

THE OPTIMAL ALLOCATION OF BEHAVIOR

UTILITY AND ADAPTATION TO CONSTRAINT

The last chapter looked at feeding as a typical motivational system and examined the effects of variables such as palatability, diet dilution, and brain damage on the amount eaten and how hard a rat will work for food. These effects were quite well summarized by a simple, static feedback model whose central assumption is that animals regulate¹ eating rate. Chapter 6 did not deal directly with interactions among multiple motivational systems, nor did it deal with the regulation of activities like running or lever pressing that do not fit into the traditional motivational trinity — hunger, thirst, and sex. It turns out that both these omissions — motivational interactions and “weakly” motivated activities — can be handled in the same way; that is, by an extension of the regulatory approach.

The last chapter argued that animals seek to regulate their rate of eating. This is hardly controversial: Food is essential to life and failure to regulate food intake, at least over the long term, does not favor a large posterity. But experimental results suggest that animals do more than this: They regulate even over periods of time too short to pose metabolic dangers, perhaps because of the need to anticipate the possibility of deprivation in the future. Moreover, eating is not the only thing regulated in this way: Essentially everything the animal does tends to occur at a certain rate, and animals will take action to restore this equilibrium rate if it is perturbed in some way. Before I can explain how this works, I need first to explain some basic economic concepts, such as utility (value) and marginal utility, and the idea of a constraint.

I begin by describing a typical experimental situation used to study multiple activities, then describe a fundamental constraint to which all behavior is subject: limitation of time. I go on to recount the pioneering work of David Premack, who first experimented with “weak” reinforcers and proposed a general definition of reinforcement based on this work. Optimality (economic) models of operant behavior derive, in large part, from his proposal. The remainder of the chapter describes the optimality approach and shows how a wide range of experimental results can be accounted for by two principles: that the marginal utility of any activity decreases with the amount available (the *principle of diminishing marginal utility*), and that animals act to maximize utility, subject to constraints of time and reinforcement schedule.

The constraints of time and behavioral competition

Imagine a rat restrained in a semi-natural environment that allows it to run in a wheel, eat, drink from a water bottle, and engage in “comfort” activities and “exploration.” At a given season and time of day, the proportion of time that the animal spends in these various activities is likely to be quite stable from day to day. A number of experiments² have revealed two important properties of this mix of activities: (a) competition among activities for the available time; and (b) regulatory processes that tend to maintain the activities in fixed proportions.

Competition is most easily and directly demonstrated simply by limiting access to one activity and measuring the change in the amount of time spent on others. For independent activities, such as running and drinking, the effect of restricting one is generally to increase the other. For reasons not yet well understood, competition appears to be especially acute under the conditions typical of operant-conditioning experiments, that is, intermittent reinforcement given to a highly motivated animal. The effects of preventing an activity are then especially clear and nearly always facilitatory: Hungry rats pressing a lever for food will press faster if denied the

opportunity to run, for example.³

In one sense, competition for time is a necessity. Amount of time is always fixed: If activities are measured in time-based units, and if everything is considered to be an activity — the animal is never allowed to “do nothing” — then the total must remain constant. If activities are defined in this way, an increase in one forces a decrease in at least one other. Few situations have been studied carefully enough to meet these conditions: Not everything the animal does is recorded, activities are most often measured as number of occurrences per unit time (rather than in terms of total time taken up), and response topography is not measured in a way that allows one to go back and forth with confidence between rate and time measures.⁴ The usual procedure

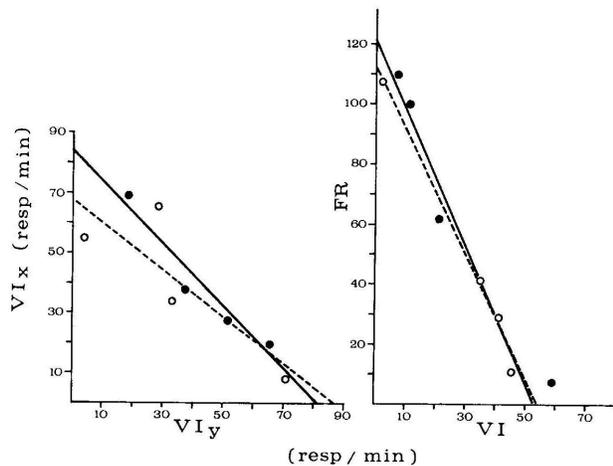


Figure 7.1. *Left panel* shows response rate on one key plotted against response rate on the other key for several pairs of variable-interval schedules in a two-choice experiment. Open and closed circles indicate data from two different pigeons (Herrnstein, 1961). *Right panel* shows similar data from a single pigeon in a two-choice variable-interval, fixed ratio experiment. Open and closed circles show data from conditions with or without a change-over delay (Bacotti, 1977).

is just to measure two or three “standard” activities (in addition to eating) and look for relations between them.

If less than all activities are recorded, there is no assurance that decrease of one activity will be accompanied by increase of any other. Nevertheless, in experiments with very hungry animals pressing a lever or pecking a key for food, it is safe to assume that the animal does little else but press or peck, and eat. Consequently, such experiments provide good evidence for behavioral competition enforced by the time constraint. For example, Figure 7.1 shows data from a choice experiment by Herrnstein (left panel) in which hungry pigeons could obtain food by pecking at either of two simultaneously available response keys. The delivery of food was scheduled by independent, time-based (variable-interval) schedules associated with each key (a *concurrent VI-VI* schedule). Herrnstein varied the average minimum interfood interval (i.e., the VI value) associated with these schedules

and determined the equilibrium rate of pecking on each alternative at each pair of VI values. Herrnstein was interested in the relation between the proportion of pecks on a key and the proportion of payoffs for pecking that key. But reinforcement for pecking a key can also be regarded as a way of affecting the number of pecks made on that key, independently of other activities; so we can look just at the relation between rates of pecking on the two keys, as the proportion changes under the influence of differential reward. Figure 7.1 shows that pecks on the two keys were reciprocally related — according to a straight line with unit slope — in Herrnstein’s experiment: as rate of pecking one key increased, rate of pecking the other decreased by an equal amount.

Data in the right-hand panel of Figure 7.1 show that this reciprocal relation does not depend either on constancy of total reinforcement rate or similarity of the schedules for each key (both of which were conditions of Herrnstein’s experiment). Bacotti used a variable-interval schedule for one key and fixed ratio for the other (*concurrent VI-FR*), and made no effort to keep total reinforcement rate constant. Nevertheless, his data show the same kind of reciprocal, linear relation between pecks on the two keys.

If we identify the number of pecks per unit time on the two keys as x_1 and x_2 , then these two experiments show them to be related by the linear equation

$$ax_1 + x_2 = b, \quad (7.1)$$

where a and b are constants: b is a scale constant that depends on the units of measurement, and a is the ratio of durations of the two alternative responses. Thus, animals typically peck at a faster rate on ratio schedules (see note 4), so that in Bacotti's experiment a pigeon will peck more times during a fixed period when it is responding on the ratio key than during a comparable time responding on the interval key (all else being equal). Consequently, if x_1 denotes the ratio schedule, then a is greater than unity; in the symmetrical situation of Herrnstein's experiment, on the other hand, a is close to one.

The functions in Figure 7.1 show the invariant relation between the rates of two activities as the balance between them is altered by external factors: here, different rates of reinforcement for each response, in other situations, variation in the properties of the stimuli controlling one or both activities. In the language of an earlier theoretical era, these functions describe *response-response* (rather than *stimulus-response*) relations.

Behavioral competition is of interest in its own right. It is still imperfectly understood, and experimental work continues. However, the bare fact of competition for available time has implications that can be traced even in the absence of detailed understanding of the process. In Chapter 10 I discuss implications for stimulus generalization and discrimination. Here I show how competition for time can be considered as one of the *constraints* to which animals adapt as they allocate behavior so as to maximize utility.

THE ALLOCATION OF BEHAVIOR

Imagine a rat in a semi-natural environment in which it can do several things, each of which we can record and make available or unavailable. There are two main kinds of experiment we can do with such a situation: (a) Add or remove the opportunity to do various things and look at the effect on the amount and temporal order of the remaining activities. (b) Impose *reinforcement contingencies* between pairs of activities and see if the instrumental response increases in rate - make the animal run in order to get access to water, and see if he runs more, for example. David Premack was the first to see a connection between these two kinds of experiment. In studies with rats and monkeys, he showed that the effect of a contingency of this sort depends upon the levels of the activities under free conditions (so-called *paired-baseline* conditions). I next describe one of his experiments and then explain two ways in which the imposition of a reinforcement contingency can be expected to change the levels of the activities involved.

In one experiment, Premack studied the reinforcing relations among the activities of Cebus monkeys. The monkeys were in individual cages and there were four things that they could play with: a lever (L), a plunger (P), a hinged flap (F), and a horizontally operated lever (H). Premack's idea was that "reinforcement" is not an absolute property of a "goal activity," but rather is just the relation between a more probable (preferred) activity and a less probable one: An animal should engage in a less probable activity (lever pressing, for example), for the opportunity to engage in a more probable one (say eating). His situation with the monkeys was more convenient for testing this idea than the usual Skinner box, both because more activities were available and because different animals showed different preferences for the different activities. He compared the frequencies of activities under paired-baseline conditions with the effect of imposing a 1:1 contingency between the several possible pairs of activities, that is, a contingency requiring t -s of activity A for t -s access to activity B; A here is termed the *instrumental response* and B the *contingent response*. Premack wrote:

The clearest predictions possible were those for Chicko, who in the first procedure [free access to all activities — paired baseline] showed three reliably different response probabilities [proportions of time spent]. Indeed, Chicko's protocol made possible three kinds of contingencies: contingent response [i.e., reinforcer] higher than, less than, and, in one case, about equal to the free [instrumental] response...the outcomes for the three types of contingencies were as follows: (1) contin-

gent response higher than free response produced...an increment in the free response; (2) contingent less probable than free response produced...a decrement...(3) the one case in which the responses were about equal produced little or no change, increment or decrement. (Premack, 1965)

Thus, a monkey that spent 10% of its time playing with L, say, and 20% playing with P under free conditions, increased its level of L when L, the instrumental response, had to occur for t sec in order for P, the contingent response, to be available for t sec. What can we learn from this about the regulatory nature of reinforcement, and about the difference between reinforcers and other events?

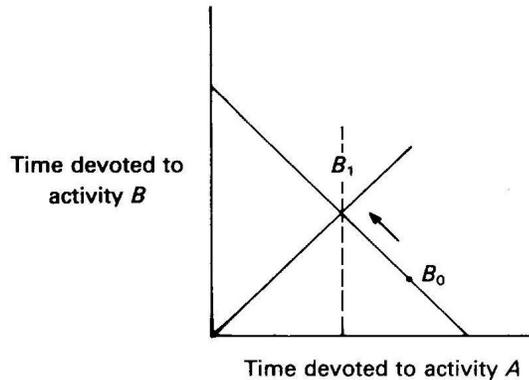


Figure 7.2. Schedule and time-allocation constraints with only two available activities.

of contingent activity, hence the unit slope). This is, of course, the feedback function for a fixed-ratio-one (FR 1) schedule. The arrow indicates the forced change (from B_0 to B_1) in B, the point representing the proportions of the two activities. It is clear that the two constraints — time allocation plus schedule — exhaust the degrees of freedom available to the animal in this simple situation. Because, by our initial assumptions, he must engage in one or other activity, the added constraint of the ratio schedule forces the initially low-probability activity (here the instrumental response) to increase above its free value, as the level of the contingent response decreases.

This increase is obviously *not* a “reinforcing” effect of the more probable activity on the less probable. It is not an adaptive response, but merely a forced change, caused by the two constraints: Because the schedule makes instrumental and contingent activities occur equally often, and because the contingent activity occurs more often than the instrumental under free conditions, the first effect of the schedule constraint is to restrict artificially the time devoted to the contingent activity. Since the animal must engage in one activity or the other, restriction of the contingent activity forces an increase in the instrumental response. This change is called a *restriction effect*.⁶

When the number of constraints is not greater than the number of activities, restriction effects are the only possible effects of a reinforcement schedule. When there are more activities than constraints, both restriction and *contingent* effects are possible — more on contingent effects in a moment.

