hence stimuli can have both inhibitory and excitatory effects. But, as we saw in Chapter 7, when an animal is not doing one thing it is probably doing something else. Moreover, animals are highly "aroused" under the conditions typical of operant conditioning experiments — hunger, combined with frequent access to small amounts of food. Behavioral competition is then especially intense: the animals have a lot to do and limited time in which to do it. A stimulus that signals the absence of food (S-) not only lets the animal know that he need not act in ways related to food, it also tells him that other activities are free to occur.

Since different activities compete for the available time, it is difficult to decide whether a stimulus that suppresses activity *A* acts directly on *A*, or indirectly by facilitating some other, antagonistic activity. Perhaps the question is not even a useful one (although it has exercised heavy thinkers in this area for at least two decades). It *is* clear that a stimulus that acts to suppress activity *A* changes the balance of behavioral competition in favor of a complementary class of activities, ~*A*. As we shall see, it is possible to demonstrate direct, excitatory control by *S*- of activities that compete with the food-related behavior controlled by *S*+.

Feature effects

There are also striking experimental results showing how difficult it is to establish inhibitory stimulus control when competing activities are weak and stimulus generalization tends to facilitate the response to *S*-. For example, in a series of classic studies, Jenkins and Sainsbury (1970; see also Hearst, 1978; and Hearst & Jenkins, 1974) trained pigeons on a simple discrete-trials discrimination procedure of the following sort. At random intervals averaging 60 s a response key is illuminated with one of two stimuli (*S*+ or *S*-). Four pecks on *S*+ produce food; four pecks on *S*- turn off the stimulus but have no other effect. In either case, if the animal fails to make four pecks, the stimulus goes off after 7 sec.

The birds quickly learn not to peck during the dark-key intertrial interval. Do they also learn not to peck *S*- (which amounts to inhibitory control by *S*-)? Not always, and the conditions when they do and don't suggest the importance of behavioral competition.

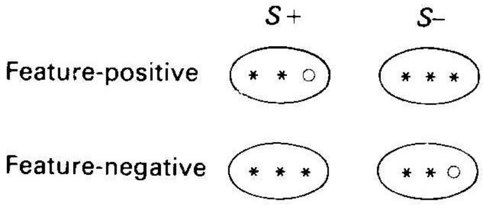

Figure 11.1. Examples of feature-positive and feature-negative stimulus displays.

Figure 11.1 shows two kinds of *S*+, *S*- stimulus pair that Jenkins used: The upper pair is termed *feature-positive*, because the feature distinguishing *S*+ from *S*- appears on *S*+. The lower pair is *feature-negative*, since the distinctive feature appears on *S*-. Pigeons easily learn the feature-positive discrimination, pecking on *S*+ and not on *S*-. But they *fail* to learn the feature-negative discrimination, pecking indiscriminately on *S*+ and *S*- (this is termed the *feature-negative* effect). Why?

There seem to be three reasons for the feature-negative effect. One is perceptual: the displays in Figure 11.1 consist of separate elements; their dimensions are separable, not integral (see Chapter 10). Moreover, *S*+ and *S*- share elements in addition to those in the display: they are usually on for about the same time, and they both appear on a white-lighted key (as opposed to the dark key during the intertrial interval). Separability means that the elements are responded to separately; *S*+ and *S*- are not treated by the animal as wholes, but rather in terms of the elements of which they are composed. Since *S*- contains two of the three separable elements in *S*+, any tendency to respond to *S*+ must generalize powerfully to *S*-.

The second reason for the feature-negative effect is that pigeons naturally peck at features that signal food (*sign tracking*); this is the basis for the phenomenon of autoshaping, briefly mentioned in the previous chapter. Sign tracking is a special case of hill climbing: going for the thing that best predicts reward. Thus, in the feature-positive case, the animals learn to peck at the distinctive element in the *S*+ display (the o in Figure 11.1) even before they show much decline in pecking *S*-. In the feature-negative case, however, the stimulus element that signals food is present in both the *S*+ and *S*- displays. Restriction of pecking to this feature is therefore incompatible with not-pecking *S*-. Pigeons can master a feature-negative discrimination if the display elements are very close together, so that *S*+ and *S*- are perceived as a whole, that is, they become integral rather than separable (Sainsbury, 1971). Pecking then is directed at the whole *S*+ complex, not to individual elements within it.

A third reason for the feature-negative effect may be the weakness of competing activities in *S*+ and *S*-. The discrimination procedure used in these experiments allows much time during the intertrial interval for activities other than key pecking: the ITI averages 60 sec, whereas *S*+ and *S*- are at most 7 s in duration. The birds soon learn not to peck during the ITI, so that the time is free for other activities. Hence (for the reason elaborated in Chapter 7 and later in this chapter: diminishing marginal value) the animal's tendency to engage in non-key-pecking activities during the relatively brief *S*+ and *S*- periods must be low. Thus, any contribution to discrimination performance made by the facilitation of antagonistic activities in *S*- will be small. In support of this idea are reports (see Note 18 and Figure 11.10) that the feature-negative effect is not obtained in more conventional successive discrimination procedures, where *S*+ and *S*- simply alternate at perhaps 60-sec intervals, with no intervening ITI.

The competition argument is diagrammed in Figure 11.2. The upper diagram shows the stimulus-control factors acting in the feature-positive case. *S*+ is made up of two kinds of element, E_1, the distinctive feature (o in Figure 11.1), and E_2, the element common to both *S*+ and *S*-. *S*- is made up of just E_2, the common element. E_1 controls T, the *terminal response*, of pecking; E_2 controls the *interim responses*, the collective term for activities other than the food-related terminal response (more on terminal and interim activities in a moment). The horizontal lines symbolize the reciprocal inhibition (competition) between T and I activities. Obviously, in the feature-positive case there is nothing to facilitate key pecking (T) in the presence of *S*-.

It is otherwise in the feature-negative case. Pecking is facilitated by E_2, the common element, so that suppression of pecking in *S*- must depend on reciprocal inhibition from interim activities. If these are weak, there is every reason to suppose that pecking will occur in *S*- as well as *S*+, as it does.

The general point of this argument is that discrimination performance is determined by two factors: (1) The nature of the stimuli: How discriminable are they? Are they integral or separable? (2) The availability of competing (interim) activities: How strong are they? What aspects of *S*- are likely to control them?

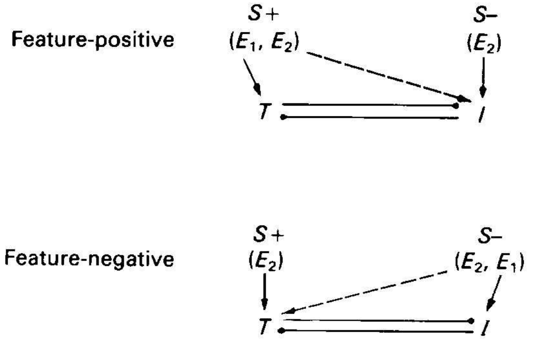

Figure 11.2. Excitatory control of terminal (T) and interim (I) activities by stimulus elements in feature-positive and feature-negative discriminations.

Several predictions follow from the idea of behavioral competition as an essential component of discrimination. First, discrimination performance should be better when a substantial competing activity is available than when it is not. Second, the operant response rate in one component of a successive discrimination (multiple schedule) should generally increase when reinforcement rate in the other component is reduced. This is the much-studied phenomenon of *behavioral contrast*. Third, inhibitory generalization gradients (shortly to be described) should generally be less steep than their excitatory counterparts. Fourth, discrimination gradients obtained following training where *S*+ and *S*- are on the same dimension should differ in predictable ways from gradients obtained after training with *S*+ and *S*- on separate dimensions.

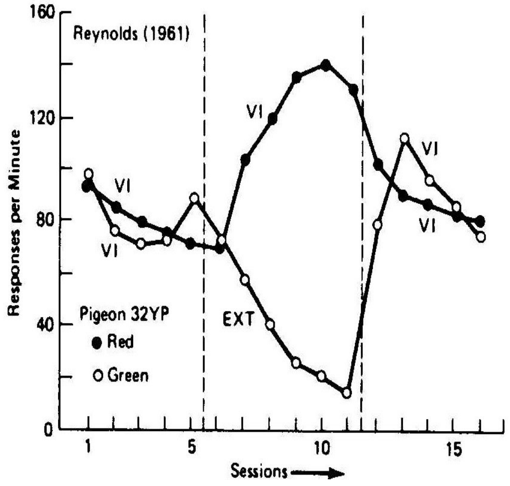

Figure 11.3. Positive behavioral contrast. Filled circles: response rate in the unchanged $(VI \rightarrow VI)$ component. Open circles: response rate in the changed $(VI \rightarrow EXT)$ component. Right third of the figure shows that the rate changes produced by the shift to extinction in the changed components are reversible. (From Reynolds, 1961b.)

Behavioral contrast and discrimination performance

When the operant-conditioning movement was enjoying its first flush of success, stimulus control was thought of in a simple way: A stimulus came to control an operant response when the response was reinforced in its presence. Each stimulus was thought of as an independent entity, maintaining behavior strictly according to the conditions of reinforcement associated with it. Available data supported this view. For example, if the schedule associated with one stimulus was variable-interval, and with an alternating stimulus, fixed-interval, then the behavior in each stimulus soon became appropriate to the schedule in force. During the fixed-interval stimulus (*component*) the animal would show "scallops" of accelerating responding between food deliveries; during the variable-interval component, a steady response rate would prevail.

It was therefore quite a surprise when George Reynolds in 1961[1] published a simple experiment that violated the rule of stimulus independence. His experiment had two conditions. In the first (the *pre-discrimination phase*), hungry pigeons were trained to peck a key for food reinforcement delivered according to a moderate (e.g., 60-sec) VI schedule. The stimulus on the key was either a red or a green light; the stimuli alternated every 60 sec. In the second condition (the *discrimination phase*), the schedule in one stimulus (green, say) was changed from VI to extinction (*multiple VI, EXl*). Conditions in the other stimulus remained unchanged. The result was a reliable and substantial increase in response rate in the unchanged (VI) component (*positive behavioral contrast*). This result is illustrated in Figure 11.3, which shows response rate in the two components before and after the shift from VI-VI (VI reinforcement in both components) to VI-EXT (VI reinforcement in one components, extinction in the other). Before the shift, response rate is roughly the same in both components; afterwards, as rate decreases in the extinction component, it increases in the (unchanged) VI component.

Reynolds also demonstrated the converse effect, negative contrast, which occurs when conditions of reinforcement in one component are improved, rather than degraded, e.g., when a multiple VI 60 VI 60 sec schedule is changed to multiple VI 60 VI 20 sec: the usual result is a decrease in response rate in the (unchanged) VI 60 component.

Behavioral contrast is a widespread but not universal effect. For example, it is much larger and easier to obtain with pigeons pecking keys than rats pressing levers — or pigeons pressing a treadle. If reinforcement rate in the VI-VI condition is very high (e.g., VI 15 sec), contrast effects are not obtained.

There are two obvious alternative explanations for contrast: the decrease in response rate in the component shifted to extinction, and the decrease in reinforcement rate in that component. The response-change account rests on the unstated hypothesis that the animal has only so many responses to “spend”, so that if he spends fewer in one component (because responses are no longer reinforced there), he will have more to spend in the still-reinforced component. The reinforcement-rate-change account rests on the general notion that response rate is guided by relative, rather than absolute, reinforcement rate.[2]

Reynolds and others attempted to discriminate between these two hypotheses by experimentally separating the response-rate and reinforcement-rate drops in the changed component. For example, in one experiment, pigeons were reinforced for *not* responding for 6 s — all periods of 6 s without a key peck ended with the delivery of food. This very effectively abolished pecking, but it failed to produce an increase in responding in the unchanged component — no contrast. In other experiments, food in the changed component was delivered independently of pecking or its availability was signaled; both maneuvers reduce or abolish pecking, but neither reliably produces contrast.

The general conclusion is that a change in reinforcement rate is usually sufficient to produce contrast, but a change in response rate unaccompanied by a change in reinforcement rate is generally ineffective.

Schedule-induced behavior

These various effects can all be brought together by the idea of response competition, provided one further thing is conceded: that time is *always* taken up by food-related activities, even when no explicit response is required to procure food. There is good evidence for this in general, although not much in some of the specific situations used to study behavioral contrast.