The test for a restriction effect is to ask whether the increase in the instrumental response associated with the imposition of a contingency is greater than, or merely equal to, the increase produced by just restricting the proportion of time the animal can devote to the contingent activity. The effect of restriction is illustrated in Figure 7.2 by the dashed vertical line through point B_1 . Clearly, restriction of the contingent response to the level attained under the contingency condition yields the same increase in the instrumental response, despite the lack of any contingent relation between the two activities.

Experimentally, a restriction effect can be demonstrated by what is known as a *yoked-control* procedure. Two animals are involved: Animal A is given contingent access to an activ-

ity. In equilibrium, the contingent activity for animal A occurs at a certain rate with a certain temporal pattern. The activity is then permitted to animal B, the yoked animal, at exactly the same time as animal A makes it available for himself. Thus, both animals get the same frequency and temporal pattern of access to the activity, but the activity is dependent on an instrumental response in one case (A), but independent of the animal's behavior in the other (B): Animal A is on a FR 1 schedule; animal B on a *variable-time* (VT) schedule with interactivity times determined by the behavior of animal A. Under these yoked-control conditions, if animal B shows the same increase in the instrumental response as animal A, we can be fairly sure that the increase is a restriction effect, nor a contingent (i.e., reinforcing) effect.⁷

The essence of *reinforcement*, as an adaptive process, is an increase in the level of the instrumental response beyond the increase attributable to restriction. Such an increase is called a *contingent effect*. A contingent effect is possible only if the degrees of freedom in the animal's behavior exceed the constraints imposed by time and schedule: Given two constraints, at least three mutually exclusive and exhaustive activities must be available for an animal to show a contingent effect of the schedule.⁸ Even if there are enough degrees of freedom, yoked controls, or comparison with a comparable variable-time schedule, are necessary to be sure that an increase in the instrumental response goes beyond the effects of restriction.

Premack's early experiments lacked such controls. Nevertheless, later work has borne out his thesis that a more probable activity will generally reinforce a less probable one in which it is one-to-one contingent, that is, produce an increase greater than that expected merely from restriction. The one-to-one proportion is critical, for both restriction and contingent effects, however. For example, if we choose a ratio schedule different from 1:1 whose feedback function passes through B_0 in Figure 7.2, then there will be no increase in the instrumental response.

Premack thought of reinforcement as being particularly associated with activities — with eating rather than the intake of food, with running rather than the availability of a running wheel. This is not a necessary part of the regulatory view of reinforcement. Food-in-the-mouth, hence food naturally eaten, is a more effective reinforcer than food delivered intravenously, because of the incentive properties of taste (see Chapter 6). With or without taste, however, it is the rate of access to food that seems to be the critical variable.

Negative reinforcement makes the same point. Receiving electric shock is not an activity. Nevertheless, animals clearly have a preferred rate for such an event: zero. With this sole quantitative difference, electric shock can be treated as just the reverse of positive.

Requiring an animal to engage in a low-probability act, or experience an event that is generally avoided, for access to one of higher probability, has a reinforcing effect on the low-probability act, but it has a punishing (negatively reinforcing) effect on the high-probability act. Reinforcement and punishment are thus two sides of the same contingent relation. The only asymmetry is a practical one: Commonly used negative reinforcers, like shock, are usually events that can be delivered to the animal without its cooperation. This is not an essential difference. Animals can be punished by forcing them to engage in more of a low-probability activity than they would like (the schoolboy required to write 100 "lines"), and they can be positively reinforced by hypothalamic brain stimulation which requires as little cooperation as foot shock.

Thus, the essential difference between positive and negative reinforcement is just in the

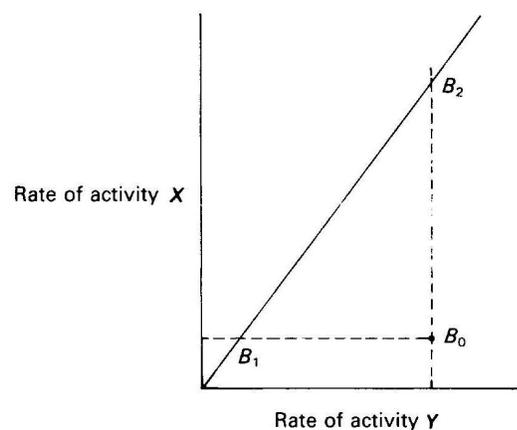


Figure 7.3. Potential adaptations to a ratio-schedule constraint.

relation between the free levels of the two activities and the direction in which these levels are driven by the contingency. For example, in Figure 7.3 (a redraw of Figure 6.1) the free levels of two mutually exclusive but non-exhaustive activities are indicated by B_0 , the free-behavior point. A ratio contingency is illustrated by the ray through the origin. Clearly, if the animal continues to engage in X at its free level, activity Y will be severely curtailed (point B_1). Hence, the contingency is potentially punishing with respect to Y , but potentially reinforcing with respect to X .⁹ The *actual* effect of the contingency depends on whether the activities change in the expected way. For example, it is logically possible either for X to show no change (behaviors at B_1 : a punishing effect on Y but no reinforcing effect on X), or for X to increase so much that Y is not reduced (behaviors at B_2 : a reinforcing effect on X , but no punishing effect on Y). In practice, either extreme result is rare.

Optimal allocation

What, then, determines the compromise struck by an animal constrained by a reinforcement schedule? Perhaps the simplest possibility is that he adjusts the levels of the instrumental and contingent responses so as to be as close as possible, in some sense, to the mixture under free conditions, the free-behavior point. In the following discussion I will argue that if activities are independent, quite general properties of schedule behavior can tell us how animals must assess the cost of deviations from their preferred mix of activities.

Cost and Value

First, a necessary digression on *cost* and *value*.¹⁰ If the free-behavior point represents the most desirable mix of activities, then *deviations* from that point must entail some cost. For example, suppose that the animal prefers to run 20% of the time but is prevented from running more than 10% of the time. Increments in running have value because they reduce the deviation of actual from preferred running level, and the amount of value they have depends on the cost of (negative) deviations from the 20% level. Conversely, if the animal is forced to run 30% of the time, reductions in running have value proportional to the cost of (positive) deviations from the preferred level of running. Thus, the value of a given activity level is the same as the cost of the deviation of that level from the preferred level, but with opposite sign.

Two beginning simplifications: First, I will suppose that the value associated with a given level of one activity is *independent* of the level of the other. If so, the total cost to the animal of given levels of the two activities is just the sum of the separate costs. Since B_0 , the free-behavior point, represents the most desirable mix, we can guess that the cost of a level of activity different from B_0 can be expressed as some function of the difference. If the coordinates of B_0 , for two activities X and Y , are (x_0, y_0) (i.e., X and Y occur with frequencies x_0 and y_0 during a unit time period), then the total cost to the animal of a level of joint activity (x, y) will be some increasing function of the difference, $x_0 - x$ and $y_0 - y$. Formally,

$$C(x, y) = f(x_0 - x) + g(y_0 - y), \quad (7.2)$$

where f and g are increasing functions. In optimality theory, $C(x, y)$ is known as the *objective function*, the quantity to be minimized.

Second, I assume that f and g are functions of the same type, for example, both linear, or quadratic, or power functions, or whatever. Thus, Equation 7.2 becomes

$$C(x, y) = f(x_0 - x) + f(y_0 - y). \quad (7.3)$$

Given these two simplifications, we can restrict function f in several ways. The most obvious possibility, is that f is just a linear function, that is, $C(x, y) = a(x_0 - x) + b(y_0 - y)$, where a and b are constant multipliers representing the relative costs of X and Y deviations. This can be eliminated as follows. First, look at what the linear cost function predicts under free conditions.

Obviously C can be made as small as we like by increasing x and y to large values, which is not plausible. Thus, a simple linear function will not do because it does not have a minimum for C at the free-behavior levels, x_0 and y_0 . The linear cost function fails to predict two basic properties of ratio response functions: that they invariably involve a compromise, and that they are non-monotonic, turning down at very high ratio values (see Figure 6.8). Conclusion: the cost function must be nonlinear.

Perhaps we should look just at the *absolute value* of deviations, a simple kind of nonlinearity. In this case, Equation 7.3 becomes

$$C(x, y) = a|x_0 - x| + b|y_0 - y|. \quad (7.4)$$

This does have a minimum at the free-behavior point under free conditions, but it predicts funny things when we introduce a schedule constraint between x and y , as in Figure 7.3. Although the argument is tedious, it is possible to show that with a ratio-schedule relating x and y ($x = my$, where m is the ratio value) Equation 7.4 predicts one of two possibilities: Depending on the relative values of a , b , and m (the ratio value), the animal should always choose either point B_1 or B_2 in Figure 7.3: if $b/a > m$ he should choose B_2 , if $b/a < m$, he should choose B_1 (I show this graphically later in the chapter). That is, he should perfectly regulate either x or y , never compromise. This never occurs.

To make things more realistic, suppose we add a constraint on total time, together with the possibility of a third activity, Z , to add a degree of freedom. The objective function is now

$$C(x, y) = a|x_0 - x| + b|y_0 - y| + c|z_0 - z|, \quad (7.5)$$

where c is the cost-of-deviation parameter for the third activity, Z . Again, the cost function predicts extreme results, depending on the relative values of a , b , c , and m . If $c > a$ or b , the animal should regulate Z perfectly. If not, he should never engage in Z . In either case, in the time remaining he should attempt to regulate either X or Y perfectly, as before, up to a limit imposed by the total time available. If the ratio value is so high that Y cannot be obtained at a rate y_0 , even if the animal spends all the available time engaged in X , then rates x and y will be constrained by a line of negative slope (Equation 7.1), representing the time-allocation constraint when all available time is taken up by X and Y . As we saw in Chapter 6, such linear response functions hold only for moderate (not extreme) ratio values. So an absolute-value cost function doesn't work either.

The problem is that both the linear and absolute-value cost functions violate a venerable economic principle: *diminishing marginal utility*. Everyone knows that the appeal of most things declines as one has more of them: An occasional dish of caviar is a treat; one a day may take the edge off one's enthusiasm; one with every meal is a bore. For most goods all the time, and for any good at some time, more and more, in quantity, means less and less, in added satisfaction.

There are some exceptions, of course. To a miser, every additional cent is worth having, no matter how many he already has. Increments in some things become more valuable as we have more of them: To a seamstress, enough material for a dress is worth a great deal more than too little, so the same increment is here worth more when added to a larger base.

These three possibilities are illustrated in Figure 7.4, which shows value plotted as a function of quantity of a good. Curve A fits the mi-

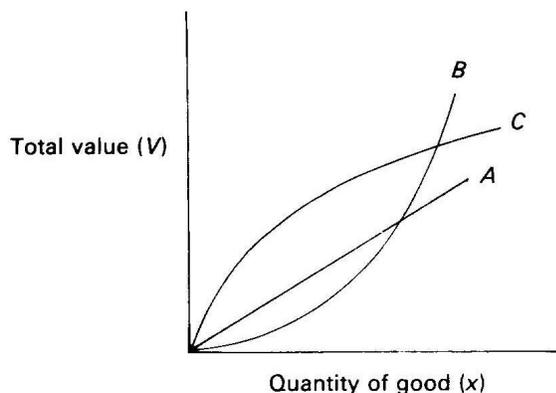


Figure 7.4. Linear (A), positively (B), and negatively (C) accelerated utility functions.

ser; value is simply proportional to amount. Hence, marginal value (technically, dV/dx , the slope of the curve) is constant and independent of amount. Curve B corresponds to the seamstress; a large piece of material is of more value than a small, but value increases dramatically when the size is sufficient for some new use: The curve is positively accelerated and illustrates increasing marginal value (dV/dx is a positive function of x). Curve C is negatively accelerated, fits almost everything else, and illustrates diminishing marginal utility (dV/dx is a negative function of x). While many utility (value) functions show regions that are linear or positively accelerated, essentially every good shows diminishing marginal utility at very high values - it is, after all, possible to have too much of practically anything.