The clearest data come from the simplest situation. When hungry pigeons are daily exposed for an hour or so to a periodic-food (*fixed-time*, FT) schedule, they spend a good portion of their time near the feeder. If food delivery is frequent and they are sufficiently hungry, they will peck the wall in front of the feeder, even though the pecks are both unnecessary and ineffective. This activity is nicely synchronized with the delivery of food (although this aspect is not critical to the present argument).

Some typical results are shown in Figure 11.4. The figure shows the behavior of a single, well-trained pigeon averaged across three 30-min daily sessions of a fixed-time (FT) 12-sec schedule, i.e., one in which brief access to food was offered every 12 s. Each curve refers to a separate activity of the animal, assessed by visual observation. Activity R7 is "pecking the feeder wall", R8 a pacing movement from side to side in front of the wall, and so on. The curves show the probability of occurrence of each activity in each of the 12 seconds between food deliveries. Thus, R7 (pecking) never occurred during the first two seconds after food, but thereafter it occurred with increasing frequency, reaching a probability of almost one by the end of the interval.

All the animals showed the same pattern, namely, two classes of behavior within each interfood interval, a single activity that increases in frequency up to the time that food is delivered — this is called the *terminal response*; and other activities that occur during earlier parts of the interval — these are termed *interim activities* (*adjunctive behavior* is another term). Thus, R7 (pecking) is the terminal response; R8 (pacing) and two other movement patterns are interim activities in Figure 11.4. This division into terminal and interim activities has subsequently been demonstrated with rats and hamsters as well as pigeons (Staddon & Ayres, 1975; Anderson & Shettleworth, 1977; see also Staddon, 1977a). The particular activities that fall into the terminal and interim classes differ from species to species; all activities become less vigorous at longer interfood intervals; and the terminal response need not always be the same as the response normally elicited by food; but the general pattern is well established.

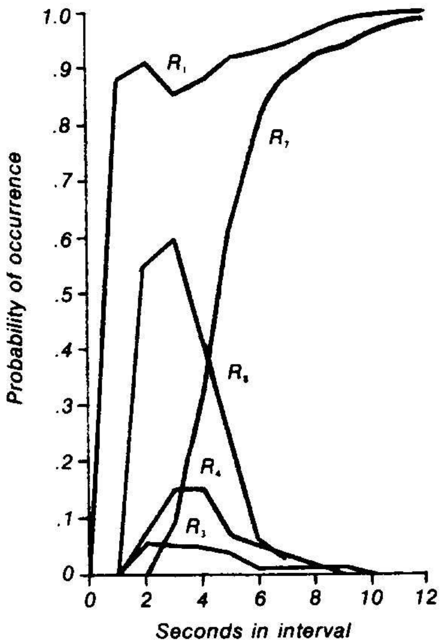

Figure 11.4. Average behavior of a single hungry pigeon receiving food every 12 sec. Each curve shows the probability that the numbered activity will occur in each of the 12 sec in between food deliveries. RI is "being in the vicinity of the feeder," R7 is "pecking the feeder wall"; other activities are described in the text. (From Staddon & Simmelhag, 1971.)

Terminal responses are obviously food-related, and occur at times when food delivery is likely. They are induced by processes usually studied as classical or Pavlovian conditioning. Interim responses are not generally food related, and occur at times when the food delivery is unlikely. This is the basis for earlier labeling activities controlled by *S*+ as "terminal" and those controlled by *S*- as "interim" in Figure 11.2.[3]

When food delivery depends upon a response (a fixed-interval rather than fixed-time schedule), this response follows the same time course as the terminal response in Figure 11.4. The *time taken up* by food-related activity is little affected by whether or not a response is necessary to get food — although, of course, the *type* of response shown is likely to depend upon the response contingency. When food delivery is aperiodic, as in a variable-time schedule, rather than periodic, the same division into terminal and interim activities can often be seen, especially if there are definable periods when the probability of food delivery is zero. For example, if the shortest interval between food deliveries is 5 sec (say), then interim activities will typically occur during the first two or three seconds after food delivery. When food deliveries are random in time, no definable post-food period is available for interim activities and they take on a vestigial form, occurring in alternation with food-related activities. Thus, on a VI schedule, key pecks often alternate with brief turning-away movements.

Whether food delivery is periodic or aperiodic, whether a response is required or not, animals spend much time in food-related activities. The rest of their time is spent in activities that tend to occur at times, or in the presence of stimuli, that signal the absence of food. And the competition between these two classes of activity seems to be especially intense in conditioning situations.

Intertemporal effects

The existence of competing interim as well as terminal responses, together with the property of diminishing marginal competitiveness (the higher the rate of an activity, the weaker its tendency to displace other activities), sets the stage for contrast effects. The argument is as follows: In the pre-discrimination condition, with the same VI schedule operative in both components, interim activities must occur in both. Because of competition for available time, key-peck rate must therefore be at an intermediate level. When one component is changed to extinction, the tendency to peck is weakened; interim activities are consequently free to increase their time allocation in that component. The more total time allocated to interim activities, the less competitive they become (because of diminishing marginal utility). Since the factors tending to produce pecking in the unchanged component have not altered, but its competitor has become less effective, pecking must increase in the unchanged component — which is positive contrast.[4]

A similar argument accounts for negative contrast. When the reinforcement rate in one component is increased, key pecking takes up more time. Hence, interim activities have less time and (because of diminishing marginal competitiveness) consequently become more competitive in the unchanged component, suppressing pecking there — which is negative contrast.

There are some obvious limitations on this argument. For example, when the VI-reinforcement rate is very high, very little time may be taken up by interim activities in the pre-discrimination phase. Consequently there is little room for further increases in the terminal response, should the interim activities become less competitive. Thus, contrast should be less at high reinforcement rates, as it is. There is a similar limitation when the rate of the terminal response is very low, rather than very high. In this case, the animal may be spending as much time as it needs in the interim activities, even in the pre-discrimination phase, so that the opportunity to spend even more time in the discrimination phase actually makes no difference. Thus, contrast effects might well be reduced, as they are, whenever the rate of the terminal response in the pre-discrimination phase is low.

The competition explanation for negative contrast depends upon displacement of competing activities from the changed to the unchanged component caused by an increase in the time allocated to terminal responding in the changed component. The VI response functions in Figure 7.15 show that such an increase is not to be expected over the whole range of reinforcement rates. If reinforcement rate in the pre-discrimination phase is already high, a further increase (in the changed component) might reduce, rather than increase, the proportion of time allocated to the terminal response in that component.

Procedural details also enter in ways that are rarely made explicit. For example, an increase in reinforcement rate in the changed component means that unless the component duration explicitly excludes eating time, the time available for responses other than eating is reduced, whether or not the level of terminal responding also increases. This factor effectively shortens the changed component and will always tend to promote negative contrast by displacing interim activities to the unchanged component.

Finally, contrast should depend upon the factors tending to strengthen or weaken interim activities. For example, rats are relatively inactive in a Skinner box when they are not actually pressing the lever; without environmental support (a running wheel, for example), rats show little evidence of interim activities. Pigeons, on the other hand, spend much time in the kinds of pacing and turning movements described earlier (Figure 11.4).[5]

These limitations account for the limitations on contrast described earlier: its absence at high reinforcement rates, with treadle pressing by pigeons (a very low-rate activity), and (usually) in rats. The fact that time is taken up by terminal responding, even if no instrumental response is required, accounts for failures to find positive contrast when food is signaled or presented independently of responding in the changed component: these manipulations abolish the measured response, but may not free up time for additional interim activities.

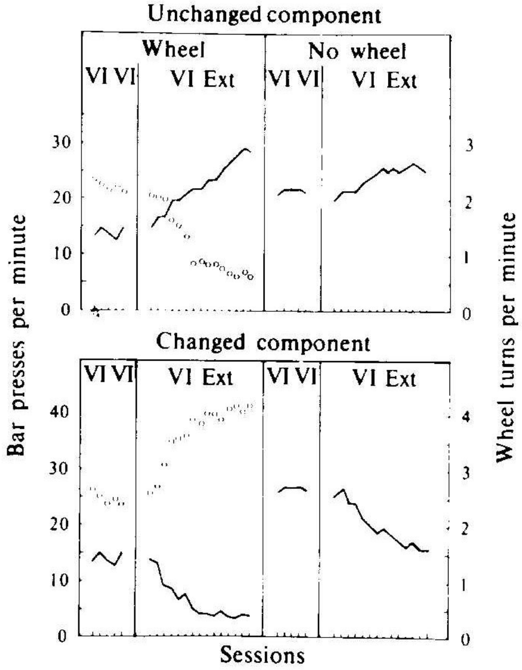

Figure 11.5. Reallocation of competing behavior during behavioral contrast. Mean daily rates of bar pressing (solid line) and wheel turning (open circles) for four rats in changed and unchanged components, with and without a running wheel available. (From Hinson & Staddon, 1978.)

Figure 11.5 shows the results of a simple experiment that demonstrates directly the role of competition in contrast and discrimination. Rats were run in the standard two-condition contrast paradigm: first trained with VI 60 s in both components of a multiple schedule, then shifted to VI EXT in one component. The experiment was done in two ways, either with or without a running wheel available to the rats. The top panels in Figure 11.5 show the levels of lever pressing and wheel running in the unchanged component. With no running wheel available (right quarter), lever-press rate is relatively high and the contrast effect (proportional increase in response rate in the unchanged component) small. With a wheel (left quarter), lever-press rate (solid line) is lower, because of competition from running, but the contrast effect following the shift to VI EXT is substantial. As response rate increases in the second panel, rate of wheel running (open circles) decreases. The lower panels show the concomitant changes in the changed component. With no wheel (right quarter), lever-press rate decreases little in the extinction component — discrimination performance is poor. With a wheel (left quarter), wheel running increases and lever pressing decreases substantially in the extinction component — discrimination performance is good. Thus, the availability of a strong interim activity can both improve discrimination performance and facilitate behavioral contrast.

The competition account of contrast is mechanistic not purposive. It refers only to the relative competitiveness of each activity, where the competitiveness of the terminal response in each component is assumed to be directly related to the reinforcement obtained, and the competitiveness of the interim activity a negatively accelerated function of its overall rate. In a moment, I will make these assumptions explicit with the aid of a simple model. It is worth noting, however, that the competition hypothesis does a pretty good job of *intertemporal utility maximization*. This is economese for a process that allocates resources among different temporal periods so as to maximize benefit to the individual. For example, the efficient executive will read reports or dictate memos on his commuter train, when he cannot engage in other species-typical activities such as telephoning, chairing committees or eating at the Four Seasons. In like fashion, the rat trapped in *S-* might as well "do" his interim activities so he can put in more time pressing the lever in *S+*. The mechanism of behavioral competition solves this problem in a way that requires no foresight and puts no load on memory.

COMPETITION AND MATCHING

Chapter 7 showed how a particular model for diminishing marginal utility could account for response functions obtained on numerous simple schedules. I now apply a simplified version of this approach to multiple schedules. It leads both to the qualitative predictions of behavioral contrast just described and to quantitative predictions of matching and undermatching on concurrent and multiple schedules.

The math in this section has three functions. The most important is to ensure that an explanation is clear: if an explanation cannot be framed in a formal way, then it is not clear — the idea of "competition" is not an obvious one; it is as well to be precise about we mean by it. Second, even if clear, verbal explanations do not lend themselves to the derivation of complex predictions. When there are many steps in the argument, or its outcome has a quantitative aspect (prediction of the shape of a gradient, for example), verbal argument is unsatisfactory. (I vividly remember being present some years ago at an argument between an eminent behavior theorist and a student about what the theorist's heavily published but non-quantitative theory predicted in a certain case. Someone suggested that they should settle the argument by implementing the theory as a computer program. The theorist bridled at this suggestion; but I have never understood why, since a theory whose predictions are a matter of opinion is worse than useless.) Third, and probably least important for biology and psychology, mathematical models make possible quantitative, rather than merely qualitative, predictions.

There is of course a cost in making formal models — even beyond its chilling effect on book sales[1]. At an early stage of knowledge, theorizing is necessarily imprecise. To insist on quantitative rigor may be just to exchange an exact irrelevance for a useful hunch. I believe that the experimental facts on discrimination performance are sufficiently solid and coherent to justify something more than verbal description.