In terms of classical utility theory, cost and value are on the same dimension, differing only in their sign: Values are positive, costs negative, and it is a matter of convenience which we use. Intuition suggests that the cost of an increment in the deviation of x (say) from x_0 (i.e., of a reduction in x) is likely to be greater the greater the existing deviation: A given reduction in rate of eating will be more costly if subtracted from a starvation than an ad libitum level, for example. This implies that function f in Equation 7.3 is *positively accelerated*. Positively accelerated cost is just the same as negatively accelerated value. Hence a positively accelerated cost function is equivalent to diminishing marginal value in terms of x , for values of x below the preferred level.¹¹

We now have all the ingredients for a simple economic analysis of behavioral allocation. Response-independence and a positively accelerated cost function allow us to estimate the value to the animal of a particular mix of activities. Time allocation and, most importantly, the reinforcement schedule, constrain the mix to values less than the maximum possible. These are the two components of standard optimality analysis: an *objective function*, and a *constraint set*.

THE MINIMUM-DISTANCE MODEL

The simplest nonlinear cost function is known as the *minimum-distance hypothesis*, for reasons which will be obvious in a moment. Figure 7.5 illustrates its application to ratio schedules. The figure shows just two activities: x , the rate of the instrumental response, and $R(x)$, the rate of the contingent response (reinforcement) actually obtained (a function of x). The two rays through the origin represent two ratio schedules, I and J. The points B_i and B_j represent the steady-state response and reinforcement rates settled on by the animal as it adapts to these schedules. Compare Figure 7.5 with Figure 7.3; note that both B_i and B_j are compromise adaptations, falling in between the two extreme options. The rule by which they were selected is the simplest one consistent with diminishing marginal utility: points B_i and B_j are just the *closest ones* on the ratio-constraint line (feedback function), to the free-behavior point, B_0 . The two dashed lines are therefore at right angles to the two ratio lines.

The minimum-distance model predicts a contingent effect of imposing a schedule. Point B_k in Figure 7.5 shows the response rate, x , expected from restriction alone — when reward rate, $R(x)$, is held to the value obtained on the higher ratio schedule. The response rate at B_k is lower than the rate at B_j showing that the animal responds faster when responding is necessary for reward than when it is not, a result different from the constrained case illustrated earlier in Figure 7.2.

Expressed in the form of Equation 7.3, the objective function for the minimum-distance

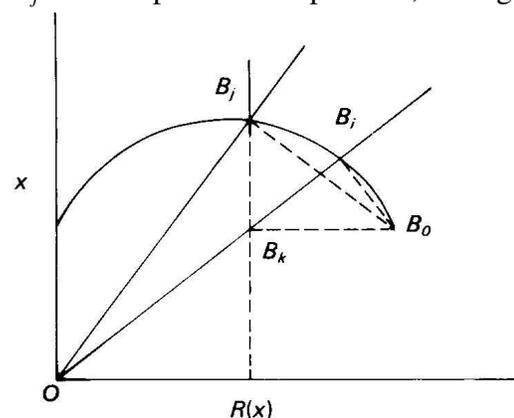


Figure 7.5. Minimum-distance adaptation to ratio constraint.

model is

$$C(x, R(x)) = [(x_0 - x)^2 + (R_0 - R(x))^2]^{1/2}, \quad (7.6)$$

where (x_0, R_0) are the coordinates of the free-behavior point, B_0 . Since the extrema (maxima and minima) of a power of any function are the same as those of the function itself, the overall square root can be dispensed with, leaving a quadratic cost function. The minimum of this function, subject to the ratio-schedule constraint, just gives the minimum (Euclidean) distance from B_0 to the constraint line. Hence this particular model, which is just one of many that are consistent with the independence and diminishing marginal value conditions, is termed the *minimum-distance* model. Clearly, the set of points representing adaptations to a range of ratio values will fall on a curve, which is just the ratio-schedule *response-function*, discussed in Chapter 6. Elementary geometry shows that the response function for these assumptions is a circle with diameter OB_0 , as shown.

The minimum distance-model correctly predicts that ratio response functions should be bitonic. Examples of empirical ratio response functions (similar to those already shown in Figure 6.8, but from a different experiment) are shown in Figure 7.6. The model fails in one respect, namely, the scales of x and $R(x)$. It is obvious from Figure 7.6 that if the two axes were to the same scale (and even allowing for the fact that x and $R(x)$ are not usually measured in equal-time units), the resulting curve would be very far from a circle, and closer to an ellipse much elongated in the vertical direction. This deficiency of the model is remedied by the very reasonable assumption that the *costs* of deviations in x and $R(x)$ are not equal. In the usual operant-conditioning situation, the reinforcer is something of considerable importance to the animal, such as eating or electric shock, X , on the other hand, will usually be a low-cost activity such as pecking or lever pressing. Indeed, low cost is almost the defining characteristic of most “standard” operant behaviors. Consequently, cost parameters need to be inserted in Equation 7.6, thus:

$$C(x, R(x)) = a(x_0 - x)^2 + b(R_0 - R(x))^2, \quad (7.7)$$

where a and b are the weighted costs of X and $R(x)$ deviations from B_0 : For a contingent response such as eating and an instrumental response such as lever pressing, a is much less than b .

One major deficiency remains. Ratio-schedule response functions generally have a pronounced declining limb and a relatively abbreviated increasing one (see Figure 7.6): They are not symmetrical, as Equation 7.7 requires. Indeed, over a considerable range, the ratio response function can be approximated by a straight line of negative slope, as we saw in the preceding chapter. It turns out that if we include the time-allocation constraint — and permit a third class of activities so as to have sufficient degrees of freedom for a contingent effect — this alters the minimum-distance prediction in the appropriate way. The necessary time-allocation constraint for three mutually exclusive and exhaustive activities is

$$x + R(x) + y = 1, \quad (7.8)$$

where all activities are measured in time units and y is the level of a third activity class. (Equation 7.8 is just an extension of Equation 7.1 to three activities.) In the usual operant-conditioning situation, y might be the time devoted to comfort activities (grooming, etc.), walking around or,

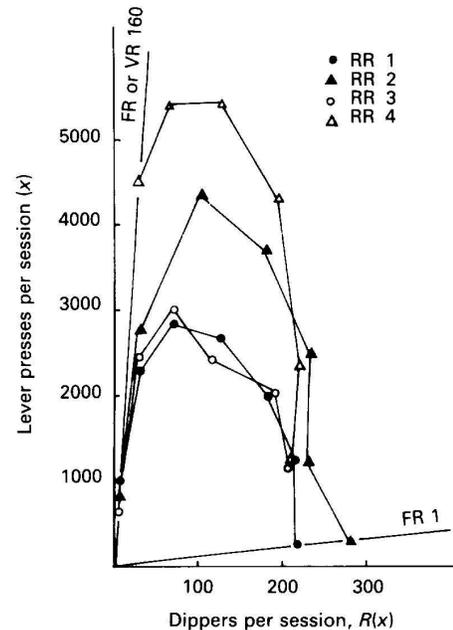


Figure 7.6. Response functions for two individual rats over the indicated range of ratio schedules. Open symbols: VR, closed: FR. Sessions were one hr long (Motheral, unpublished).

for a rat, running in a wheel. This third activity adds a third term to Equation 7.7, thus:

$$C(x, R(x)) = a(x_0 - R(x))^2 + b(R_0 - R(x))^2 + c(y_0 - y)^2, \tag{7.9}$$

and Equation 7.8 imposes a second constraint.¹²

The minimum of Equation 7.9, subject to the constraints of schedule and time allocation, can be found by standard calculus techniques.¹³ The predic-

tions of the model are illustrated graphically in Figure 7.7. The level of the third activity, Y , is represented by a third axis, at right angles to the X and R ($R(x)$) axes. The restriction that the three activities must occupy the total time (i.e., add up to a constant, see Equation 7.8) takes the form of a *behavior plane*, XYR , on which the point representing the levels of the three activities must lie. The ratio-schedule constraint is represented by a second plane, POY , the *ratio plane*, passing through the Y axis (since the ratio constraint relates x and $R(x)$, and is independent of y). The intersection of the behavior plane, XYR , and the ratio plane, POY , is the line PY . Clearly, the point representing allowable combinations of x , $R(x)$, and y , the *representative point*, must lie on line PY . In general, B_0 , the free-behavior point, will not lie on PY . The quadratic cost function implies that, for equal-cost activities, the representative point will lie at the point on PY closest to B_0 .

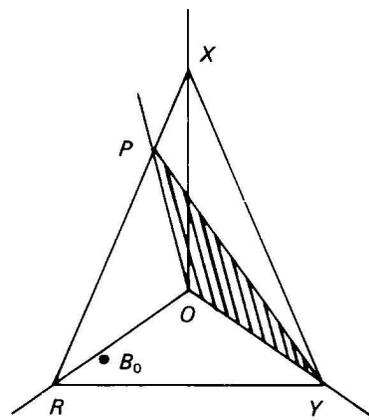


Figure 7.7. Joint effect of ratio constraint (plane POY) and time-allocation constraint (plane XYR) when three activities are available. B_0 , the free-behavior point, lies in the behavior plane XYR , but will not generally lie on the line of permissible behavior PY .

The minimum-distance prediction for three activities is illustrated in Figure 7.8, which shows the behavior plane, XYR , full on. Line PY is the set of allowable behavior mixes, as before, and B_1 is the minimum-distance solution. The figure shows two things. First, that the minimum-distance assumption predicts (as it should) that x should be higher when $R(x)$ is a positive function of x (e.g., a ratio schedule) than when $R(x)$ is independent of x . In other words, there will be a contingent effect, defined as the difference between the change in x produced by imposing a schedule and the change produced simply by restriction of $R(x)$. Line $P'Q$ represents the control case (this is just the three-dimensional equivalent of the vertical line — now a plane intersecting XYR — through B_1 in Figures 7.2 and 7.5). The minimum-distance adaptation to this condition is at B_1' , which represents a lower rate of x than B_1 (higher levels of an activity are represented by points closer to the vertex for that activity¹⁴).

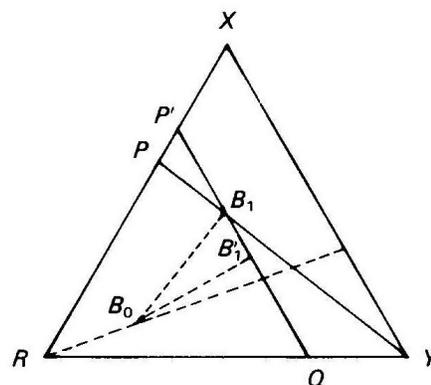


Figure 7.8. Plan view of behavior plane, illustrating restriction vs contingent effects in the three behavior case.

Second, Figure 7.8 shows that as one activity, such as $R(x)$, is progressively restricted, independently of the others (i.e., presented on a variable-time schedule), the other two do not remain in the same fixed proportion. Thus, as line $P'Q$ is moved progressively closer to R , the proportion of x (x -direction) to y (y -direction) changes as the representative point approaches B_0 along line B_0B_1' . Only in the special case where B_0 lies on the line joining R and B_1 will the proportions of x and y remain constant and independent of the level of $R(x)$. This point is worth making because some theories of choice¹⁵ say that the relative value of a given pair of alterna-

tives should be independent of other alternatives that may be available.

Prediction of response functions

Figure 7.9 shows in a three-dimensional graph the response function derived from the minimum-distance assumption under both schedule and time-allocation constraints. The response function is generated in the behavior plane; the relation between x and $R(x)$, with y eliminated (which is what is usually measured, as in Figure 7.6) is just the projection of this onto the XOR plane. Figure 7.10 shows theoretical ratio response functions. Parameters a and b are the costs of x - and $R(x)$ -deviations, respectively; the five curves represent the effects of different costs of y -deviations. As this cost increases, one effect is to tip the theoretical response functions to the left in the way required by empirical results. The reason for this tipover is not obscure. At high ratio values, the trade-off between x and $R(x)$ is very unfavorable: The increment in $R(x)$ produced by an increment in x is small. Correspondingly, the time taken away from y by the increment in x necessary to yield a fixed increment in $R(x)$ is, at high ratios, substantial. As the cost of y -deviations increases, therefore, the high rates of x associated with high ratios are especially depressed; the time is better spent on y than on x , given the small return for the latter in terms of $R(x)$.

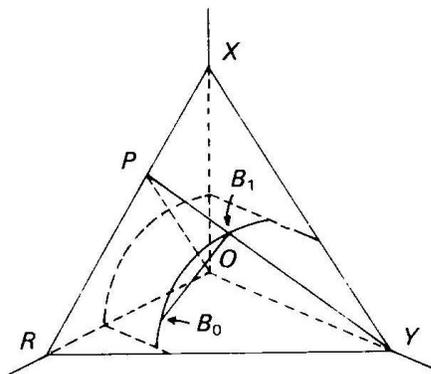


Figure 7.9. Geometrical derivation of the ratio-schedule response function from the minimum-distance assumption.

Response functions can be derived in this way for any kind of feedback function, that is, any kind of reinforcement schedule. For example, the feedback function for interval schedules (see Chapter 5) is negatively accelerated rather than linear, with a fixed maximum rate of reinforcement. The minimum-distance response function for interval schedules is consequently different from the corresponding function for ratio schedules: It is lower for a given rate of reinforcement, and tipped to the right rather than the left. Again, the reason is fairly obvious intuitively. At long interval values (low maximum-reinforcement rates), as response rate increases, the animal very soon approaches the diminishing-returns region of the feedback function where increments in responding produce smaller and smaller increments in reinforcement rate. There is, of course, no such region in the linear ratio feedback function. At the same time (as in the ratio case) more and more time is being taken up by the instrumental response so that other activities are forced further and further from their free-behavior values — at ever-increasing incremental cost, because of the positively accelerated cost function. The increasing cost of these deviations, together with the diminishing payoff entailed by a negatively accelerated feedback function, means that the compensatory increase in response rate as reinforcement rate falls (i.e., as the interval value increases) is soon overwhelmed, so that response rate and reinforcement rate fall together at long interval values. The small increment in reinforcement rate that can be obtained by large increases in response rate is not, at long interval values,

The increasing cost of these deviations, together with the diminishing payoff entailed by a negatively accelerated feedback function, means that the compensatory increase in response rate as reinforcement rate falls (i.e., as the interval value increases) is soon overwhelmed, so that response rate and reinforcement rate fall together at long interval values. The small increment in reinforcement rate that can be obtained by large increases in response rate is not, at long interval values,

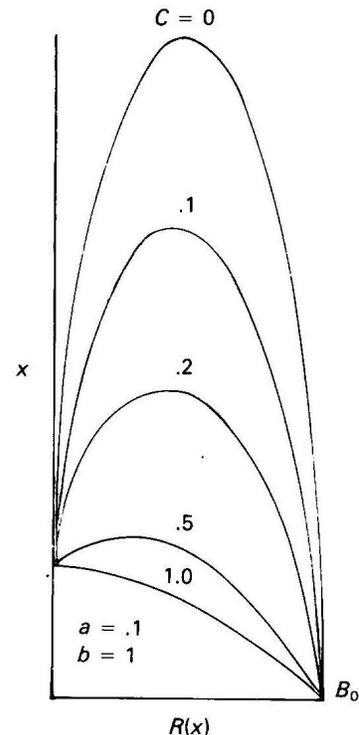


Figure 7.10. Theoretical effects on the ratio response function of increasing the cost of “other” behavior. Cost of instrumental response = a , contingent response = b , other = c .

worth the added costs - of responding and in time taken away from other activities. Hence, the increasing limb of the interval-schedule response function dominates, whereas the decreasing (regulatory) limb is the dominant feature of the ratio response function (see Figure 7.15).

This discussion of the predictions of economic models has necessarily been rather abstract. It is important, therefore, not to lose sight of the essentials, which are that the basic properties of molar adjustments to ratio, interval (and, as we will see, numerous other) schedules are derivable from just three general considerations: The values of activities are independent; activities compete for the available time; and activity values are subject to diminishing marginal utility.

The minimum-distance model is convenient because it lends itself to graphical illustration and because conclusions true of it are also true of almost any model that incorporates diminishing marginal utility. It is one way of generating what in economic parlance may be termed the standard preference structure for partially substitutable goods. Since many of the things I have discussed in terms of costs of deviations and so forth can also be (and often are) described in the economists' language of substitutability, complementarity, and indifference curves, the next section looks at the relations between the two forms of analysis.

Substitutability and indifference curves¹⁶

Psychologists and philosophers differ on whether it makes sense to give numbers to utilities (some psychologists say you can, most philosophers say you can't), so the analysis I have just

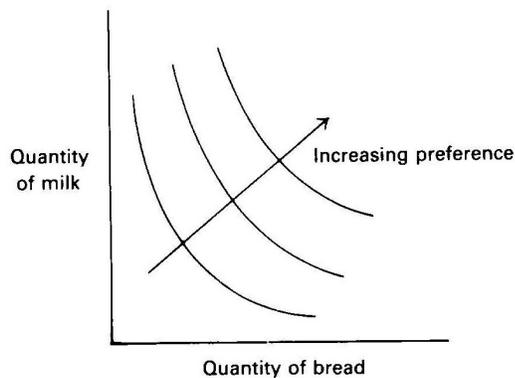


Figure 7.11. Conventional indifference curves.

described, in which behavior is derived from an assumption about the way the utility depends on the quantity of something, is no longer the mode in economics. No matter, we can get to the same place by a different route. Consider two goods, such as bread and milk. Even if we can't give numbers to the utility of a given quantity of milk and bread, everyone seems to agree that we can always *equate* the value of bundles of goods. For example, 2 liters of milk and 3 loaves of bread may be judged equal to another bundle with 4 liters of milk and 2 loaves, in the sense that we are indifferent as to which bundle we get. There will be a whole set of bundles of this sort, differing in the proportions of bread and milk, but the same in that we are indifferent among them. This set defines an individual *indifference curve*: a set of points in a space whose axes are the amounts of the goods being compared. Other names for indifference curve are *preference isocline* and *iso-preference curve*.

All agree that bundles can be rank-ordered, A being preferred to B, B to C, and so on. Bundles so ranked obviously cannot lie on the same indifference curve; consequently, to rank-order bundles is also to rank-order indifference curves. The *preference structure* for bread and milk, or any other pair of goods, can therefore be represented by a sort of contour map of indifference curves, such as those illustrated in Figure 7.11. The arrow indicates the direction of increasing preference which, unsurprisingly, is generally in the direction of

the proportions of bread and milk, but the same in that we are indifferent among them. This set defines an individual *indifference curve*: a set of points in a space whose axes are the amounts of the goods being compared. Other names for indifference curve are *preference isocline* and *iso-preference curve*.

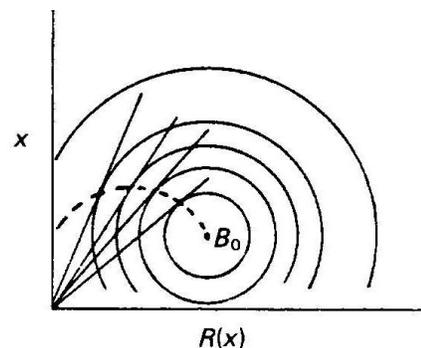


Figure 7.12. Indifference curves derived from the minimum-distance model. Rays through the origin are ratio constraints, and the dashed curve is the expected response function.

more of both goods.

It is obvious that a utility or cost function, such as Equation 7.9, defines the form of the corresponding indifference curves (although it is not possible to go in the reverse direction and derive a unique utility function from a given set of curves). For example, suppose that for goods X and Y, the utility functions are $f(x)$ and $g(y)$. An indifference curve is then defined by the relation $f(x) + g(y) = \text{constant} = K$, each value of K being associated with a different curve. Thus, for the quadratic cost function, the equation for an indifference curve is $a(x_0 - x)^2 + b(y_0 - y)^2 = K$, which is an ellipse (for $a = b$, a circle) centered on x_0, y_0 .

Knowledge of the preference structure is not by itself sufficient to specify behavior — in the absence of any constraint, the organism has no reason to stray from the point of maximum value. Once a constraint is specified, the optimal solution is clearly to settle on the highest indifference curve consistent with it. For linear, or curvilinear, constraints, this implies an equilibrium at the point where the constraint line is tangent to the highest indifference curve. This is illustrated for the circular indifference curves of the equal-cost quadratic model in Figure 7.12. Here the constraint is simply the linear feedback function for a ratio schedule. The dashed line is the ratio-schedule response function. As you can see, indifference-curve analysis gives the same result as the direct minimum-distance method shown earlier in Figure 7.5.

The indifference-curve method sometimes makes it easy to see things that are otherwise not obvious. For example, Figure 7.13 shows the indifference curves generated by the absolute-value model discussed earlier (Equation 7.4). The ratio of height to width of the diamond-shaped indifference contours is equal to the relative costs of the instrumental and contingent responses.

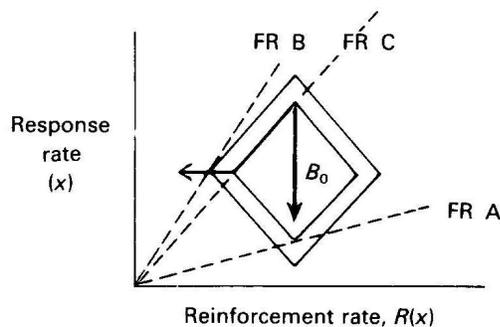


Figure 7.13. Prediction of ratio response functions from indifference curves generated by the absolute-value linear model (Equation 7.4). See text for details.

The rays through the origin are three ratio feedback functions. Optimal responding is indicated by the point where each feedback function is tangent to an indifference curve. The critical ratio value is the center one, FR C. It is pretty obvious that for any ratio value less than C (such as A), the animal should regulate the contingent response perfectly: The response function lies on the vertical, arrowed line. Whereas for any ratio value greater than C (such as B), he should regulate the instrumental response perfectly: The response function for these ratios lies on the horizontal arrowed line. For ratio C, any point on the feedback function connecting these two lines is optimal. Thus the indifference-curve analysis shows at once the extreme properties of a linear cost function.

Marginal value and substitutability

Economists talk a lot about the *marginal value* of things. The reason is that usually the way to maximize value is to allocate resources (e.g., money) so that the *increment* in value associated with the last little bit of resource is the same no matter to what we allocate it: Given \$100 to spend on any mixture of three commodities, bread, milk, and beer, maximum satisfaction is assured if we are indifferent about whether to spend our last dime on more milk, bread, or beer. Indifference means that we have equated the marginal values of the three commodities. Those familiar with differential calculus will recognize that equating marginal value is just the mathematical operation of equating the partial derivatives of the value (or cost) function with respect to each commodity.

The concept of marginal value is useful even without getting into mathematics. For example, it explains why negatively accelerated utility functions (positively accelerated cost func-

tions) turn up a lot. Consider how one should allocate a fixed amount of time between two activities whose values are each simply proportional to time spent: $V_1 = aT_1$; $V_2 = bT_2$. Because the value functions are linear, marginal value is constant in each case. Hence, no pattern of allocation can equate marginal values, and the best course is to devote all the time to one activity or the other, depending on whether $a > b$, or the reverse. When utility functions are negatively accelerated, on the other hand, the more one does of something, the smaller the marginal benefit, so that by adjusting the amounts of different activities it is always possible to equate marginal utilities. The fact that we usually observe *behavioral diversity*, animals doing several things within a given time period, implies negatively accelerated value functions.