So, how to apply the optimality analysis of Chapter 7 separately to each successive component of a multiple schedule, where the components may differ in the value of VI schedule associated with them? Let's begin with a concurrent schedule, which can be thought of as a multiple schedule in which the component durations are brief and, most importantly, under the control of the animal. The components share interim activities (for simplicity I will assume a single interim activity), since these can occur in both. I assume that in each component, the animal adjusts its level of responding so as to minimize cost (maximize value), in the sense defined in Chapter 7. Given the single further simplifying assumption that the cost of the instrumental response is negligible, minimum-distance-type models reduce to simple reinforcement maximization. The condition for optimal behavior is then that the marginal change in reinforcement rate be equal to the marginal change in value of the interim activity, formally

$$dR(x)/dx = dV(z)/dz, \qquad (11.1)$$

where $V(z)$ is the value of interim activity as a function of its level, z, and x is the level of responding in one component and $R(x)$ its rate of reinforcement. If we assume that the marginal value of interim activity is *constant* across both components of the schedule, the optimality condition is that

$$dR(x)/dx = dR(y)/dy, \qquad (11.2)$$

where y is the response rate in the other component.

The expressions in Equation 11.2 can be evaluated if we know the feedback functions, $R(x)$ and $R(y)$, for interval schedules, but it is not necessary to know these functions with exactness. In Chapter 7 I pointed out that many negatively accelerated functions (including two, the power and hyperbolic forms, that have been proposed as interval schedule functions) have the property that $dF(x)/dx = G(F(x)/x)$, that is, the marginal of the function can be expressed as a

---

[1] Theoretical physicist Steven Hawking reports that when he was writing *A Brief History of Time* his publisher warned him that each equation in the text would halve his sales.

function of the ratio of the function and its argument. Thus, for the hyperbolic function, $R(x) = ax/(a+x)$, $dR(x)/dx = a^2/(a+x)^2$, which can be rewritten as $[R(x)/x]^2$. Since $R(x)$ and $R(y)$ are interval schedule functions, we can rewrite Equation 11.2 as

$$G(R(x)/x) = G(R(y)/y),$$

which is obviously equivalent to

$$R(x)/x = R(y)/y, \qquad (11.3)$$

that is, matching of response and reinforcement ratios — a result discussed in Chapter 8 in connection with concurrent schedules.

To go further and derive contrast from this result, we need to incorporate two other things into the analysis: (a) a time-allocation constraint, and (b) a model for the diminishing-marginal-utility property of $V(z)$, the function relating the value of z (the interim activity) to its level.[6]

As in Chapter 7, I assume that activities x, $R(x)$, y and z are scaled in time units, so that the time-allocation constraint is just Equation 7.8 generalized to four activities, namely

$$x + R(x) + y + z = 1. \qquad (11.4)$$

The simplest way to incorporate diminishing marginal utility into $V(z)$ is to assume the same type of negatively accelerated relation between z and $V(z)$ as between x and $R(x)$ or y and $R(y)$. The argument for *some* kind of negatively accelerated function is that under free conditions, animals do more than one thing within the time period typical of an experimental session. As we saw in Chapter 8, nonexclusive choice implies diminishing marginal utility and a negatively accelerated value function. The defense for picking a function the same as the VI feedback function is simplicity — the conclusions I draw would not be much affected by choosing other negatively accelerated functions, however.

We are left, then, with two relations: Equation 11.4, and (if the cost-of-deviation of the instrumental response, X, is also assumed negligible) Equation 11.3 generalized to three activities,

$$R(x)/x = R(y)/y = V(z)/z. \qquad (11.5)$$

Consider now three cases of increasing complexity: (1) The single-response case (i.e., no response y); (2) the 2-response concurrent schedule; and (3) the 2-response multiple schedule.

Simple VI (single-response case).

In this case we have

$$R(x)/x = V(z)/z \qquad (11.6)$$

and $x + z + R(x) = 1$; hence $z = 1 - x - R(x)$. Substituting in Equation 11.6 yields

$$R(x)/x = V(z)/[1 - x - R(x)],$$

which reduces to

$$x = \{R(x) - [R(x)]^2\}/[V(z) + R(x)], \qquad (11.7)$$

which is the response function for VI given the simplification that the cost of the instrumental response is negligible.

If the time taken up by $R(x)$ is small, the squared term can be neglected and Equation 11.7 reduces to

$$x = R(x)/[V(z)+R(x)], \qquad (11.8)$$

which is a well-known response function for simple VI schedules proposed by Herrnstein.[7]

2. Concurrent VI-VI (two-response case)

We have already dealt with this case in the derivation from Equation 11.2; the result is the prediction of matching, which here requires only that the CoDs of the two instrumental activities be negligible, and assumes nothing about the feedback function for the interim activities beyond its constancy. It is also pretty obvious that matching can be derived directly from a generalization of the single-response case (Equation 11.8), namely

$$x = R(x)/\Sigma R \text{ and } y = R(y)/\Sigma R, \qquad (11.9)$$

where ΣR is the sum of reinforcement rates for all activities ($R(x)+V(z)$ in Equation 11.8, $R(x)+R(y)+V(z)$ for two instrumental responses), which reduces to

$$x/y = R(x)/R(y),$$

that is, matching of response and reinforcement ratios.

Contrast is derivable from this analysis just by comparing response rate, x_{pre}, in the unchanged component in the pre-discrimination condition (VI in both components) with rate, x_{post}, in the post-discrimination condition (VI in one component and EXT in the other). From Equation 11.9, $x_{\text{pre}} = R(x)/[R(x)+R(y)+V(z)]$, and $x_{\text{post}} = R(x)/[R(x)+V(z)]$, hence the *ratio* of pre and post (a measure of contrast) is just

$$x_{\text{post}}/x_{\text{pre}} = C_o = [R(x)+R(y)+V(z)]/[R(x)+V(z)]. \qquad (11.10)$$

Equation 11.10 has some of the properties described earlier for the informal competition model of contrast. For example, if $R(x)$ (and $R(y)$, since they are equal) is small relative to $V(z)$, contrast effects will be small, since the ratio C_o will be dominated by $V(z)$; thus, the analysis predicts small contrast effects when the absolute reinforcement rate for the instrumental responses is small (i.e., reinforcement is infrequent or the animal is only weakly motivated).

If $R(x)$ and $R(y)$ are large relative to $V(z)$, Equation 11.10 predicts large contrast effects, and as we have seen this is contrary to fact. But Equation 11.10 does not take into account the time constraints that affect multiple but not concurrent schedules, that is, the fact that responding in each component cannot take up more than the proportion of total time devoted to that component: in multiple schedules, component duration, hence the maximum possible disparity between numbers of responses in each component, is set by the experimenter. In concurrent schedules, it is set by the animal.

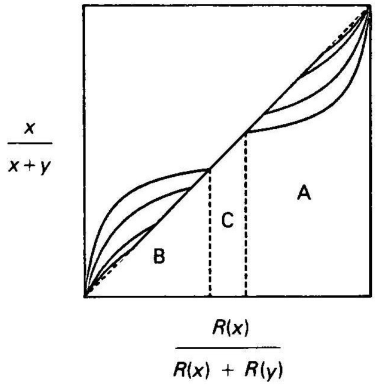

Figure 11.6. Matching relations predicted by competition model. Region C indicates range of $|R(x) - R(y)|$ for perfect matching for the outermost curve; as size of $V(z)$ relative to $R(x)$ and $R(y)$ increases, region of perfect matching also increases, as shown by the other curves.

3. Multiple VI-VI

Responses x and y in these equations are expressed as proportions of the total time. This raises no difficulties in either the single-response or concurrent cases, because responses x, y and z, can occur at any time. But on multiple schedules, the time available for x and y is limited to the durations of their respective components: in the equal-component case, to 50% of the total time each. This additional constraint introduces a discontinuity into the matching function, since neither x nor y can exceed .5, no matter what the reinforcement proportions. The result is shown in Figure 11.6, which is a plot of reinforcement proportions vs. response proportions. Perfect matching corresponds to the diagonal, but the constraint owing to "saturation" of responding in a component causes the function to deviate towards *undermatching* (i.e., too much responding in the minority component) when *R*(*x*) and *R*(*y*) are very different.[8] Each S-shaped function is for a different value of *V*(*z*): the larger *V*(*z*) relative to *R*(*x*) and *R*(*y*), the longer the linear part of this function — in accord with experimental results showing that matching on multiple schedules of food reinforcement *improves* when body weight is allowed to approach normal levels, that is, as the animals become *less* motivated.[9]

Figure 11.7 shows results from an early experiment on matching in multiple schedules.[10] The S-shaped form of the functions is clear, but they lack the discontinuities shown in Figure 11.6. There are several possible reasons for this: The data are averaged, which must blur sharp transition points if they show any variability. The hypothesized feedback relation, V(*z*), may not be identical to the VI feedback functions for *x* and *y*; the marginal of the VI feedback function may not be expressible as the ratio of the function and its argument; the assumption that instrumental responses X and Y are costless is only an approximation. These deviations from the simple mathematical model (should they turn out to be real) are not surprising. It is more important to attend to the fact that the general form of the deviations from matching in Figure 11.7 is just what we would expect from competition, and saturation.

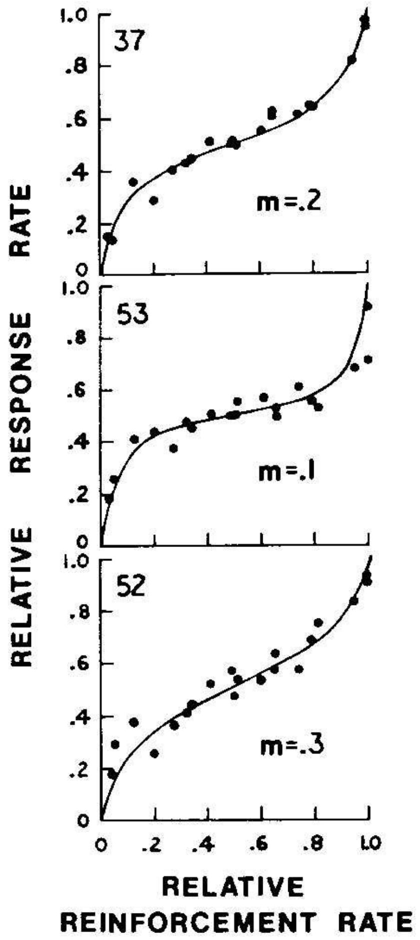

Figure 11.7. Relative response rate(*x*/[*x* + *y*]) vs. relative reinforcement rate (*R*(*x*)/[*R*(*x*) + *R*(*y*)]) from a multiple-schedule experiment by Reynolds (1963a) in which relative reinforcement rate was varied in three ways: One component held at VI 3 min, the other VI varied; the other component held to 1.58 min, the other varied; both varied. (From Herrnstein, 1970.)

Saturation also means that contrast is reduced at high as well as low values of *R*(*x*), as shown in Figure 11.8, which shows the contrast ratio, C_o, as a function of *R*(*x*) for a fixed value of *V*(*z*): contrast is a maximum when *R*(*x*) = *V*(*z*) (given the assumptions about *V*(*z*) just described), and declines asymptotically to 1 (no contrast) as *R*(*x*) increases beyond that value.

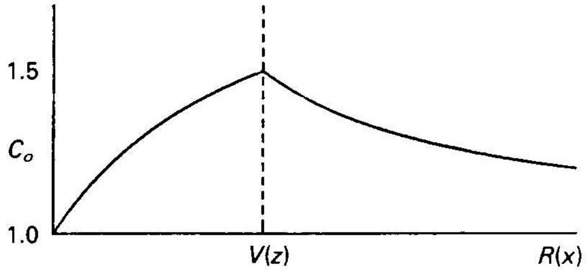

Figure 11.8. Contrast ratio, C_O (Equation 11.10), as a function of VI reinforcement rate, *R*(*x*), predicted from competition and diminishing marginal competitiveness. Contrast in multiple schedules is at a maximum at an intermediate value of reinforcement rate

This analysis of contrast and matching shows that behavioral competition provides a simple way for animals to allocate their time efficiently. It's also notable that one arrives at very similar conclusions by looking at behavior in competition terms, or in terms of marginal utility. In addition to accounting for behavioral contrast and its limitations, the approach also provides a rationale for Herrnstein's equation for the VI response function, for matching, and for the general form of deviations from matching on multiple schedules. The theoretical basis for these phenomena is far from settled, however, and there are several other accounts for positive contrast and some aspects of matching[11] — although none which attempts to relate all these effects to one another.

Inhibitory Generalization Gradients

Performance on reinforcement schedules is determined by the joint effects of external stimuli and competition. The message of the preceding section is that on variable-interval schedules the competitiveness of an activity is directly proportional to the reinforcement for it, and inversely proportional to its overall rate. Can we find a comparably simple way to describe the effects of external stimuli on behavior?