Suppose we have a certain amount of one thing and are now offered the possibility of a trade: give up some of good A in exchange for some of new good B. The corresponding behavioral problem is to restrict the level of one activity and ask how the others trade off against one another to fill up the time made available. Another way to look at the same question is in terms of price: If milk becomes more costly, less will generally be bought, and more of other things, like beer, will be bought. If all the reduction in spending on milk is taken up by spending on beer, then beer and milk are acting as perfect *substitutes*. In an operant experiment, a rat may be responding on one lever for cherry cola, and on another for Tom Collins mix;¹⁷ if the price (ratio value) for one commodity increases, by how much will responding for the other increase? The experimental answer tells us to what extent the two drinks are substitutes for one another.

The extent to which one thing is a substitute for another is measured by the *marginal rate of substitution*, or *substitutability*. Given a certain quantity of any good, an individual will be prepared to give up a small amount in return for some amount of another good. The marginal rate of substitution is just the ratio of amount given to amount received, evaluated at a point, that is, as the absolute quantities traded become very small.¹⁸ It is pretty obvious that this ratio will usually vary as a function of the amount of the first good. For example, given a stock of only 1/10 of a liter of milk, and no bread, one might be prepared to give up only 1/20 liter in exchange for a loaf of bread; but given a stock of 5 liters, 1 liter of milk might willingly be exchanged for the same quantity of bread. This can be looked at either as an indication of diminishing marginal value of milk, or of diminishing marginal rate of substitution between milk and bread.

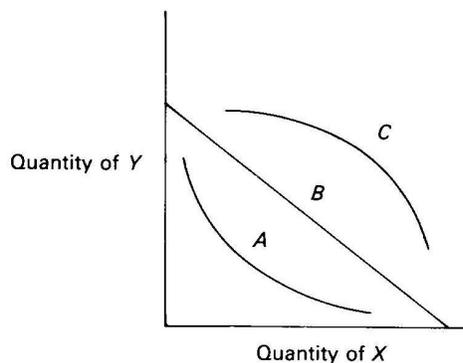


Figure 7.14. Indifference curves showing decreasing (A), constant (B), or increasing (C) marginal rates of substitution.

Marginal rate of substitution can also increase or remain constant as a function of amount of a good, corresponding to the cases of increasing and constant marginal utility already described. These three possibilities are illustrated in Figure 7.14: Curve A is the familiar case of decreasing marginal substitution; curve B shows constant marginal rate of substitution — cases of partial and complete substitutability, respectively. Curve C corresponds to increasing marginal rate of substitution.¹⁹

We have already seen that because of the universal time-allocation constraint, if one of an animal's activities is forced to decrease in frequency, at least one other must increase in frequency. How all other activities change when the frequency of one is controlled depends on substitutability relations. For example, with just three activities, if one decreases, at least one of the others must increase; but the remaining activity may either increase or *decrease*. In the latter case, the two activities that change in the same direction are termed *complements*. Everyday examples of complementary goods are meat and meat sauce or cars and gasoline: If consumption of meat goes down (perhaps because of an increase in its price, perhaps

because of a diet fad), consumption of meat sauce is also likely to decrease. There are comparable pairs in the behavior of rats on schedules. For example, given periodic, brief access to food (i.e., a controlled food rate, a fixed-time schedule) rats will drink after each food pellet. Over much of the range of food rates, drinking rate and food rate covary — drinking and eating thus act as complements. How can we decide on the direction of change of the free activities when one is restricted?

It turns out that for the quadratic model — indeed for any model of the form $C(x) = a(x_0 - x)^m$, where x_0 is the free-level parameter, a is the cost, and the exponent m is the same for all activities²⁰ — all activities are partial substitutes, none are complements. For all such models, as one activity is progressively restricted, the levels of any pair of other activities will both increase along a linear function whose slope is the ratio of their costs. If the requirement for equal exponents is relaxed, then restriction of one activity can lead to a decrease in some others. Since complementary activities exist, the equal-exponent simplification proposed during the earlier discussion of Equation 7.2 is not always justified.

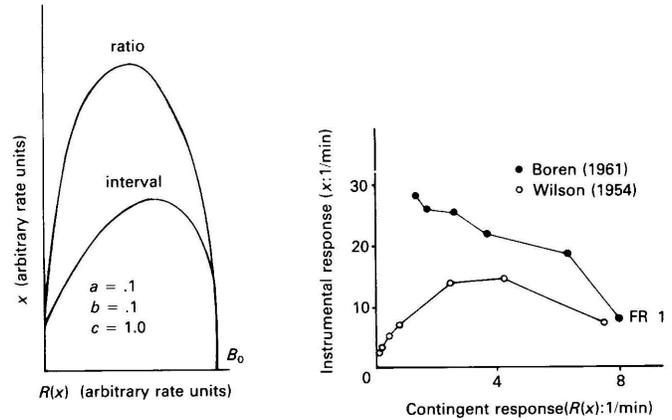


Figure 7.15. Left panel: theoretical response functions for interval and ratio schedules. Right panel: empirical interval and ratio response functions. (From Staddon, 1979a; redrawn from Boren, 1961, and Wilson, 1954).

The contours of an indifference-curve representation, in as many dimensions as there are independent activities, are a map of an organism’s “motivational space” (*preference structure*, in economic language). All that is assumed by optimality analysis is that such a space exists (preferences are not irreconcilably inconsistent) and that it is *invariant* (preferences do not change capriciously). This is known as the doctrine of *revealed preference* (see note 10). To go beyond this and attempt to explain the properties of the preference structure, that is, to derive the set of indifference curves, will generally entail some assumptions about cost or utility. The minimum-distance model is one such assumption. Thus, utility is a theoretical notion but indifference curves are theoretical only in presupposing the stability of preferences.²¹

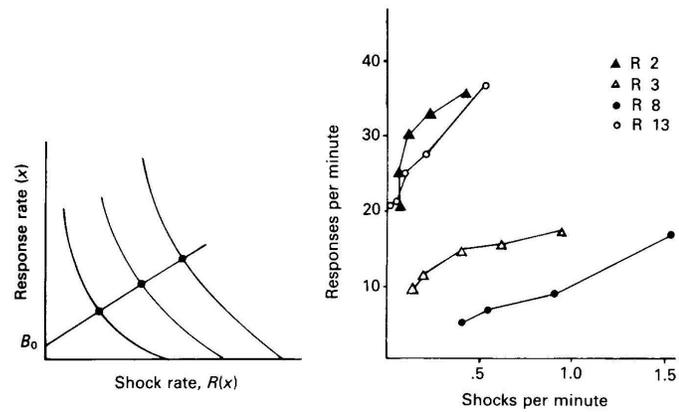


Figure 7.16. Left panel shows hypothetical feedback functions (descending curves), for a shock-postponement schedule, and an estimated response function (ascending curve). The right panel shows empirical response functions from four rats exposed to such schedules; each point is the average of an ascending and descending run (de Villiers, 1974). See text for details.

Goods that are partial or complete substitutes — beer and wine, cars and trucks, for example — generally have some property or properties in common. Similar activities are better substitutes than dissimilar ones. It is intriguing to speculate that all observed activities partake to different degrees of different amounts of a much smaller number of what might be termed (by analogy with ecology) *resource axes*, that are

truly independent of each other. Activities are substitutable, then, because they lie on neighboring (non-orthogonal) axes within this hypothetical space. This is an appealing but largely untested possibility (see note 12).

EXPERIMENTAL APPLICATIONS

The power of the optimality approach lies in its ability to predict the differences among the molar response functions associated with different feedback functions (different schedules of reinforcement). Available data from “standard” situations are consistent with the three assumptions I have already discussed: value independence, competition, and diminishing marginal value.²² For example, the left panel of Figure 7.15 shows the response functions for interval and ratio schedules predicted by the minimum-distance model. The right panel of Figure 7.15 shows response-function data from two similar experiments with rats pressing a lever for food reinforcement. The differences between the curves for ratio and interval schedules are of the sort predicted: The ratio curve is higher at all but the very highest reinforcement rates (corresponding to very short interval values, when an interval schedule cannot be distinguished from FR 1 by the animal); and the ratio response function is tipped to the left whereas the interval function is tipped to the right.

Data from other studies of positive reinforcement, such as interlocking, conjunctive, and tandem schedules (see Chapter 5), seem also to be consistent with these assumptions, to the extent that systematic data are available, and granted the difficulties in arriving at exact, molar feedback functions.²³

Experimental results from schedules of shock postponement also fit into the present analysis. In avoidance schedules, the feedback function is a declining one, that is, as response rate increases, shock rate decreases (see Chapter 5, Figure 5.11). The free-behavior point is close to the origin, since unconstrained animals will choose to receive no shocks and will make few instrumental responses. The response function expected from minimum-distance-type models will be an increasing one: As the free shock rate (i.e., shock rate in the absence of instrumental responding) increases, response rate should also increase, as shown in the left-hand panel of Figure 7.16. The right-hand panel shows data from an experiment in which rats received shocks at a constant rate if they failed to press a lever. A lever press cancelled the next scheduled shock; additional presses had no effect until the time for the avoided shock had passed. The rats were permitted to establish a stable rate of lever pressing at each shock rate. As expected, equilibrium response rate is a positive function of shock rate. Results from numerous other shock-postponement experiments fit this general pattern.

We saw earlier that if activities other than the instrumental and contingent responses are neglected, the minimum-distance model predicts ratio response functions that are too symmetrical. When these “other” activities are included (that is, their cost-of-deviation and free level are allowed to be greater than zero), the response function tips to the left — the declining (regulatory) limb of the function becomes shallower. Moreover, the greater the cost of “other” activities, the shallower the regulatory limb (see Figure 7.10). We can

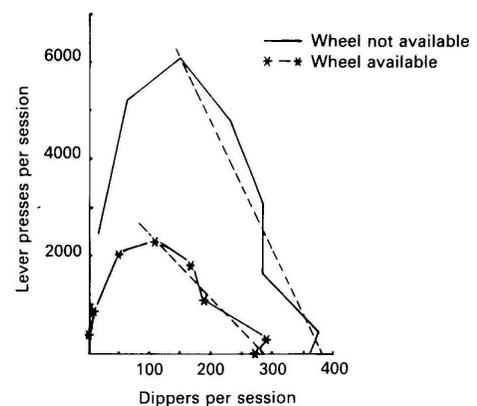


Figure 7.17. Effect on the ratio-response function of introducing a valued competing activity. Upper curve shows the response function for a group of four rats pressing a lever for food with no running wheel available. The lower curve shows the response function when a running wheel was available (Motheral, unpublished data).

test this prediction by looking at the effect on the ratio response function of adding a third activity.

Figure 7.17 shows the response functions obtained from a group of four rats responding on a range of ratio schedules. The upper curve was obtained in a standard Skinner box, with no opportunity provided for competing activity. The lower curve shows the response function obtained with a running wheel concurrently available. Not only is it depressed, as would be expected just from competition, but the regulatory limb is also shallower as minimum-distance requires (compare the slopes of the dashed lines drawn through the declining limbs of the two functions). A similar shift can be produced by giving the rat constant access to an alternative,

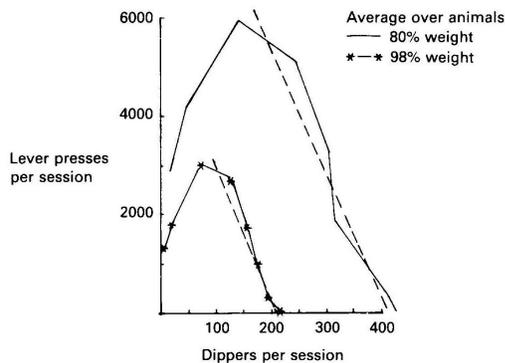


Figure 7.18. Effect on the ratio response function of level of motivation. Upper curve is the ratio response function for four rats maintained at 80% of their normal weight. Lower curve is the response function for the same group maintained at 98% of their normal weight (Motheral, unpublished data).

but less preferred, food (Lea & Roper, 1977). The point is that the slope of the declining limb of the response function depends on the *relative* cost-of-deviation of the contingent response versus other activities. If these activities have high cost-of-deviation, then the rate of the contingent response will be less well regulated (have a shallower declining limb) than if they have low cost-of-deviation. Figure 7.17 shows that the availability of running impairs the regulation of eating. Lever pressing (by which eating is regulated here) and running compete for the available time, so that both cannot, at the same time, be perfectly regulated. If running is a relatively preferred activity (which it is), then adding the opportunity to run must impair the regulation of eating rate.

I argued in Chapter 6 that food deprivation should improve the regulation of eating rate. Data from cyclic-ratio experiments support this contention

(Figure 6.9). These results show only the regulatory limb of the response function, and the slope is indeed steeper for the more deprived animals. This result agrees with the commonsense argument that cost-of-deviation in eating rate should be higher for a food-deprived animal than for one close to its normal body weight. This result is not always found. Figure 7.18 shows ratio response functions obtained from a group of rats at 80% and 98% of their free-feeding weights. The difference between the response functions is almost entirely in the value of the free-level of eating (R_0).²⁴ In this experiment, each ratio was in effect for one session or more (i.e., the cyclic-ratio procedure was not used). Whether this difference in experimental method is responsible for the different effects of food deprivation in the two studies, or whether something else is involved, remains to be determined.

Reinforcement constraints

In the history of reinforcement theory much discussion has been devoted to limitations on the reinforcement principle.²⁵ Why, for example, is it easy to train pigeons to peck for food, but hard to get them to peck to avoid electric shock? Why is it relatively difficult to sustain treadle pressing with food reinforcement and almost impossible to reinforce wing flapping with food? A popular solution is to invoke physiological “constraints” that prevent the animal from making certain kinds of “associations.” Since physiology underlies everything, and associations cannot be observed directly, this kind of account is not satisfying. It serves merely to push the problem back from a region where direct test is possible, observable behavior, to a different level, where, at present, it is not. Always be suspicious of a behavioral scientist invoking physiology: The ap-

peal is likely to be motivated less by faith in physiology than by a reluctance to inquire further.

A different tack is to turn to the ecology of the animal. Pigeons, it is argued, have evolved in an environment where pecking at a potential food item is usually the best way both to test and to eat it. Wing flapping, on the other hand, is useful for avoiding predators and getting from place to place (especially over long distances; it is too expensive in energy cost to waste on short trips, as anyone watching pigeons scuttle out of the way of a car will have noticed). There has been no prior selection for wing flapping to get food. Perhaps, therefore, special links have evolved to “prepare” pigeons to peck for food.

There is much to this argument, although it runs into difficulties in explaining numerous tricks that animals can learn for which they have no obvious evolutionary preparation — seals balancing balls, bears riding bicycles, not to mention human language and other symbolic activities. It also suffers from the weakness of any purely evolutionary argument about behavior: There is little to restrain it; it makes few testable predictions, and does not relate things that an animal finds easy to learn to things that it finds difficult. The ecological approach gives no clue to the mechanisms that must underlie both constrained and unconstrained learning. We will have to do better than just label things easy to learn “prepared” and things hard to learn “unprepared.”

There is an alternative: that most of these effects can be explained by the properties of the animal’s preference structure, by differences in the cost-of-deviation (CoD) of different activities. The argument is straightforward: The effectiveness of one activity as reinforcer for another is directly related to their relative CoD values: the greater the CoD of the contingent response relative to the instrumental, the more the instrumental response will change in response to a given deviation of the contingent response from its preferred level. The CoDs of various activities are presumably the outcome of past natural selection, so this account is not in conflict with a functional one; it is just more explicit. Learning (reinforcement) constraints follow from this idea because a low-cost activity makes a poor reinforcer for a high-cost one, both in terms of the absolute rate changes that can be achieved and, more importantly, in terms of the slope of the regulatory limb of the ratio response function. The effects of CoD value can be shown analytically in the following way.

Imagine the animal to be at its free-behavior point; consider the effect on the rate of the instrumental response of an increment in the ratio requirement. If, for a given ratio increment, the increment in instrumental responding is large, the sensitivity of the response to this reinforcer is obviously high; conversely, if an increment in the ratio requirement produces little change in the level of the instrumental response, the reinforcing effect is small — the response is insensitive to this particular reinforcer. Thus, the slope of the response function in the neighborhood of B_0 provides a measure of the *sensitivity* of a particular response to a particular reinforcer. Formally, it is convenient to define the quantity

$$S(u, v) = \partial x / \partial R(x) |_{u, v}, \quad (7.10)$$

that is, the partial derivative of x (the rate of instrumental response) with respect to $R(x)$ (the reinforcement rate), evaluated at a point u, v (here, the free-behavior point, so that $u = x_0$ and $v = R_0$) as a quantitative measure of sensitivity. When the absolute value of S is high, x is sensitive to reinforcement; when S is low, x is insensitive.

For the minimum-distance model, the expression for sensitivity turns out to be

$$S(x_0, R_0) = -[R_0(b + c) + cx_0] / [x_0(a + c) + cR_0], \quad (7.11)$$

where a , b , and c are the CoDs of the instrumental, contingent (reinforcing), and other activities, respectively. Since all the terms are positive, the expression as a whole is negative, indicating the regulatory character of the system. If c , the cost of other activities, can be neglected, Equation 7.8 reduces to

$$S(x_0, R_0) = -bR_0/ax_0, \quad (7.12)$$

so that the sensitivity measure is approximately proportional to the relative costs of the contingent and instrumental responses. Consequently, relative cost is a tolerable measure of the “reinforcibility” of one activity by another.

Hard-to-reinforce activities may, therefore, be high-CoD activities — but in that case they should be good reinforcers themselves. Conversely, low-CoD activities should be easily reinforced by many other activities, but should be poor reinforcers themselves. Presumably activities such as lever pressing, perch hopping, and key pecking, commonly used as instrumental responses, all fall into this category. Indeed, the cost dimension has precisely the properties of the older distinction between *appetitive* or *precurrent* activities and *consummatory* activities advanced by Sherrington and, independently, by the American comparative psychologist Wallace Craig. Appetitive activities are the variable, sometimes exploratory, activities that lead up to consummatory activities such as eating or copulation. Appetitive activities are variable and labile (that is, reinforceable), whereas consummatory activities are stereotyped and intractable. Thus, *appetitive* corresponds to “low CoD” and *consummatory* to “high CoD.”

Are all apparent constraints on the effects of reinforcement attributable to an invariant preference structure, or are there indeed special links between some activities and special limitations on others? Such links or limitations should show up as inconsistent relations between sensitivities. For example, consider the sensitivity relations among three different activities, A_1 , A_2 , A_3 . There are six possible relations, since three activities can be paired up in three ways, and each can serve as both instrumental and contingent response. Thus, in principle, a set of six S values could be obtained, yielding six equations of the form $S(i, j) = -K_{ji}c_j/c_i$, where K_{ji} is the ratio of free levels of activities j and i , and c is the unknown cost of each activity.²⁶ Thus, there are three unknown costs and six equations. Inconsistencies will be revealed if the costs differ substantially depending on which three equations are solved.

The analysis is particularly simple if instead of sensitivities we use *elasticities*, the ratio of proportional, rather than absolute, changes. The *cross elasticity* for two goods x and y is defined as $(\partial x/\partial y)(y/x)$. Thus, from Equation 7.12, the elasticity of two activities on the ratio response function evaluated at B_0 is just b/a , the ratio of their costs. A reinforceable response, one for which b/a is large, will be highly elastic, showing a large proportional increase in response to an increase in ratio value. Conversely, an inelastic response will be poorly susceptible to reinforcement, showing only small increments. For our three activities, two comparisons, between A_1 and A_2 and A_1 and A_3 , yield two elasticities: c_1/c_2 and c_1/c_3 . From them, a third, c_2/c_3 , can be obtained and compared with the results of a third experiment comparing A_2 and A_3 .

Unfortunately, I have not been able to find any experiment in which systematic comparisons of this sort have been made. Nevertheless, there are numerous less-systematic studies showing apparent constraints of a sort quite compatible with a fixed preference structure. For example, male Siamese fighting fish, *Betta splendens*, will regulate their food intake almost perfectly when all their food must be obtained on a ratio schedule for which the response is swimming through a short tunnel.²⁷ Here, swimming is a highly elastic response. Conversely, *Bettas* increase their frequency of swimming through the tunnel only slightly when the reinforcement is sight of another male *Betta* (a stimulus to which they display): Regulation is poor and here swimming appears to be highly inelastic. Both these results are consistent with a lower CoD for display reinforcement than for food reinforcement. To show real inconsistency, a further experiment is necessary in which some other instrumental response, say fin spreading, is highly elastic for display reinforcement, and inelastic for food. In the absence of such a test, the differences between food and display reinforcement are easily handled by quantitative differences in a cost parameter.²⁸

The differences between food and water as reinforcers seem also to be largely in the cost

parameter. Lever pressing, as instrumental response, is generally less sensitive to water than to food reinforcement, for example. Heat and access to nest material by mice also behave like low-CoD reinforcers. Whereas access to a running wheel is, for a rat, a relatively high-CoD activity almost as effective as food in maintaining lever pressing.

Perhaps the clearest demonstration of real constraints comes from attempts to reinforce “comfort” activities such as grooming. For example, Shettleworth (1975) studied the behavior of golden hamsters (*Mesocricetus auratus*) under free and constrained conditions and found that face washing occurred at a high level under free conditions, yet was easily displaced by other activities. Easy displacement implies low CoD, hence high sensitivity to food reinforcement. Yet another experiment showed face washing to be *insensitive* to food reinforcement.

There is anecdotal evidence that birds will readily take flight to avoid an electrified grid, even though they can be taught to peck to avoid electric shock only with great difficulty. Yet pecking is easily entrained by food reinforcement, and wing flapping (if not actual flight) is poorly reinforced by food. Thus, although more systematic data would be desirable, there is little doubt that real constraints, not easily handled by a simple economic model based on diminishing marginal utility, do exist. Most effects termed *constraints* are mislabeled, however, being better explained as cost differences.

The very idea of a constraint — on reinforcement, learning, or whatever — has meaning only in relation to some norm. It is only if a given reinforcer fails to affect a given instrumental response in the expected way that appeal is made to a “constraint.” At one time it was thought that all reinforcers should act with equal facility on all responses. This idea was shown to be incorrect almost as soon as people began to look at responses other than the prototypical lever press and key peck, and at reinforcers other than food. We have just seen that many of these limitations may reflect nothing more than CoD differences — constraints on learning that might well have been anticipated, had psychologists been thinking in economic terms. Other limitations remain that cannot be so accommodated, however: They are “real” constraints, relative to the norm of the minimum-deviation model — although they may not be constraints in relation to some more comprehensive economic model.

Still, no matter how elaborate the model, some anomalies will probably always remain. Future work will have to decide whether these can usefully be accommodated within economic analyses — or whether a point has been reached where economics must give way to an understanding of the *processes* involved. The kind of analysis described in this and the preceding chapter is neither dynamic nor causal. It says nothing about the time course of adaptation to a particular procedure, or about the “decision rules” that animals may follow once they have adapted; nor does it account for metastable adaptations or the process of acquisition. For example, the time lag between the delivery of a reinforcer and the behavior on which it is contingent is vital, yet such delays are not adequately represented in a molar feedback function. At long VI values, the molar feedback function for a variable-interval schedule that delivers food 10 sec after the effective response is almost the same as the function for one that delivers the food at once, yet animals may not learn to respond when there is a delay and, having learned, will respond much more slowly. Observed behavior, whether optimal or not, is always the result of particular, and imperfectly understood, mechanisms for detecting correlations, computing time averages, and assessing stimulus covariation. No matter how subtle such mechanisms may be, there will always be some combination of circumstances for which they will be inappropriate. At some point, probably when they become excessively complex, economic models are no longer useful and must give way to a detailed understanding of the particular dynamic processes that allow reinforcers to exert their effects.

I take up these processes in later chapters, but before leaving this topic some other departures from simple optimal behavior should be mentioned.