In the steady-state concurrent and multiple-schedule situations I have been discussing, discrimination is perfect or near perfect: the stimuli are easy to tell apart, and changes in relative response rate reflect not failures of discrimination, but efficient allocation of behavior. Even when a rat continues to respond in *S*- (as in Figure 11.5, bottom right, for example), there is little doubt that this is not because he cannot tell the difference between *S*+ and *S*-, but rather because the cost of lever pressing is very low and he has nothing more appealing to do. In a generalization test, response rate falls off gradually rather than abruptly; yet given appropriate training, extremely sharp gradients can be produced. Presumably the psychophysical discriminability of the stimulus continuum (i.e., the cognitive and perceptual properties discussed in Chapter 10) has not altered, so that this change reflects a change in behavioral competition. What form does this change take?

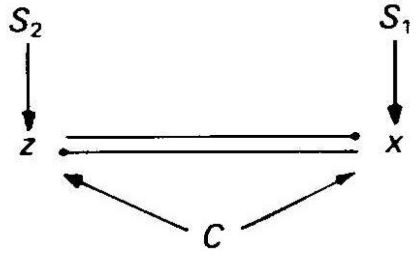

Figure 11.9. Causal factors in disjoint stimulus control: each activity is facilitated by a disjoint stimulus element (S_1 or S_2) and common factors (C), and inhibited by the other activity.

The static analysis of molar behavioral allocation developed in this chapter and in Chapter 7 provides a relatively simple way to deal with these interactions among stimuli, competing activities and differential reinforcement for those activities. Figure 11.9, a simplified version of Figure 11.2 (and similar in form to Figure 2.11 which was used to model reciprocal inhibition) provides a starting point. The figure shows the factors affecting terminal (x) and interim (z) activities in a successive discrimination of the sort already discussed. S_1 and S_2 represent *disjoint* (i.e., perfectly discriminable — another term is *orthogonal*) elements that differentiate *S*+ and *S*-: for example, if *S*+ is a dark pecking key with a white star in the center, and *S*- a green key, then we can assume both that there is zero perceptual overlap between these two stimulus elements (green and star) and, since they are the only aspect of the key that is different in the positive and negative schedule components, that they must be the basis for any discrimination. Thus x is controlled (*facilitated* might be a more accurate term[12]) just by S_1 and z is facilitated just by S_2; in addition, each activity is inhibited (because of competition for time) by the other. *C* in the figure refers to all those effective stimulus factors that have effects on both x and z (i.e., common factors). Let's consider how the causal diagram in Figure 11.9 applies to an experiment with a simple, successive discrimination procedure.

In this experiment[13] two groups of pigeons were exposed to two alternating stimuli, one associated with VI reinforcement, the other with extinction (a multiple VI EXT schedule, of the type by now familiar). For one group of animals, *S*+ was a vertical line, *S*- a blank key; for the other group, *S*- was a vertical line, *S*+ a blank key. Since the vertical line is the distinctive feature here, the first group is *feature-positive*, and the second *feature-negative*, in the sense these terms were used earlier.

After training sufficient to produce good discrimination between *S*+ and *S*- (i.e., few *S*-responses), both groups of animals were given generalization tests in which peck rate was measured in the presence of different line tilts. The results are shown in Figure 11.10 from a classic experiment. The result for the feature-positive group is the familiar excitatory gradient, with a peak at *S*+. The feature-negative group shows an *inhibitory* or *incremental* gradient, however, in which response rate increases as the test stimulus departs from *S*-; moreover (and this is typical of many similar subsequent studies) the inhibitory gradient is somewhat shallower than the excitatory one, even though the two experiments used the same procedure.

The diagram in Figure 11.9 implies that in the feature-positive group, the vertical line will facilitate *x*, the instrumental response, whereas *z*, the interim activities, will be facilitated by the blank key. Variation in line tilt during a test should therefore weaken *x* much more than *z*, yielding the usual decremental gradient for *x*. In the feature-negative group, however, the vertical line is *S*-, hence must facilitate *z* much more than *x*. Consequently, variation in line tilt must weaken *z* much more than *x*, allowing *x* to *increase* (since its facilitating factors — the blank key background to the line — are still present), producing the incremental, *inhibitory* gradient shown in Figure 11.10.

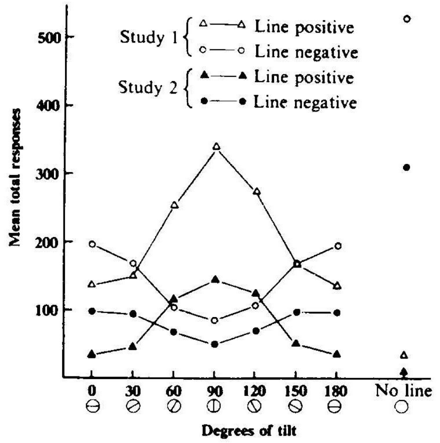

Figure 11.10. Inhibitory generalization gradients. *Open triangles*: excitatory line-tilt generalization gradient for a group of pigeons trained on multiple VI EXT, with a vertical line as *S*+, blank key as *S*- (feature-positive group). *Open circles*: inhibitor line-tilt generalization gradient from a group trained with a vertical line as *S*-, blank key as *S*+ (feature-negative group). *Closed symbols* are from a repeat experiment. (From Honig, Boneau, Burstein, & Pennypacker, 1963.)

The shallower slope of the inhibitory than the excitatory gradient is also easily explained. The excitatory gradient directly measures the effect of stimulus variation on the activity it facilitates, *x*, (pecking). However, the inhibitory gradient is an indirect measure of the effect of *S*- variation on *z* (an interim activity), the activity it facilitates. If we were to measure this interim activity directly (which, unfortunately, was not done in these experiments) — as in the contrast experiment shown in Figure 11.5 — we would expect to see an excitatory gradient as the *S*- feature is varied. The effect on the measured instrumental response, however, depends upon what proportion of the interim activities are directly facilitated (controlled by) the *S*- feature varied in the test. If the activity under the control of the *S*- feature comprises essentially all the interim activities, and time-allocation conservation holds (i.e., $x + z = 1$), then any decrease in *z* implies an equal increase in x. But if, as seems more likely, some interim activities are not controlled by *S*-, then any decrease in *z*, caused by variations in its controlling dimension, must be shared between *x* and this other activity. Consequently, a given decrease in *z* will generally produce a smaller increase in *x* (see the discussion of substitution relations in Chapter 7). Under any but the most restricted conditions the existence of a third class of activities, not controlled by either *S*+ or *S*-, is highly likely. Hence, we would expect inhibitory gradients to be generally shallower than excitatory ones, as they are.

Conjoint Stimulus Control and Peak Shift

The cases discussed so far, behavioral contrast and inhibitory generalization gradients, have been based on successive-discrimination training with disjoint stimuli. In all these experiments, *S*+ and *S*- are so easy to tell apart it is unlikely that the animals ever get confused as to which is which. When this is not true, when *S*+ and *S*- differ slightly on only a single dimension, the resulting generalization gradient is different in informative ways.

Figure 11.11 shows the result of a classic experiment on generalization after *conjoint* discrimination (i.e., discrimination with similar stimuli). Pigeons were trained on the familiar multiple VI EXT procedure; *S*+ was a monochromatic green light of 550 nm, *S*- was a greenish-yellow light of a slightly longer wavelength, 570 nm. In the control condition of the experiment, a group of pigeons was trained with *S*+ alone (simple VI), then given a generalization test. The result is the lower gradient, peaked at *S*+, in Figure 11.11. The other group was trained on the multiple VI EXT procedure, and then tested. The result is the higher gradient with its peak at about 540 nm, that is, shifted from *S*+ away from *S*-.

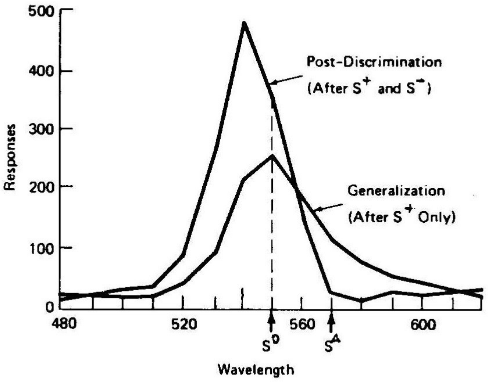

Figure 11.11. Wavelength generalization gradients after exposure to *S*+ only and after exposure to successive discrimination involving *S*+ and *S*-. The differences between the two gradients illustrate positive contrast and peak shift. (After Hanson, 1959.)

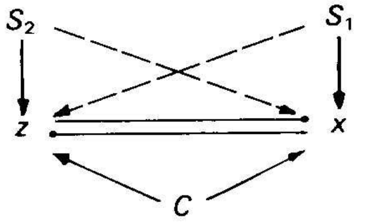

Figure 11.12. Causal factors in conjoint stimulus control. The stimulus element controlling the instrumental response, *x*, also has some facilitatory effect on interim activity *z*, and vice versa.

The increased responding in *S*+ in the second, postdiscrimination, gradient is just positive behavioral contrast, which we have already analyzed. The *peak shift* is a related phenomenon. It represents an effect of conjoint stimulus control. Figure 11.12 shows the causal model for conjoint control. The figure is the same as Figure 11.9, except for the added diagonal dotted lines, which represent the partial confusion caused by the similarity of *S*+ (S_1) and *S*- (S_2). Because the two stimuli cannot be discriminated perfectly by the animal, S_1 has some tendency to facilitate *z*, the interim activity, as well as a stronger tendency to facilitate *x*, the terminal response; and S_2, similarly, has some tendency to facilitate *x*, together with a stronger tendency to facilitate *z*.

To show that this partial confusion between S_1 and S_2 can cause peak shift, it is necessary to give the pictures in Figures 11.9 and 11.12 some quantitative properties. A simple way to do this is to let the effects depicted in Figure 11.12 be linear and additive.[14] This means writing two equations, one to describe the factors acting to facilitate activity *x*, another for those acting to facilitate *z*, thus

$$x = S_1 - k_{zx}z + Cx, \qquad (11.11A)$$

and

$$z = S_2 - k_{xz}x + C_z, \qquad (11.11B)$$

where *S* represents the additive facilitating effect of the two controlling stimuli, *k* represents the subtractive effect of *z* on *x* and vice versa and *C* is the additive effects of stimulus factors not varied during a generalization test.

Notice that if by convention we take *x* to be the measured response, then Equation 11.11A represents the relation between *x* and *z* in the feature-negative case (S_2 varied) and Equation 11.11B represents the relation in the feature-positive case (S_1 varied).

Because these subtractive effects are due to competition for available time, the gain in all other activities associated with unit decrement in activity *z*, must be equal to one — there cannot be a total gain. But when three or more activities are involved, k_{zi} can even be negative for some activities — if *y* strongly inhibits *x*, and *z* strongly inhibits *y*, for example, then an increase in *z* might well lead to an *increase* in *x*, because of reduced inhibition of *x* by *y*; this is termed *disin-*

hibition.

To illustrate the kinds of interaction I'm thinking of, imagine that there are three activities involved, x, y and z, measured in time proportions so that $x + y + z = 1$. Suppose that the free levels of these three activities are 0.2, 0.3 and 0.5, say, but that for experimental purposes, z is progressively restricted to values of 0.4, 0.3, 0.2 and 0.1. The linear inhibition assumption predicts that the levels of x and y will increase in such a way that plots of x against y or z, or of y against z will all be linear. Corresponding to the given values of z, the values of x might be 0.25, 0.30, 0.35 and 0.40 and of y, 0.35, 0.40, 0.45 and 0.50 (for convenience, these are all listed in Table 11.1). In this instance, the decrements in z are shared equally between x and y, but this need not always be the case.

Table 11.1.

	Free	z Constrained			
x	.2	.25	.30	.35	.40
y	.3	.35	.40	.45	.50
z	.5	.40	.30	.20	.10
Total time	1.0	1.0	1.0	1.0	1.0

In this example, $x = 0.45 - 0.5z$, and $y = 0.55 - .5z$ so that $k_{zx} = k_{zy} = 0.5$, and $S_1 = 0.45$. To find k_{xz} and S_2, it is necessary to hold x to various levels and measure the level of z.

In the disjoint case, where variation in a stimulus element or dimension affects only one activity directly, Equations 11.11A and B imply that there should be a linear relation, with slope generally between -1 and 0, between the activity whose controlling stimulus is being varied and any other activity. Since we can expect the controlled activity to show an excitatory gradient, this means that measurement of an activity not directly controlled by the varied dimension should show an inhibitory gradient of generally shallower slope, as I argued in the preceding section.