Intermittent, brief, predictable feedings often induce food-related activities such as peck-

ing, and even lever pressing. These effects are variously termed *autoshaping* or *superstitious* behavior, depending upon procedural details. They represent departures from a simple economic analysis, in the sense that these activities behave as complements to food, rather than substitutes, since they covary with the frequency of food delivery. They can be accommodated within regulatory models only by assuming that the free levels of some activities are dependent on the levels of others, or by changing other properties of the objective function.

There are several other kinds of activity that are induced or facilitated by intermittent reinforcer delivery, most notably *schedule-induced polydipsia* (excessive drinking) by rats on food schedules, and schedule-induced attack (on a conspecific or model thereof) by pigeons and other animals. These activities all behave like complements, rather than substitutes, for the reinforcer.

Animals can be induced to behave in maladaptive ways by special training procedures. For example, dogs given recent exposure to severe, inescapable electric shocks are unable without special help to learn a simple response to avoid shock (so-called *learned helplessness*). Restrained squirrel monkeys, trained to avoid shock by pressing a lever, can be brought to actually produce intermittent shocks. These violations of optimality theory depend upon limitations in the learning mechanism. As I have already pointed out, every adaptive mechanism has its limitations. They tend to be more elusive in the case of food reinforcement, perhaps because food is a natural reinforcer. Electric shock is obviously remote from animals' normal experience, and they are accordingly less flexible in dealing with it. Avoidance contingencies ("do this to prevent the future occurrence of that") are also intrinsically more difficult to detect than positive ones ("do this to obtain that"). These issues recur when I discuss the process of acquisition in later chapters.

SUMMARY

It used to be customary to discuss *motivation and learning* in different books, or at least in separate sections of the same book. This separation is not possible any longer. Regulatory processes are involved in the molar rates even of activities like wheel running and lever pressing, not usually thought of as "motivated" in the same sense as "consummatory" activities like eating and sexual activity. These regulatory processes can be described by relatively simple economic models that embody three major assumptions: that activities are constrained by reinforcement schedules and the availability of time; that the marginal value of an activity diminishes with its rate of occurrence; and that animals allocate time to activities in such a way as to maximize total value.

I have taken the *rate* of an activity as the dimension that is subject to diminishing marginal value, because lever-press or key-peck rate is usually the only thing the animal can vary to affect his eating rate, but this is not essential to the optimality approach. The fundamental idea is that whatever the animal can do that has consequences for him is subject to diminishing returns. For example, in some experiments by Collier and his associates (e.g., Figure 6.4), rats had two ways to affect eating rate: Frequency of access to the feeder was determined by the rate at which the rat pressed a lever on a ratio schedule; and duration of each meal was determined by how long the animal continued to feed once the feeder operated. Granted that both these activities are subject to diminishing returns, one might expect the animal to vary both, according to their relative costs, so as to maintain eating rate relatively constant.

This point, that an optimality analysis must consider *everything* that is under the animal's control, might seem trivial except that its neglect can lead to absurd conclusions. For example, in an experiment by Marwine and Collier,²⁹ one thirsty rat drank an average of 14.7 ml of water per bout on a FR 250 schedule. Yet on FR 10 the same rat drank only about .4 ml per bout. By so doing, the animal failed to get as much water as it physically could (given the capacity demonstrated on FR 250) and also failed to get as much as its free baseline level. This result can be taken as contrary to an optimality analysis only if the animal is assumed to be indifferent to drinking-bout length. Once we grant that the marginal cost of a drinking bout must increase with

its length, the longer bouts on the higher ratio schedules make perfect sense.

Optimality analysis explains reinforcement in terms of relative costs: Activities with low costs-of-deviation from their optimal values are allowed to vary so that activities with high costs-of-deviation can remain approximately constant. Many learning constraints can be explained by CoD differences, without the necessity to invoke special links between activities. Some phenomena, such as schedule-induced behaviors, and easily displaced but unreinforceable activities such as grooming, cannot be explained in this way. No doubt more comprehensive optimality models could be found to accommodate these effects. But since the main virtue of economic models is that they explain numerous facts by a small number of assumptions, complexity diminishes their usefulness. At some point other approaches may prove better.

This chapter has intentionally dealt with a restricted range of situations: no more than three activities, and only one reinforcement schedule. The next two chapters extend the optimality approach to situations with more than one reinforcement schedule and look at applications to natural foraging behavior.

NOTES

1. *Regulation.* Motivational theorists spend much time arguing over the meaning of the term *regulation*, although it's not really clear why. A regulated variable is simply one that is subject to negative feedback. Thus, suppose variable y is a positive function of another variable x , $y = f(x)$, as body weight is a positive function of average eating rate. Suppose y is altered in some way — perhaps by a forced change in x (food deprivation), but perhaps by some other means that affects y directly (a metabolic drug, for example). When the perturbing influence is removed (ad-lib feeding is restored, the drug injections cease), if y is regulated with respect to x , the change in y will produce a change in x in the opposite direction — dy/dx is negative. If body weight has fallen, eating rate will rise; if body weight has risen, eating rate will fall.

The concept of *regulation* is precisely the same as the idea of *stable equilibrium*, described in Chapter 5.

2. For experimental results showing behavioral competition see reports by Dunham (1971), Henton and Iversen (1978), Hinson and Staddon (1978), Premack and Premack (1963), Staddon and Ayres (1975). For a theoretical discussion see Staddon (1977b), and Chapter 10. For theoretical discussions of regulation in the sense of this chapter see McFarland and Houston (1981), Staddon (1979a,b), Timberlake and Allison (1974), Timberlake (1980), and an earlier paper by Premack (1965).

3. See, for example, Staddon and Ayres (1975). Other experiments showing competition are Dunham (1977), Reid and Dale (1983), and Reid and Staddon (1982).

4. *Rate and time measures.* The time constraint for two recurrent activities can be stated formally as follows: Let each instance of activity one take an amount of time t_1 , and similarly for activity two. If all the time is taken up with one or other of these activities, then, in an experimental session of unit length, $N_1 t_1 + N_2 t_2 = 1$, where N_1 and N_2 are the number of instances of activities one and two. For m mutually exclusive and exhaustive activities, the constraint is obviously

$$\sum_{i=1}^m N_i t_i = 1. \quad (\text{N7.1})$$

This equation assumes that the time taken up by each instance of a recurrent activity is constant. This is true for many commonly studied activities such as licking, wheel running and lever pressing by rats, and key pecking by pigeons. It breaks down if the activity has different topographies

at different rates, as in free locomotion, where animals show different gaits — walking, trotting, cantering, galloping — at different speeds. Some procedures (e.g., ratio schedules) induce pigeons to peck at very high rates, and then the topography of the response differs from pecking at lower rates: Each peck is briefer and less forceful, and movement elements present under normal conditions drop out. Time taken up by an activity is still proportional to its rate, but the constant of proportionality changes with changes in response topography.

If the response topography (“gait”) is constant, then time taken up by a response will be proportional to the rate of the response.

5. Activities generally occur in relatively brief bouts, so that even a rare activity (in the sense that it occupies little of the total time) may nevertheless occur quite frequently. A high-probability activity need not be the first to occur, nor does it have to be “exhausted” before lower-probability activities can occur. Activities are generally intercalated; higher-probability activities are simply “on” for a larger fraction of the total time. If activities were not intercalated in this way, the concept of *average rates*, and the notion of a free-behavior point defined in terms of them, would have no meaning. The dynamics of behavioral sequences are discussed at greater length in Chapter 11.

6. This terminology, and the major theoretical arguments in the rest of the chapter, are presented more formally in Staddon (1979a).

7. There are some theoretical problems associated with the yoked-control procedure, because no two animals are identical (see Church, 1964). Even within-animal yoking is open to the standard objections to within-animal experiments in general (see Chapter 4), but as a practical matter yoking is often a useful way to tackle contingency problems (Dunham, 1977).

8. *Income and substitution effects.* The effect of a schedule constraint can always be partitioned into a portion due to restriction and the remainder, which reflects the contingency alone. The restriction component is the change in the rate of the instrumental response, x , that would have taken place had the level of the contingent response, $R(x)$, been artificially held to the value attained under the contingency. The true contingent effect is any increase above that level. If we consider available time as analogous to the *income* of a consumer, and the ratio value as equivalent to the *price* of a commodity, then this distinction between restriction and contingent effects is essentially the same as the economist’s distinction between *income* and *substitution* effects. Other relations between economic concepts and optimal time allocation are taken up later.

9. This argument assumes a bidirectional contingency; that is, x units of B are required for access to y units of A , but these y units must then be used up before access to B is restored: For example, the animal must press a lever five times to get a pellet of food, but then the pellet must be eaten before he again has access to the lever. Under usual conditions only one of these conditions will matter: A starving animal needs no urging to eat the small portions of food delivered by a schedule of food reinforcement, but it will not press the lever much unless lever pressing is required for access to food. In most experiments, therefore, only one half the contingency need be enforced.

10. *Value, reinforcement, and Darwinian fitness.* The term *value* is used here in several senses. To an economist of the old school it meant subjective utility and connoted some kind of pleasurable response in terms of which the value of all “goods” was assessed. It proved impossible to arrive at a satisfactory cardinal scale of utility in this sense. More recently, economists seem to have agreed (one can never be sure with economists) that although magnitudes cannot properly

by assigned to utilities, things can at least be rank ordered in terms of value: We can agree that half a loaf is better than no bread, even if not *how much* better. Much ingenuity has gone into showing that most of the predictions of classical utility theory are still derivable from this more parsimonious assumption. Useful reviews of classical utility theory are Hicks (1956), Friedman and Savage (1948), and Stigler (1950).

The modern theory of *revealed preference* treats value as something that can be inferred from (is “revealed” by) the subject’s pattern of choices. In other words, value is not measured directly but inferred from behavior, under the general assumption that behavior is guided by a consistent, fixed *preference structure*. I have followed this kind of reasoning in the text, where it led to the notion of a positively accelerated cost function. The approach is scientifically useful when it is possible to derive the preference structure from a limited set of behaviors, and then use it to explain a much larger set. Samuelson (e.g., 1965) is responsible for the concept of revealed preference in economics.

For the biologist, utility can only mean “contribution to Darwinian fitness,” but this elegant notion conceals much. Neither organism nor biologist can assess fitness directly. The organism can only evaluate outcomes in terms of some intermediate *currency* selected by past evolution as a good predictor of fitness. For example, in choosing among food grains, a bird cannot know directly which type will most favor future offspring. But if size is correlated with nutritional value, then an inbuilt tendency to select the larger grains will suffice. For such a bird, utility equals grain size. The indirect nature of the process is revealed by allowing the animal to choose between artificial grains, made equal in food value but of different sizes. If the animal retains its size preference, we can be sure of the signal value of size. If he does not, or if he gradually loses his preference, we must conclude that size is not the only signal available — perhaps the animal can learn to correlate the appearance of grain with its nutritional value (as chicks seem to be able to do under some conditions — see Chapter 14). If so, then nutritional value is the fitness signal — *prima facie* a more reliable one than size, perhaps, but indirect just the same.

The things we identify as reinforcers (more accurately: The preference structure discussed later in the chapter) are just the signals that evolution has identified as the most useful proxies for fitness.

In terms of the regulatory analysis, deviations of different activities from their free-behavior values should be scaled in terms of their probable importance for future reproduction, as estimated by past natural selection. For example, a given reduction in rate of feeding will generally be more of a threat to reproductive success than the same reduction in grooming; this difference might well be reflected in a greater weighted cost for eating-rate deviations.

For further discussion of the strengths and weaknesses of optimality theory applied to evolutionary problems see Maynard Smith (1978) and several chapters in the collection edited by Krebs and Davies (1978).