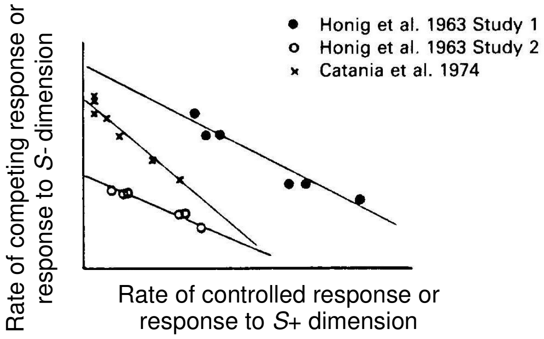

Figure 11.13. Linear relations between levels of interim and terminal responses in successive discrimination. Response rates in inhibitory and excitatory gradients from the studies shown are plotted against each other. (Data from Honig et al., 1963; Catania, Silverman, & Stubbs, 1974.)

There are few direct tests of the linear prediction. However, Figure 11.13 shows three sets of data that provide one direct and two indirect tests (see also Figure 7.1). The figure shows plots of one response, controlled by the stimulus dimension being varied (abscissa), versus a competing response whose controlling stimulus remains constant (ordinate). The experiment most directly related to the present analysis is by Catania, Silverman and Stubbs (1974).[15] They trained hungry pigeons to peck at two keys for food delivered on independent VI schedules. The stimulus on the right key was a vertical line, on the left key a color. Line tilt was varied in the generalization test, and the result was an excitatory generalization gradient for pecking on the right (line-tilt) key and an inhibitory gradient for pecking on the left (color) key. The *x's* in Figure 11.13 represent average response rate on the left plotted against average response rate on the right at each line-tilt value: the points are tolerably well fitted by a straight line, as predicted by Equation 11.11. The other two lines are from two similar experiments by Honig et al., already shown in Figure 11.11. These data are a less direct test, just because they represent a between-groups, rather than within-group or within-animal, comparison. Nevertheless, when response rate for the feature-positive group (abscissa) is plotted against rate for the feature-negative group (ordinate) at the same line tilt, the points again are approximately collinear. The slope for the Catania et al. data is close to -1, suggesting that pecking on the two keys occupied most of the available time; the slope for both the Honig et al. experiments is between -1 and 0.

So far we have established that both terminal responses and interim activities have controlling stimuli; that variation in these stimuli yields excitatory generalization gradients; and that because of competition, an excitatory gradient in terms of one response will usually be associated with a (shallower) inhibitory gradient in terms of other responses. To show in addition that these processes can lead to peak shift we need to make some assumption about the way in which variation in a controlling stimulus dimension affects the stimulus contribution in Equation 11.11.

The terms S_1, S_2 in Equation 11.11 represent the excitatory contributions of controlling stimuli to the level of each activity. If the properties of the stimulus are varied from the *S*+ value, then this contribution will decrease — presumably in proportion to the *similarity change* associated with a given change in the physical properties of the stimulus. If we make some definite assumption about the relation between physical change and similarity change, it becomes possible to predict generalization effects. I first give the equation for the disjoint case (inter-dimensional discrimination); then derive from them the necessary equation for the conjoint case (intra-dimensional discrimination), from which peak shift can easily be derived.

The form of the competition equations in the *disjoint* case (Figure 11.9) is

$$x = S_1.\mathbf{S}(u_0,u) - k_{zx}z + C_x, \tag{11.12A}$$

and

$$z = S_2.\mathbf{S}(v_0,v) - k_{xz}x + C_z, \tag{11.12B}$$

where u and v are physical dimensions controlling x and z independently, u_0 is the value associated with *S*+ and v_0 is the value associated with *S*-, and $\mathbf{S}$ is a function representing the similarity between u and u_0 and v and v_0. The properties of $\mathbf{S}$ are straightforward. For example, when $u = u_0$, $S = 1$ (i.e., $S = 1$ means identity, the highest value of similarity), so that Equation 11.12A becomes equal to Equation 11.11A; when $u>u_0$ or $u<u_0$, $S<1$. In addition, we might expect that the changes in S as u deviates increasingly from u_0 will be gradual rather than abrupt and that as the deviations become very large similarity will approach zero. Many functions satisfying these rather minimal conditions are sufficient to predict the properties of inhibitory gradients and peak shift. For example, a simple one is

$$\mathbf{S}(u_0,u) = 1/[1+(u_0-u)^2],$$

and another is the familiar bell-shaped Gaussian curve,

$$\mathbf{S}(u_0,u) = \exp[-(u-u_0)^2/D], \tag{11.13}$$

with mean u_0 and standard deviation D, which is the one I shall use.

Conjoint (in this case *intra-dimensional*) control incorporates both feature-negative and feature-positive cases, since variation in the same stimulus dimension affects both x and z directly. Hence Equation 11.12B must be rewritten as

$$z = S_2.\mathbf{S}(u_1,u) - k_{xz}x + C_z, \tag{11.14}$$

where u_1 is the value of *S*- on the varied dimension. Eliminating z from Equations 11.12A and 11.14 and rearranging yields response rate x as a function of the stimulus value, u:

$$x(u) = [S_1.S(u_0,u) - k_{xz}S_2.S(u_1,u) - k_{zx}C_z]/(1-k_{zx}k_{xz}), \tag{11.15}$$

which is just the weighted difference of the two similarity functions: $S(u_0,u)$, which is centered on the *S*+ value, and $S(u_1,u)$, which is centered on the *S*- value, i.e.,

$$x(u) = A_1S(u_0,u) - A_2S(u_1,u) + A_3, \tag{11.16}$$

where A_1, A_2 and A_3 are lumped constants made up of the various constant terms in Equation 11.15. Substitution of Equation 11.13 for the similarity functions in Equation 11.16 then allows prediction of the postdiscrimination generalization gradient.[16]

Figure 11.14 shows the kind of prediction that results: the two identical bell-shaped curves are the Gaussian similarity functions for *S*- and *S*+; the higher curve is the postdiscrimination gradient derived from Equation 11.16. There are three things to note about the predicted postdiscrimination gradient: (a) It is steeper than the underlying similarity functions; this underlines the conclusion we came to earlier, that behavioral competition aids discrimination performance. The hypothetical animal in Figure 11.14 discriminates perfectly between *S*+ and *S*- (in the sense that no responses are made to *S*-) even though the similarity functions for the two overlap considerably. (b) The postdiscrimination gradient is higher than the similarity functions; this is just behavioral contrast. (c) The peak of the postdiscrimination gradient is shifted away from *S*+ in a direction opposite to *S*-; this is *positive peak shift*.

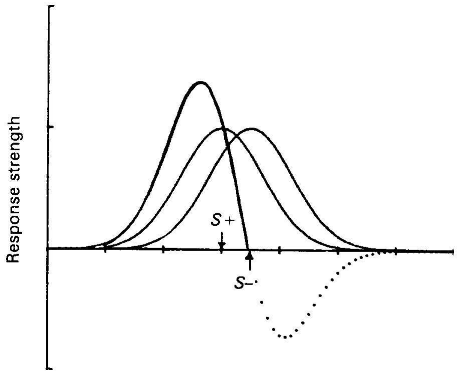

Figure 11.14. Positive and negative peak shift predicted from linear reciprocal inhibition and Gaussian similarity gradients. Two Gaussian curves are the similarity gradients; heavy curve is the predicted postdiscrimination gradient, showing positive peak shift; dotted portion shows negative peak shift.

The term C_x (the contribution to x made by factors common to *S*+ and *S*-) in Equation 11.14 acts like an additive constant; when C_x is high (early in training, for example, before the discrimination is fully developed), the horizontal axis in Figure 11.14 may be displaced downwards. Response x then occurs during *S*- as well as *S*+, so that a *negative peak shift* (Guttman, 1965) — the dotted curve — can sometimes be observed.

Thus, the basic properties of discrimination performance, inhibitory control and generalization gradient peak shift seem all to reflect behavioral competition and direct excitatory control of behavior by stimuli.

Several other effects also fit into this general picture. For example, some years ago Terrace did an experiment in which pigeons were first trained to peck a distinctive *S*+ (e.g., a key illuminated with white vertical line) for food presented on a VI schedule. Once this pecking was established, Terrace occasionally presented a brief *S*- with a very different appearance (e.g., a dark key). Pecks on this stimulus did not lead to food, but in fact most animals never pecked it. Terrace then progressively increased the duration of *S*- and faded in a color. In this way he was able to establish discriminations without the animals making any "errors" (*S*- responses). The animals turned away from the key as soon as *S*+ went off, so that presumably no activity came under the explicit control of *S*-. Perhaps for this reason, these animals failed to show inhibitory generalization gradients (responding was close to zero at all values of the varied *S*- dimensions). They also failed to show behavioral contrast. Other work has confirmed that contrast and inhibitory gradients seem usually to go together. The present analysis suggests they should, because both depend upon stimulus control by *S*- of activities antagonistic to the terminal response. Contrast and peak shift also tend to go together, although there are some dissociations: contrast can occur without a peak shift, and vice versa, for example. These dissociations are not surprising, because although peak shift and contrast both depend on the strength of inhibitory interactions (the k values in Equation 11.11) as well as the similarity difference between *S*+ and *S*-, the quantitative form of the dependence is quite different.[17]

The static view of free-operant discrimination performance just presented is undoubtedly much too simple: it assumes linear competition, additive stimulus effects, and invariant similarity relations — and all of these things are probably true approximately at best. Nevertheless, this view brings together a range of experimental facts not easily related in any other way. It is likely to require modification in at least three respects: to accommodate nonlinear effects; to incorporate acquisition processes — *how* discrimination is learned as well as *what* is learned; and to accommodate dynamic effects. I have little to say about nonlinear models, and acquisition is dealt with later. I end this chapter with a brief account of dynamic effects.

Dynamic Effects

Most of this chapter has been concerned with the effects of a very simple successive-discrimination procedure in which two stimuli, *S*- and *S*+, are presented in alternation. I have said little about the significance of the duration of each component, although in fact it makes a great deal of difference to the outcome. For example, suppose we compare the results of two contrast experiments, Reynolds' original, in which 60-s components are presented in alternation, and another in which *S*+ is presented for 60 *minutes* each day for several days, followed by *S*- for a comparable time, followed by a return to *S*+. What differences might we expect between these two experiments?

The two studies will be similar in one respect: animals will in both learn to respond when *S*+ is present and not when *S*- is present. But the similarities end there. In the second study, response rate in *S*+ on its second presentation is unlikely to be higher than during its first — no behavioral contrast. If anything, rate in *S*+ the second time around is likely to be *lower* than at first (this is sometimes termed *induction*[18]). In a generalization test, there will be no peak shift nor will it be possible to demonstrate inhibitory generalization gradients around *S*-.

There is an obvious lack in the static scheme that accounts for these differences. Recall that the basic explanation for contrast is that *z*, the competing, interim activity, is reallocated to the extinction component (or the component with lowered reinforcement rate if reinforcement rate is merely reduced in *S*-); with more total time spent doing *z, z* becomes less competitive in *S*+ allowing the rate of the terminal response to rise (positive contrast). The fine print here is the word *total*: over *what period* is this total to be measured? In explaining the feature-negative effect, I assumed that the length of the 1-min ITI relative to the 7-sec S- was sufficient to reduce the competitiveness of interim activities in *S*-. We saw in Chapter 6 that animals seek to regulate the rate of activities not just over days, but even over periods of a few minutes. Yet the static model assumes that the competitiveness of *z* (which is proportional to *V*(*z*) in the model) is constant and depends upon its average across a whole experimental session, or at least across one *S*+ — *S*- cycle. This is obviously wrong when component duration becomes longer than a few seconds. For example, if component duration is 30 minutes, say, who can doubt that the competitiveness of interim activities at the end of a "rich" component, when there has been little time to engage in anything but the terminal response, is likely to be considerably higher than at the end of the alternating "lean" component, when there has been much free time. Thus the assumption of constant *V*(*z*) will increasingly be violated at long component durations. What effects might this have?

To answer this question we need to know something about the second-by-second processes that determine the strength[19] of an activity. Our knowledge here is unsatisfactory because it is practically impossible to study activities in isolation. Anything we can measure is likely to be multiply determined (cf. Note **17**), making identification of individual causal factors a formidable task. Nevertheless, we can make some reasonable guesses based on the analysis of choice in Chapter 8. Recall that nonexclusive choice (at the level of individual responses) depended on negative feedback. In patch foraging, for example, the longer the animal works one patch, the smaller the incremental payoff. The result is that he switches from patch to patch, rather than remaining fixed on one. Similarly, on concurrent VI-VI schedules, the longer the animal continues to respond to one alternative, the better the payoff for responding to the other; hence he spends some time on both. Similar negative feedbacks must underlie the animal's allocation of time to different activities under free conditions. The only difference is that the feedback must then be internal, rather than via external contingencies of reinforcement. The effects in both cases are the same: an activity becomes less attractive the longer it is engaged in. I will term this the *principle of satiation* by analogy (but only analogy, no other similarity is intended) to satiation for food and water.[20]

There is of course a converse principle, *deprivation*: the longer the time since the last occurrence of an activity, the stronger its tendency to occur. Both principles are combined in the generalization that *an activity becomes relatively less attractive as more time is spent on it*.

These principles do not apply to every activity at all times. Under restricted conditions, even opposite effects can be demonstrated, such as *sensitization*, when the occurrence of an activity for a while *increases* its strength. Conversely, some activities may become less likely if they are prevented from occurring. Still others apparently occur at random, with no dependence on time at all. Nevertheless, satiation-deprivation has the same kind of generality as the principle of diminishing marginal utility: without it, there could be no preferred mix of activities and no regulatory response to contingencies that force a deviation from the preferred mix.

In multiple schedules, there are three activities to consider: the terminal responses in each component, x and y, and the interim activities, z. Since animals on these schedules are generally very hungry, we might expect relatively little weakening within a session of the tendency to make the food-related response (i.e., x or y). The interim activities are presumably less urgent than eating, however, so should be more subject to satiation. If the rates of deprivation and satiation are relatively rapid, we might then expect to see changes in the competitiveness (strength) of z from component to component of the sort depicted in Figure 11.15. The center panel shows the changes in the strength of z, on the assumption that it occurs only at a low rate in $S+$ (the rich component) so that deprivation causes an increase in strength; z can presumably occur at a higher rate in $S-$ (the lean component), so that satiation then causes a decrease in strength. In $S+$ the strength of the terminal response is high, in $S-$ it is low, as shown in the top panel. Since $z + x$ or y is constant, because of time allocation, the actual values of x and y will be the complements of z, as shown in the bottom panel. We have seen this general pattern before; it is the *exponential lead* discussed in Chapter 3, and represents a direct relation between the controlled variable (here response rate, x or y) and the *rate of change* of the controlling variable (here rate of reinforcement).[21]

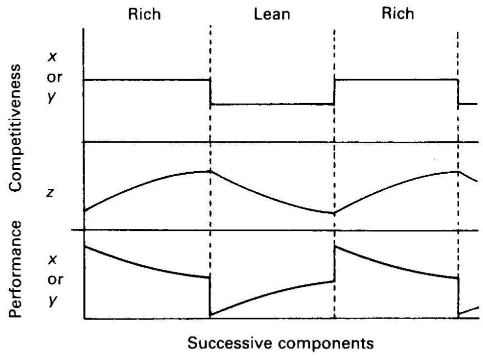

Figure 11.15. Model for dynamic changes in interim activities. *Top panel* shows step changes in competitiveness of the terminal response, associated with stimuli in each multiple-schedule component. *Middle panel* shows satiation-deprivation effects on interim activity — becoming less competitive in lean component, more competitive in rich component. *Bottom panel* shows local contrast effects predictable from the difference between top and middle panels.

Figure 11.16 shows experimental data confirming that *local contrast* effects of this sort actually occur. The figure shows response rate for two individual pigeons during successive 30-sec periods within 3-min components of a multiple VI 6 VI 2-min schedule. As in the theoretical figure, response rate is highest at the beginning of $S+$ components and lowest at the beginning of $S-$ components, subsequently converging towards an intermediate level. Other experiments have shown that the cause of these effects is indeed the relative richness of the two schedule components: there is negative local contrast in a lean component preceded by a rich one, and positive local contrast in a rich preceded by a lean.

Relative response and reinforcement rates in multiple schedules conform better to matching when component duration is short — a few seconds — than when it is long — a few minutes (Shimp & Wheatley, 1971; Todorov, 1973). Satiation-deprivation dynamics suggest why: the longer the components, the less valid the assumption (necessary to the derivation of matching) that $V(z)$ is constant throughout. $V(z)$ will tend to increase during a long, rich component, and to decrease during a long lean one. Hence, z will be more competitive during the rich component than during the lean, so that (terminal) response rate will be lower during the rich component than required by matching; the result will be undermatching, which is the usual deviation from matching on multiple schedules.

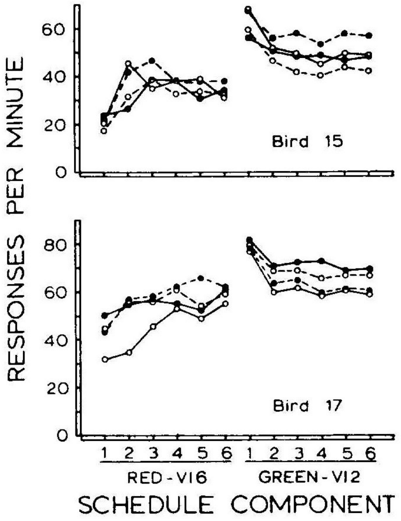

Figure 11.16. Local contrast effects in a multiple VI 6 VI 2-min schedule; each curve is an average of a single experimental session. (From Nevin & Shettleworth, 1966.)

An intriguing fact about these local contrast effects is that they are usually *transient*, occurring during early exposure to the multiple VI-VI schedule, but dissipating after a few days. On multiple FI-FI schedules, however, they persist indefinitely; and they can be reinstated on multiple VI-VI by manipulations that make the stimuli harder to discriminate or remember.[22] Under other conditions they may fail to occur at all, or occur in a different way — being correlated with the succeeding rather than the preceding component, for example.[23]

The transience of local contrast effects has not yet been satisfactorily explained, but one possibility is as follows. The magnitude of local contrast obviously depends upon the relative importance as causal factors of competition and facilitating stimuli: if the effect of stimuli is large ($S_1 >> k_{zx}z$ in Equation 11.11A, for example), then changes in the strength of competing activities will obviously have little effect; conversely, if the effect of stimuli is relatively weak, changes in competition will have large effects — of course of the effects of stimuli are zero, there can be no effects at all, since changes in response rates will not be correlated with stimulus changes. It turns out that the things that promote local contrast are also things that should weaken the animal's ability to identify *S*- and *S*+ or weaken control by *S*+ and *S*-: making *S*+ and *S*- more similar; introducing a competing source of control (multiple FI-FI, where responding is under both temporal control, by the reinforcer, and control by the discriminative stimuli, compared to multiple VI-VI, where only the stimuli are predictive); and intermixing many different stimuli, so that memory limitations make it difficult to recognize individual stimuli. Thus the transience of these effects is consistent with some kind of additive model in which response rate is jointly determined by stimuli and competition from other activities.

There is no obvious explanation for local effects in which rate changes seem to reflect the following, rather than the preceding, component. The procedures in which these effects are found are often complex — with more than two components, for example — but this is not a sufficient explanation in itself. The effects may be *path dependent*, that is, they could depend not just on the conditions of reinforcement now, but also on prior history: the sequence of conditions leading up to the present arrangement. Metastable (see Chapter 5) patterns are relatively common on multiple schedules, and frequent switching between conditions is necessary if one is to be sure of the stability of any given pattern.[24]

Given the kind of complex individual dynamics I have described, it is far from obvious what we should expect when there are several schedule components and their durations are close to the natural frequency of the hypothesized satiation-deprivation processes. It is not too difficult to provide a formal model for the simple case in which only the single competing activity z is subject to satiation-deprivation effects. If there is more than one such activity, and the time constants for each are different, or if the instrumental responses are subject to satiation-deprivation, it is not trivial to provide an exact account. But without an exact account it is impossible to sure what form local contrast effects should take. Hence the jury is out on the significance of effects other than standard local-contrast effects.

Stimulus Effects

Local contrast effects are interactions along the dimension of time, true dynamic effects. Very similar interactions take place along stimulus dimensions in *maintained generalization gradients*. A maintained gradient is obtained when instead of two very different stimuli being used as *S*+ and *S*-, many similar stimuli are used, one or a few being associated with reinforcement, the rest being unreinforced. For example, Figure 11.17 shows elegant data from an experiment by Blough (1975; see also Catania & Gill, 1964; Malone & Staddon, 1973) in which pigeons were reinforced for pecking at wavelength stimuli from 597 to 617 nm and not for pecking shorter wavelengths, from 570 to 595.5 nm. After many days training on this procedure, the average percentage of responses to each wavelength followed the smooth curve shown in the figure. Of special interest are the positive and negative "shoulders" on either side of the boundary between *S*- and *S*+ stimuli. The shape of the curve bears an obvious resemblance to the perceptual "edge effects" discussed in Chapter 3. Like those effects, these also hint at the existence of inhibitory interactions along the stimulus dimension.

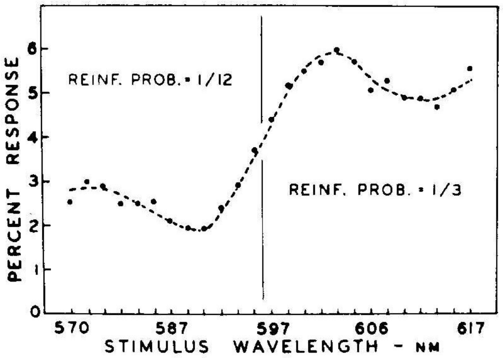

Figure 11.17. Edge effects in a maintained generalization gradient. Each point is the average response probability to the indicated wavelength; stimuli on the right signaled reinforcement, those on the left, its absence. (From Blough, 1975.)

Blough explains these and other similar data by means of a linear learning model of the type briefly discussed in Chapter 8 (see also Chapter 13) together with a generalization assumption analogous to the similarity gradients discussed in the previous section. Like most local contrast effects, these stimulus effects are transient and disappear after more or less extended training. Blough's dynamic model handles both the general form of the inhibitory "shoulders" and their eventual disappearance, but has not been applied to behavioral contrast, peak shift, matching and so on, nor to the effects of stimulus- and memory-related manipulations on the transience of local contrast.

SUMMARY

In this chapter I argued that the static, molar properties of performance on multiple (successive discrimination) and concurrent (simultaneous discrimination) reinforcement schedules can be derived from four assumptions. (a) That inhibitory stimulus control — the suppression of an ongoing activity by presenting a stimulus — is associated with excitatory control of antagonistic activities: inhibitory control of activity A is excitatory control of complementary activity ~A. The two complementary classes correspond to the *terminal* and *interim* activities reliably observed in periodic-food experiments. (b) The competitiveness of an activity decreases as its level increases (diminishing marginal competitiveness). (c) That the competitiveness of an activity is inversely proportional to its rate of occurrence and directly proportional to its rate of reinforcement; (d) and that activities are in equilibrium when all are equally competitive. Matching, on concurrent VI-VI schedules, deviations from matching, on multiple schedules, and behavioral contrast, together with the effects on contrast of species, absolute reinforcement rate and response type, are all derivable from these assumptions.

Two other assumptions allow prediction of generalization-gradient peak shift and the relations between contrast and peak shift: (a) In the steady state, response rate is additively determined by stimulus factors and competition from other activities. (b) Stimulus factors may either be independent (disjoint) so that the stimulus element facilitating one activity has no effect on its antagonist, or overlapping (conjoint), so that the same stimulus element has effects (of different magnitudes) on both activities.

These molar, static relations must depend in some way on local, dynamic processes. One obvious possibility is that as an activity continues to occur at a high rate, its competitiveness declines, and as time goes by without the activity, its competitiveness increases. A moment-by-moment satiation-deprivation process of this sort is the dynamic counterpart of static diminishing marginal competitiveness. Satiation-deprivation processes can account for local (i.e., time-dependent) contrast effects.

Nevertheless, several questions remain unanswered: What is the proper quantitative form for these dynamic processes? Will the same form do for all activities? For the same activity under all conditions? What accounts for the transience of local contrast? Memory limitations seem to be involved because local contrast revives under conditions where stimulus identification or control is weakened. How do memory mechanisms relate to these short-term dynamic processes?

Stimuli are additive terms in these equations, but of course no discriminative stimulus acts directly on behavior. To be effective, a stimulus must be recognized, and recognition depends on the properties of memory in ways that are the topic of a later chapter.