11. This conclusion can be stated somewhat more formally as follows: Diminishing marginal value implies that dV/dx is a decreasing function of x or, equivalently, that d^2V/dx^2 is negative over the range of interest. By our intuitive assumption, dC/du is an increasing function of u , where $u = x_0 - x$. By the chain rule, $dC/du = -dC/dx$; therefore, $-dC/dx$ is an increasing function of $x_0 - x$, hence, for $x < x_0$, a decreasing function of x . Since cost is just the negative of value (if you give a man \$10, the cost is added in as - 10 to your net worth), dV/dx is, therefore, a decreasing function of x , which corresponds to diminishing marginal value.

For example, let $C(x) = A(x_0 - x)^2$. Since $V(x) = -C(x)$, $dV/dx = 2A(x_0 - x)$, a decreasing function of x for $x < x_0$. These relations are illustrated graphically in Figure 7.19.

12. *Bliss point and characteristic space.* There are at least two possible interpretations of the parameters x_0 , R_0 , and y_0 in the quadratic cost equation. One is the interpretation assumed in the text, that these are just the free levels of the activities when all three are available; the point (x_0, R_0, y_0) then corresponds to the economists' *bliss point*, the ideal where all wants are satisfied. Unfortunately like all other utopian notions, this cannot be true. There can't be anything special about this point because the addition of a new activity, or replacement of one by another, will generally lead to changes in the free levels of the original activities. The free-behavior parameter must be a property of an activity in a context of other activities, not a property of each activity alone. The effects of changing the behavioral repertoire cannot be predicted from the simple minimum-distance model. An alternative interpretation, therefore, is that x_0 , say, is just a property of behavior X that can be estimated from the data. It may correspond to

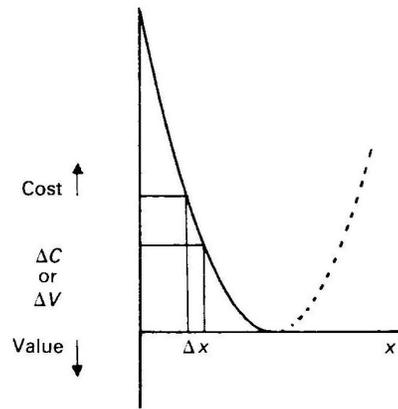


Figure 7.19. The equivalence between a positively accelerated cost function and a negatively accelerated value (utility) function.

the free level of X in the absence of any competing activity, or under paired-baseline conditions, but it need not — the value of x_0 for some activities may be negative or in excess of the total time available, for example. In this more general view, the true “bliss point” may not be physically realizable.

One way of dealing with the problem of predicting the effects of changes in the set of activities is to postulate a small set of fundamental motivational *characteristics*, V_i — a more sophisticated version of the old notion of primary drives or motives. Any given activity is

then assumed to share to some degree in one or more of these characteristics, perhaps in a simple linear way, for example,

$$X_j = AV_1 + BV_2 + \dots + KV_k,$$

where $A-K$ are weighting constants and V_1-V_k are the k characteristics that make up activity X_j . The animal is then presumed to adjust the levels of the activities available to him so as to minimize deviations from a bliss point in this *characteristic space*. This approach gains generality at considerable cost in complexity, which is why I have elected not to discuss it in the text. The notion of a characteristic space has been explored in economics by Lancaster (1968) and in behavioral psychology by Lea (1981). The ethologists McFarland and Houston (1981) have independently followed a very similar line.

13. See, for example, Chiang (1974); Glaister (1972); and Baumol (1977). The quadratic model is treated explicitly in Staddon (1979a).

14. Points in plane XYR in Figure 7.8 are in triangular coordinates, which should be read as shown in Figure 7.20. Thus, points on any line at right angles to line RC represent constant lev-

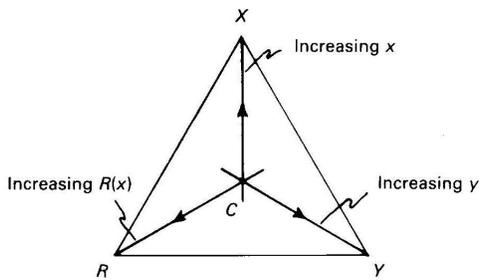


Figure 7.20. Triangular coordinates.

els of $R(x)$.

15. Most notably the *choice axiom* for independent alternatives proposed by Luce (1959, 1977). The axiom fails here because the alternatives, although independent in terms of the cost function, are not independent in terms of time taken up (the time-allocation constraint). Some results from human choice experiments are consistent with Luce's axiom, but data from behavior-allocation experiments of the type here discussed have not supported it. Figure 7.8 emphasizes the importance of thinking very carefully about what is meant by "independence" before applying theories that require it.

16. This section gives only the briefest survey of that portion of economic theory most relevant to the analysis of operant behavior. For more comprehensive treatment, see standard economics texts such as Walsh (1970), Baumol (1977), or the most recent edition of Samuelson's popular book. The classical treatment is by Marshall (e.g., 1925).

For recent discussions of the relations between economic theory and behavioral psychology see the volume edited by Staddon (1980a), especially the chapters by Rachlin and Kagel et al. Papers by Rachlin, Green, Kagel, and Battalio (1976), and Rachlin and Burkhard (1978) are recent efforts to bring together economics and behavioral psychology.

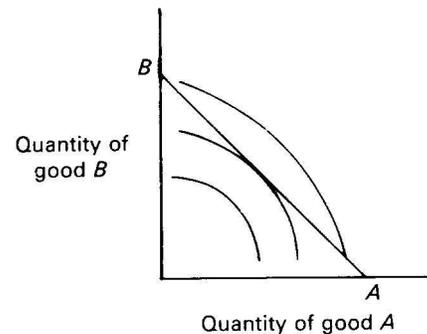


Figure 7.21. Indifference contours and the budget line.

17. These are the favored substances in a series of studies by Rachlin, Green, Kagel, and Battalio. For rats they are imperfect substitutes.

18. Technically the marginal rate of substitution of two commodities is the negative of the slope of the indifference curve at that point. A related but more useful concept is the notion of *elasticity*, which is defined as the ratio of the slope of the indifference curve to the ratio of its absolute values: $\epsilon_{yx} = (dy/dx)/(y/x)$. For cost functions of the general form $C = (x_1, \dots, x_k) = \sum a_i (x_{0i} - x)^m$, ($m = 2$ for the minimum-distance model), elasticity of substitution is always unity.

19. This third logical possibility, a curve concave to the origin, is permissible within the present analysis, but cannot occur in ordinary economic theory. The reason is illustrated in Figure 7.21, which shows the income constraint (*budget line*, AB) representing the boundary limiting the consumer's purchases of bundles composed of goods A and B . For example, if A costs \$1 per unit and B \$2 per unit, and income is \$100, then the equation of the budget line is $A + 2B = 100$, where A and B are quantities of A and B . Either preference increases away from the origin (as I have been assuming) or toward the origin. If away from the origin, then points A and B are both preferable to point T , the point of tangency between the budget line and the highest indifference contour (since indifference curves are contours, the assumption that preference increases away from the origin places the origin at the bottom of a utility "bowl," with points A and B both higher than point T). Hence, indifference curves of this form dictate that the consumer shall always purchase one of the two goods exclusively. Conversely, if preference increases toward the origin (the origin is the peak of a utility "hill"), the consumer can maximize utility by purchasing nothing at all.

If we retain the convention that preference increases away from the origin, then indiffer-

ence curves like those in Figure 7.21 will always dictate exclusive choice. As a practical economic matter, this never occurs: Consumers do not spend all their income on one thing. In the more restricted conditions of behavioral experiments this outcome is still improbable, but is not totally excluded.

There is another important difference between the economics of the human consumer and the economics of behavioral allocation. Money income can be saved or spent, but time, the currency in behavioral allocation, is always expended. Consequently, we can expect that the economic *axiom of non-satiation* — that more is always better, or at least not worse — will regularly be violated, because more of any activity always entails less of others. This means that indifference contours for behavioral allocation will generally be closed curves, and the constraint imposed by a reinforcement schedule (analogous to the budget line) will always restrict behavior to a region away from the region of highest value. For example, the indifference curves for electric shock vs. lever pressing will resemble those of Figure 7.21, with the region of highest value close to the origin. The constraint imposed by an avoidance schedule is of the same general form as the budget line in the figure, but it restricts behavior to the region away from (rather than including) the origin. Optimal performance is, as usual, the point of tangency between this constraint line and the indifference curve closest to the region of maximum value, that is, to the origin.

20. Models of this form are also minimum-distance models, but the space is Euclidean only when $m = 2$.

21. The logical relations between utility and revealed preference are one of the deeper areas of mathematical economics. Questions considered are, for example, under what conditions does a utility function exist? Under what conditions is it unique? Under what conditions are preferences transitive (A preferred to B and B to C implies that A is preferred to C)? The classical treatment of these issues is by Von Neuman and Morgenstern (1947). More recent treatments are Luce and Raiffa (1957), Hicks (1956), Henderson and Quandt (1971), and Intriligator (1971).

22. The alert reader will perhaps have noticed that in a sense only two assumptions, competition and diminishing marginal rate of substitution, are essential for these predictions. After all, many forms of indifference contours are consistent with the diminishing marginal substitution assumption, and by judicious choice among these, quite close matches to the empirical data can be obtained (see Rachlin & Burkhard, 1978). Response independence is a necessary assumption only if we wish to rely on minimum-distance or some comparable additive rule to generate the preference structure.

23. For systematic data on tandem and conjunctive schedules, see Zeiler and Buchman (1979). Susan Motheral, in my laboratory, has obtained systematic data on a schedule related to the interlocking schedule discussed in Chapter 5. Tandem schedules require a fixed number of responses (the ratio requirement) after an interval requirement has been met. An approximation to the feedback function for these schedules can be derived by extending the argument used previously to obtain the feedback function for interval schedules. Thus, on an interval schedule, given random responding, the expected time between occurrences of reinforcement is $D(x) = 1/a + 1/x$, where $1/a$ is the interval value and x is the average response rate. On a tandem schedule, the quantity m/c , where m is the ratio requirement, must be added to this, yielding the feedback function $R(x) = ax/[x + a(I + m)]$. In accordance with the data of Zeiler and Buchman (1979), this function predicts a bitonic response function; moreover, the minimum-distance model also predicts that response rate (at a given interval value) will be increased by the addition of a moderate ratio requirement, which is also in accord with their results.

A conjunctive schedule imposes a ratio requirement that can be met before the interval requirement has elapsed. It can be considered as a tandem schedule with a variable-ratio requirement smaller than the nominal value (because the animal makes contact with the requirement only when it fails to make the necessary number of responses during a particular interval). This assumption leads one to expect that the peak of the bitonic response function for this schedule should be at a lower ratio requirement than the maximum for the comparable tandem schedule, as Zeiler and Buchman report.

24. Response functions (not shown) were fitted by nonlinear regression, on the assumption that the free level of bar pressing was zero.

25. Several edited books on the general topic of “constraints on learning” have been published in recent years, as well as a number of influential review papers. See, for example, Hinde and Hinde (1973), Seligman and Hager (1972), and papers by Bolles (1970), Breland and Breland (1961), Shettleworth (1972, 1975), and Staddon and Simmelhag (1971).

26. I assume here that all activities are measured in time units. And for Equation 7.12 to be used, it is necessary that in each pairwise comparison the cost of “other” activities be negligible.

27. Hogan, Kleist, and Hutchings (1970). See Hogan and Roper (1978) for a recent review of many similar results.

28. Note that this analysis of reinforcement effectiveness is not the same as the so-called weak reinforcement hypothesis (discussed by Shettleworth, 1972, and others), which states that all reinforcers can be arranged on a continuum of strength, so that a “weak” reinforcer, like access to nest material, will behave just like food for a slightly hungry animal. The minimum-distance model characterizes each activity (reinforcer) by two parameters, a cost and a free-behavior level, not one as the strength view implies. If the equal-form assumption (Equation 7.3) is also relaxed, activities might differ in many other ways. There is no reason why the indifference contours relating one reinforcer and an instrumental response should be of the same form as those relating another reinforcer to the same response.

29. Marwine and Collier (1979). These figures, and the misplaced criticism of optimality theory, are due to Allison (1981).