NOTES

1. Reynolds (1961b). An effect similar to behavioral contrast had been shown some time earlier by Crespi (1942), with rats running in runways. But the *Crespi effect* (as it is sometimes called) depends on a single change in amount of reward in the same situation, rather than different frequencies of reward in the presence of different, alternating, stimuli. Behavioral contrast is in fact hard to show when only amount of reward is varied, and it is a steady-state effect, rather than an effect of a one-time manipulation. These differences in species and apparatus, and the different theoretical tradition within which Crespi's work was done, meant that it did little to diminish the impact of Reynolds' experiment.

Research on behavioral contrast since Reynolds' first experiment is reviewed in Schwartz and Gamzu (1977), and a there is a good brief review in the textbook by Fantino and Logan (1979).

2. The response-change and reinforcement-change accounts of contrast are conceptually rather different. The response-change account is perfectly mechanistic, and just depends upon conservation of key pecks: what is lost in one component is gained by the other. But the reinforcement-change account demands a bit more in the way of computation by the animal, since it assumes control of behavior by *relative* rather than absolute reinforcement rate.

The competition view (to be presented in a moment) is similar to the response-change account in one respect, that it emphasizes activities rather than reinforcement. But it differs in three important ways: it does not require that pecks lost in *S*- exactly equal pecks gained in *S*+ — which is fortunate, since the two are not generally equal; it considers activities other than the measured instrumental response; and it acknowledges that animals spend time in food-related activities even when food is delivered free.

3. I have argued that there are at least two different types of interim activities: *facultative* and *schedule-induced* (Staddon, 1977a). Facultative activities are things like wheel running (for a rat) or pacing (for a pigeon): they occur on food schedules but are not facilitated by them. In-

deed, the usual effect of a food schedule is to reduce the proportion of time devoted to facultative activities. Schedule-induced interim activities, on the other hand, are facilitated by a schedule. The most striking example is drinking by rats; other examples are schedule-induced attack, by pigeons, and perhaps wood-chewing by rats. Most interim activities appear to be of the facultative variety (cf. Roper, 1981). The arguments I will make apply to facultative rather than schedule-induced interim activities.

4. In support of the competition view see Estes (1950), Henton and Iversen (1978), and Hinson and Staddon (1978). Bouzas & Baum (1976) and White (1978) have shown that contrast involves time reallocation. The main arguments in this section are taken from a theoretical chapter by Staddon (1982).

5. It is true that "lying around" is as much an activity as "pacing" and similar vigorous behaviors, but it is also likely to be much more easily displaced by other activities (i.e., have a lower cost-of-deviation, cf. Chapter 7) — which is the important point for the competition argument.

6. *Matching as a Consequence of Maximizing*. Matching, and Herrnstein's equation for response functions on a single VI, can be derived from minimum-distance-type models in the following way: Recall that the objective function from the minimum-distance model (Equation 7.9) is

$$C[x,R(x),z] = a(x_0-x)^2 + b(R_0-R(x))^2 + c(z_0-z)^2, \qquad \text{(N11.1)}$$

where x is the rate of instrumental responding, z is the level of interim activity and $R(x)$ is the obtained rate of reinforcement. The constraints are due to time allocation:

$$x + R(x) + z = 1, \qquad \text{(N11.2)}$$

and the VI feedback function:

$$R(x) = Ax/(A+x), \qquad \text{(N11.3)}$$

where A is the programmed (i.e., maximum) VI reinforcement rate. Equations N11.1, N11.2 and N11.3 are then combined to form the Lagrangian:

$$L\ [x,R(x),z,\lambda_1,\lambda_2] = a(x_0-x)^2 + b(R_0-R(x))^2 + c(z_0-z)^2 - \lambda_1(1-x-R(x)-z) - \lambda_2(R(x) - \\ -Ax/(A+x)). \qquad \text{(N11.4)}$$

If we assume that the cost of x is negligible, then $a = 0$ and the first term in Equation N11.4 vanishes. Taking partial derivatives to find the minimum then yields:

$$\partial L/\partial x = \lambda_1 + \lambda_2 A^2/(A+x)^2,$$

which can be rewritten as

$$\partial L/\partial x = \lambda_1 + \lambda_2[R(x)/x]^{2,} \qquad \text{(N11.5)}$$

plus four other expressions. A similar exercise for the other instrumental response, y, yields an expression exactly parallel to Equation N11.5. By the hypothesis, these two marginals must be equal to the fixed marginal for interim response z; hence they can be equated, which yields $R(x)/x = R(y)/y$, the matching relation. Adding the time constraint yields Herrnstein's hyperbolic equation, as shown in the text.

7. Herrnstein (1970). Herrnstein's equation is (in his symbols) $P = kR/(R+R_0)$, where P is response rate, R is reinforcement rate, R_0 reinforcement rate for "other" behavior and k a constant. If response rate is proportional to time spent, this equation is equivalent to Equation 11.8 in the text, with k being the number of responses per unit time. When eating time is excluded from the rate measures, experimental results over a range of VI values that includes all but very short interfood intervals fit this relation well.

8. *Matching and Contrast: Formal Derivation and a Mechanical Model.* Figure 11.18 shows a very simple mechanical analogue for the equations in the text. The figure shows an airtight cylinder containing two free pistons that divide the unit volume into three compartments of volume x, y and z respectively. Clearly, $x + y + z = 1$, as in Equation 11.4 (the time taken up by $R(x)$ can be neglected). In equilibrium, the pressure of gas in each of the three compartments must be the same, or else the pistons would move. Hence we can assume that $P(x) = P(y) = P(z)$. From Boyle's law we know that for each compartment, $P(i).i = NT$, where i is the volume, N is the number of molecules per unit volume and T is the absolute temperature. Suppose we identify N with the reinforcement rate for activity i, and let T be an arbitrary constant, k. We can then write $P(i).i = R(i).N$, *or* $P(i) = kR(i)/i$ for all three activities. Equating the values of P leads at once to Equation 11.5 in the text: $R(x)/x = R(y)/y = V(z)/z$.

If the pistons in Figure 11.18 are completely free to move anywhere in the cylinder, the model is appropriate for concurrent VI-VI schedules, since the component duration is there completely under the control of the animal. Multiple schedules require the addition of a *detent* in the middle of the cylinder that, for equal-duration components, prevents x or y from taking up more than half the total volume. It is obvious from this analogy that perfect matching can hold in the multiple case only if $P(z)$ is sufficiently large that neither x nor y expands to more than half the total volume.

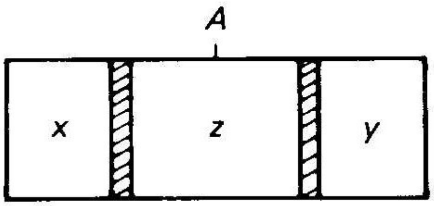

Figure 11.18. Pressure model for matching and contrast effects in multiple and concurrent schedules. Figure represents a cylinder containing two free pistons dividing unit total volume into three compartments, of volumes x, y, and z. For equal-component multiple schedules there is a detent at A preventing x or y from exceeding half the total volume.

This argument can be made formally to provide a simple account for behavioral contrast. Behavioral contrast is defined by the ratio of response rates in component X of a multiple schedule under two conditions: a *pre-discrimination* phase, in which responses X and Y receive equal VI reinforcement; and a *post-discrimination* phase in which reinforcement for X continues as before, but reinforcement for Y is abolished. The limitation on the maximum values of x and y relative to z, the level of the interim activity, means that the analysis must consider separately the two cases where $V(z)$ is less than or equal to, or greater than, $R(x)$:

1. $V(z) \leq R(x)$: In the pre-discrimination phase, $R(x) = R(y)$; hence it is obvious that z takes up equal volumes in compartments x and y. If $V(z) > 0$, neither x nor y can "saturate," so that the simple matching law applies. Thus,

$$x_{pre} = R(x)/[2R(x)+V(z)], \qquad \text{(N11.6)}$$

in the pre-discrimination phase. When reinforcement for Y is abolished in the postdiscrimination phase, however,

$$x_{\text{post}} = R(x)/[R(x)+V(z)], \qquad \text{(N11.7)}$$

but since $V(z) \leq R(x)$, this expression is ≥ 0.5, the maximum possible value for x. Consequently, $x_{\text{post}} = 0.5$ (i.e., x has saturated) and the magnitude of contrast is therefore given by

$$x_{post}/x_{pre} = C_o = [R(x)+.5V(z)]/R(x)$$

$$= 1 + V(z)/2R(x),\ \ V(z) \leq R(x), \qquad \text{(N11.8)}$$

which is a positive function in $V(z)$, and negative in $R(x)$. Contrast is at a maximum, at $C_o = 1.5$, when $R(x) = V(z)$.

2. $V(z)>R(x)$: As before, $x_{\text{pre}} = R(x)/[2R(x)+V(z)]$. However, the matching relation yields a value for x_{post} that is less than the maximum possible, because $R(x)<V(z)$; consequently, x_{post} is

given by Equation N11.7 and contrast is therefore given by

$$C_o = [2R(x)+V(z)]/[R(x)+V(z)], \quad V(z)>R(x), \tag{N11.9}$$

which is a negative, hyperbolic function of $V(z)$.

Equations N11.8 and N11.9 show contrast to be a nonmonotonic function of both $V(z)$, the reinforcement for ("strength" of) other behavior and $R(x)$, reinforcement rate in the unchanged component. Both functions are discontinuous and bitonic, with maxima at $R(x) = V(z)$.

The predictions for matching in multiple schedules must be derived in a similar piecemeal way, except that there are now three cases to consider:

(a) If $R(x) \geq R(y)+V(z)$, then $x = 0.5$ (X saturates) and

$$y = 0.5R(y)/[R(y)+V(z)], \tag{N11.10}$$

i.e., matching between Y and Z in a compartment of duration 0.5.

(b) If $R(y) \geq R(x)+V(z)$, then $y = 0.5$ (Y saturates) and

$$x = 0.5R(x)/[R(x)+V(z)]. \tag{N11.11}$$

(c) If neither (a) nor (b) is true, then $V(z)>|R(x)-R(y)|$; neither X nor Y saturates, so that overall matching applies: $x/y = R(x)/R(y)$.

If a constraint such as $R(x)+R(y) = K$ is placed on $R(x)$ and $R(y)$, then the function relating $x/(x+y)$ *to* $R(x)/[R(x)+R(y)]$ can be derived. Examples are shown in Figure 11.6. Each function has three regions corresponding to cases (a), (b) and (c). Simple matching holds in the central region defined by $|R(x)-R(y)|<V(z)$ (i.e., $V(z)$ greater than the absolute value of the difference between $R(x)$ and $R(y)$). As the value of $V(z)$ increases relative to $|R(x)-R(y)|$, region (C) increases relative to regions (A) and (B), the curvature of the function in regions (A) and (B) decreases, and the function as a whole approximates more and more closely to matching. I expand on this argument in the next chapter.

9. Herrnstein and Loveland (1974); see Reynolds (1963b) for data showing that contrast effects are small when reinforcement rates are high.

10. These data are from Reynolds (1963a); the fitted curves are from an equation proposed by Herrnstein (1970), based on his analysis of matching (see Note 8, above). In my symbols, his equation is

$$x = R(x)/[R(x)+KR(y)+V(z)],$$

where K is a constant. It is obviously similar to Equation 11.10 in the text, and so is forced to predict that contrast must always be reduced by increases in $V(z)$ — which is false. Parameter K indicates the degree of interaction between components, being high for short components, low for long. The equation must therefore predict that absolute response rate in both components should decrease as component duration decreases, which is also inconsistent with data (Edmon, 1978). Herrnstein's multiple-schedule equation is evidently an inappropriate extension of his successful rule for simple VI response functions.

11. *Theories of Contrast*. There are two other theoretical accounts of contrast that have achieved wide currency: additivity theory and the matching law. I will show in a moment that those features of the matching-law account that fit the data also follow formally from the competition view. Additivity theory (see Schwartz and Gamzu, 1977, for the most extensive account) argues that contrast experiments are also autoshaping situations, in the sense that reduction of reinforcement rate in the changed component improves the signal properties of the stimulus in the unchanged component. Since pigeons will often peck at stimuli that signal food (the theory continues), these "autopecks" add to the pecks already maintained by the response contingency signaled by $S+$ to produce contrast. Ingenious experiments have shown that if the signal for $S+$ is

presented on a key separate from the pecking key, additional (and useless) pecks are indeed delivered to it when the changed component is shifted to extinction -- and little or no contrast is then observed on the pecking key. Nevertheless, the theory has run into serious trouble because it cannot easily handle negative contrast or contrast in situations that do not produce autoshaping, such as rats pressing levers or pigeons pressing treadles. Attempts to show that autopecks are quantitatively different from "operant" pecks (shorter duration, less vigorous) are also open to the alternative interpretation that these two classes just differ in strength (see Note 19 to this chapter and the discussion of reflex strength in Chapter 2). (See also Ziriax & Silberberg, 1978; Schwartz and Gamzu, 1977; Farley, 1980; Whipple & Fantino, 1980).

A possible interpretation of autoshaping in terms of the competition view is that pecking and similar food-elicited activities will occur in situations that signal food where competing tendencies are also at a minimum. For example, the account of dynamic factors later in the chapter suggests that the most effective way to weaken competing activities is to lengthen *S*- relative to *S*+; but this is also the best way to get pigeons to peck at *S*+ (Gibbon, Locurto and Terrace, 1975).

Rachlin (1973) has proposed a non-quantitative theory of contrast similar to the dynamic model sketched out at the end of this chapter. In recent years, Nevin and his associated have elaborated a theory of behavioral "momentum" that can accommodate a wide range of contrast and choice phenomena (e.g., Grace, 1994; Nevin, 1974; Nevin, Tota, Torquato, & Shull, 1990)

12. Use of the term *controlled* to refer to the effect of a stimulus on a response is conventional, but it is not quite accurate. Behavioral competition means that the level of an activity is jointly determined by its stimulus and by competing activities (hence by their stimuli); discriminative stimuli exert less than total control over the level of their responses. Terms such as *facilitated* or *excited* are more accurate in the present theoretical context.

13. Honig, Boneau, Burstein, & Pennypacker (1963). The argument in the next few paragraphs about time allocation and the slope of inhibitory gradients was first made by Jenkins (1965). Note that in this situation, with relatively long *S*+ and *S*- components and no timeout periods in between, pigeons have no difficulty handling a feature-negative discrimination.

14. In earlier discussion, response rates *x* and *y* were involved in two other sets of equations: time-allocation conservation (Equation 11.4) and reinforcement (Equation 11.7). It is natural to ask about the relation between these two and Equations 11.11A and B, the competition equations. This is a gray area, but a possible answer is as follows: The first two equations (certainly the reinforcement equation) refer to conditions at equilibrium. They say nothing about the moment-by-moment determinants of activity. But the competition equations refer to the results of generalization tests, which are brief probes that give information on the current "forces" acting on behavior. There is no reason, therefore, why the two sets of equations should be simply related.

Animals seem to adapt to these steady-state procedures not by learning specific stimulus-response "connections" but by constructing a *routine*, analogous in some ways to a computer program, that has some stimulus "inputs" to be sure, but also involves other processes (see Staddon, 1981b). If the world changes, the effect on the animal depends on how large and how persistent the change is. If it is relatively brief, and not too large — as in a generalization test — then the routine built up by the training procedure continues to function, and the results of the test can tell us something about it. But if the change is large, other processes come into play, the old routine is partially or completely abandoned and the animal goes about constructing a new routine to cope with the changed circumstances. The temporary nature of the structure probed by the generalization-test procedure is emphasized by the changes that take place during a test (the gradient becomes steeper, for example) and the transience of phenomena like the peak shift: with repeated testing, the shift disappears and the peak of the gradient moves back to $S+$.

The competition equations are an attempt to make a model of the routine set up by successive-discrimination procedures. The parameters of these equations must be such that the balance of behavior satisfies the earlier time-conservation and reinforcement equations, but the way the system responds to stimulus changes is not itself predictable from those equations.

15. The other study in Figure 11.13 is by Honig et al. (1963).

16. *Spence's Theory of Transposition*. Equation 11.16 has a family resemblance to the most famous model offered for an effect like peak shift, that due to Kenneth Spence (1937). In my symbols, Spence's model amounts to

$$x(u) \propto S(u_0,u) - S(u_1,u),$$

i.e., responding is proportion to the simple difference between excitatory and inhibitory gradients. There are procedural and formal differences between the two models. Spence was concerned with data from simultaneous (rather than successive) discriminations, in which rats were trained to respond to the larger (say) of two squares. After training, the animals were then confronted with two new squares, the smaller of which was equal in size to the larger of the previous pair. Would the animals respond to the same-size square (absolute responding, congenial to stimulus-response theory) or to the still-larger square (*transposition*: relational responding congenial to gestalt theory)?

The clever animals usually picked the larger square, giving aid and comfort to the gestaltists and, until Spence, spreading gloom and despondency among S-R theorists. Spence's insight was that this kind of result can be derived from S-R assumptions: if reinforcing responses to $S+$ causes the development of an excitatory gradient (something widely accepted at the time), then *non*-reinforcement of responses to $S-$ should produce an inhibitory one. If these two gradients are smooth, and the slope of the inhibitory one is greater than the slope of the excitatory one (at least in the vicinity of $S+$), then the difference between the two gradients will have its peak not at $S+$, but away from $S+$ on the side opposite to $S-$. Given $S+$ and a still larger stimulus, therefore, Spence's model predicts that response strength associated with the larger stimulus might well be greater than response strength associated with $S+$. Moreover (and this was the *coup de grace*), Spence's model also predicted that if the new stimulus were too large, the animal would not transpose but would show *transposition reversal*, preferring the original $S+$. Animals indeed do this, so that Spence's model received solid support.

More than twenty years later, after the invention of the technique of free-operant generalization testing and the discovery of peak shift following intra-dimensional discrimination, Spence's theory was applied to the peak shift. But here it has several flaws: The least important is that it was devised for simultaneous situations, but peak shift does not occur in concurrent operant experiments. More critical is its requirement that the inhibitory gradient be steeper than the excitatory one, and that responding to $S+$ after the formation of a discrimination be at a lower level than before. As we have seen, measured inhibitory gradients are generally shallower than excitatory ones, and postdiscrimination $S+$ responding is generally at a higher level than before (positive behavioral contrast).

S-R theorists took Spence's success as further proof that animals do not respond to relations, but just to simple physical properties. Later work shows that they won the battle but lost the war. Experiments in which animals can learn to recognize complex stimulus categories, such as "people" or "trees" (see Note 17, Chapter 10), studies in which they learn to pick the "odd" stimulus, or one of intermediate value in a set of three, show that animals do indeed possess the complex perceptual abilities favored by the gestaltists. Spence's contribution was to show that these talents are probably not involved in some tasks that might seem to demand them.

Spence was also correct in rejecting "relational responding" as a satisfactory explanation for anything. It's fine to show that pigeons can solve tricky perceptual problems; but to term this "relational responding" and leave things there is just to label the mystery, not explain it. A proper explanation must, like Spence's, eventually get down to the physical properties of the stimulus and the process by which they are transformed into measurable behavior by the animal. We are very far from being able to do this, even for things like object recognition. Animals are able to perform feats of recognition that still cannot be duplicated by even the most sophisticated pattern-recognition programs.

17. See Staddon (1977b) for a quantitative account of the different predictions for contrast versus peak shift. For reviews of experimental results see Terrace (1966), Mackintosh (1974) and Rilling (1977).

18. This *induction* is quite different from the *induction* of Pavlov and Sherrington discussed in Chapter 2: Pavlovian induction is much closer to local contrast (shortly to be discussed) than to Skinner's induction. Why Skinner, who was certainly aware of the prior usage, chose to use the term in this very different way is a puzzle.

19. *Response Strength*. It seems impossible to talk about choice and the allocation of individual behavior without using terms like "tendency," "causal factors" and "strength" that refer to the likelihood one activity will win out over others. The term *response strength* has a long history in experimental psychology, with its meaning changing according to the currents of theoretical fashion (cf. discussion of reflex strength in Chapter 2). Applied to instrumental behavior the term has been consistently tied to the supposed strengthening effects of reinforcement. When the influence of operationism was at its height, earnest attempts were made to define the term operationally. These met with only limited success, and some alarm was caused by experiments showing that supposedly equivalent measures of strength such as *latency, vigor, probability* and *resistance to extinction* did not reliably rise and fall together: a short latency need not always go along with high vigor, probability or resistance to extinction (see Osgood, 1953, for a review of this controversy).

Skinner (e.g., 1950) added his nail to the response-strength coffin by pointing out that response properties like latency and vigor can be shaped by suitable contingencies of reinforcement. A measure of reinforcement that can itself be altered by reinforcement is about as much use as a rubber ruler, argued Skinner. He went on to advocate *response probability* (by which he meant *rate* of response) as only true measure of response strength. Unfortunately, response rate is just as amenable to shaping by means of reinforcement as response vigor, so that the special status of response probability could be preserved only by bringing in the idea of time as a discriminative stimulus. Low-rate behavior might then be deemed strong at particular points in time. For example, on spaced-responding schedules, probability of a response is low just after a response, and rises to high values at times approaching the spaced-responding value. Time clearly can act like more conventional stimuli, but response rate under these conditions ceases to be a useful quantity. The idea was widely held, but little discussed.

Nevertheless, a notion like response strength is almost indispensable. One approach to the "strength" problem has been via empirical laws relating response rate to relative rate of reinforcement — the equations of the matching law have sometimes been proposed as an appropriate measure — the strength of a response being just its relative rate of reinforcement. J. A. Nevin (e.g., 1974) has taken a different approach, suggesting that response strength be defined as *resistance to change*. The idea is that a strong response is one whose rate is changed only slightly by operations usually effective in weakening or abolishing operant behavior, such as extinction and the presentation of free food. Nevin's idea is obviously close to the concept of competitiveness, nor is it remote from the matching relations (see also Staddon, 1978, for a theoretical discussion of this view, which is developed further in Chapter 12).

These developments follow a current and I believe correct philosophical trend, namely to let the definition of a term like *strength* grow out of valid theory. Operationism put the cart before the horse. P. W. Bridgman (its physicist founder) noticed that physical terms like force and energy could be reduced to sets of measuring operations. He erroneously drew the conclusion that requiring operational definitions for fledgling terms in young sciences would be helpful to their growth — forgetting that even in physics, the theory came first, operations after. Infants grow taller as they mature, but a stretched infant is not an adult.

20. This is not a new principle. Hull's (1943) *reactive inhibition* has similar properties; sensory adaptation and reflex habituation (see Chapter 2) are also similar in some respects. More recently the theoretical consequences of self-inhibition have been most extensively explored by Atkinson and Birch (1970). See also numerous papers by Grossberg (e.g., 1981, 1982) for related mathematical discussions of inhibition and competition.

21. One form of the satiation-deprivation assumption applied to activity z is

$$dV(z)/dz = K_1 - K_2z, \qquad \text{(N11.12)}$$

where K_1 and K_2 are constants: when z (the rate of the interim activity) equals zero, the rate of growth of $V(z)$, the "value" of z, is maximal; as z increases, the rate of growth of $V(z)$ decreases, eventually becoming negative when $z > K_1/K_2$. The solution to Equation N11.12 is just the exponential lag function derived in Chapter 3; the complement of this function (i.e., the time-course of the terminal response) is the exponential lead.

22. Staddon (1969); the effects may also be persistent on multiple VI VI when the components are very long. For example, Rachlin (1973) has reported large and apparently persistent local effects with 8-min components. Conversely, local effects may be weak when components are short (Hamilton & Silberberg, 1978). The dynamic analysis suggests that local effects should be stronger at longer component durations. Effects of number of components have been shown by Malone & Staddon (1973), of stimulus discriminability by Catania & Gill (1964).

23. See Buck, Rothstein and Williams (1975) for some of these unusual effects.

24. Under some conditions, activities that happen to occur in a highly predictive stimulus may persist long after the changes that caused them have passed. For example, in an experiment by Kello, Innis and Staddon (1975) pigeons were trained on a multiple FI 1 FI 3-min schedule, in which 12 FI 1 intervals occurred in one component (green key) and 4 FI 3 intervals occurred in the other (red key). After some training, all animals showed typical local-contrast effects: elevated response rate in the first FI-1 interval, depressed responding in the first FI-3 interval. Then the green key in the FI 1 component was changed to blue for just one interfood interval. All the animals at once responded at a high rate — this is a relatively familiar finding on FI: temporal control (see Chapter 13) by the food is context-dependent, so that a novel stimulus abolishes it, leading to a high response rate. The novelty of this result was that the effect persisted indefinitely, so long as the blue stimulus occurred only during one FI-1 interval out of 12 in each component. This result shows that under certain not-very-well-specified conditions, unusual effects can be maintained for long periods on multiple schedules.