JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1985, 43, 407-417 NUMBER 3 (MAY)

INCENTIVE THEORY: IV. MAGNITUDE OF REWARD

PETER R. KILLEEN

ARIZONA STATE UNIVERSITY

Incentive theory is successfully applied to data from experiments in which the amount of food reward is varied. This is accomplished by assuming that incentive value is a negatively accelerated function of reward duration. The interaction of the magnitude of a reward with its delay is confirmed, and the causes and implications of this interaction are discussed.

Key words: incentive theory, amount of reward, delay of reward, concurrent chained schedules, concurrent schedules, mathematical models, marginal utility

Incentive theory has been successful in describing behavior under the control of concurrent chained schedules of reinforcement (Killeen, 1982b). But problems arise for that account in situations in which the amount of reward is varied. The critical experiment is that of Snyderman (1983), who studied the behavior of pigeons under concurrent chained schedules in which the amount of reward was either 2 s or 6 s of eating time, available at the ends of terminal links of different fixed delays. He concluded that my model for choice did not handle some of the basic qualitative aspects of the data. I demonstrate here that his analysis, although accurate, did not take into account a critical factor: The incentive value of a reward may be a concave function of its duration. If we assume the value of a 6-s reward to be three times that of a 2-s reward—the simplest assumption, and one consistent with my previous presentation of the model—then my predictions are seriously in error. If, however, we assume that there is a decreasing marginal utility to extended hopper durations, then the predictions align with the data.

First I review the model, and then offer a transformation to convert duration to value.

THE BASIC MODEL

The model assumes two effects for incentives, a directive effect and an instigating effect. The directive effect itself consists of two parts, one due to the direct action of the incentive on behavior ("primary reinforcement"), and one due to the effect mediated by the terminal-link cues ("conditioned reinforcement"):

$$S_d = e^{-qt} + 1/t, \qquad t > 0, \qquad (1)$$

where S_d is the directive strength, q is the slope of the delay of reinforcement gradient, and t measures the delay. Because the reward is delayed from the response that brings it about, its effect is likely to be blocked by intervening stimuli (or responses; Williams, 1982). If the probability of blocking is constant over time, an exponential decay function results; that is the first term in Equation 1.

The terminal-link cues are contiguous with the response, but their strength is weakened because they signal a state in which the density of incentives is less than maximal. I have tried a number of forms for the function relating the strength of stimuli to the delays that they signal, and have found the most consistently accurate to be the reciprocal of the delay. Inasmuch as there is always some latency in getting to food, a minimal value around 0.5 s will both represent that latency and keep S_d finite.

The second effect of incentives is to incite responses: Animals that are fed more often become aroused and emit more responses. This arousal can come under stimulus control; animals respond faster when placed into an experimental chamber that has previously provided an increased density of incentives (Killeen, 1979). It follows that the stimuli signaling concurrent schedules may differentially produce

---

I thank Michael Davison and J. A. Nevin for their help on this article. Send reprint requests to Peter R. Killeen, Department of Psychology, Arizona State University, Tempe, Arizona 85287.

arousal. Except where very low rates are involved, the level of arousal is proportional to the rate of incitement (R), measured as the reciprocal of the sum of the time in the initial link plus the time in the terminal link, for each schedule. The overall strength of a schedule of incentives is then:

$$S = RS_d. \qquad (2)$$

Magnitude and Value

Incentive theory has not been extensively developed to deal with variations in duration, amount, or quality of reward. The basic theory (Killeen, 1982a, Equation 4; McDowell and Kessel, 1979) suggests that increasing the duration of a reward should increase its value according to a concave function:

$$v = c(1 - e^{-\lambda d}), \qquad (3)$$

where v is the "value" (or, equivalently, the "reinforcement strength") of a single incentive of duration d. The parameter c accommodates differences in the type or quality of reward. Since that is not varied in the studies reviewed here, c is assigned a value of 1.0 throughout. The value of lambda determines the curvilinearity of this "utility function," with very small values linearizing the function, and with large values bringing the function to its maximum very quickly, so that incentives of long duration are little better than those of moderate duration. According to Equation 3, value does not grow indefinitely with duration (as would be the case if it were a power function), but rather approaches an asymptote of c.

Value and the Delay of Reinforcement Gradient

Variation in the value of an incentive affects behavior in two ways. As assumed above, it affects the arousal of the animal. This is taken into account by multiplying the right side of Equation 2 by v. I had first expected that this would be the only way that incentive value would enter the predictions. But as Navarick and Fantino (1976) noted, amount and delay of reward do not interact in a simple multiplicative fashion; this is also the case for any transformation of amount, such as Equation 3. Figure 7 of Killeen (1982b, based on data from Green, 1969) shows that preference for 9 s versus 3 s of grain increases from 55% to 90% as the length of equal terminal links increases from 5 to 60 s. A similar increase is seen in the data of Snyderman for the equal delay conditions. We must inevitably conclude that large amounts of food become relatively more effective at long delays. The only way to accommodate this theoretically is to have value affect the delay of reinforcement gradient, which may be done by dividing q by v. (In Killeen, 1982b, I divided by d; this addressed the interaction problem but was not a precise solution because it did not permit value to be a concave function of delay.)

I address now the question of the formulation's adequacy in the face of Snyderman's (1983) data. I evaluated Equation 3 for different values of lambda and found that the one providing the best fit to the data is $\lambda = 0.77/\text{s}$. This determination was made in the context of the model:

$$S = vR(e^{-qt/v} + v/t), \qquad t > 0, \qquad (4)$$

where v enters as a source of "bias," multiplying the rate of incitement; as a modifier of the delay gradient, dividing the parameter q; and as a factor in the strength of conditioned reinforcement. This equation is evaluated for each of the schedules, and its relative value predicts the relative rate of responding. Normally an additional "key" bias parameter (b) is used, but that adds nothing to the goodness of fit to these data. The value of q, 0.04/s, determined by an iterative least-squares procedure, is much lower than typically found (~0.13/s). Part of the difference is due to the inclusion of v in the exponential term of Equation 4. If value had been considered in the earlier treatments, the modal value for q would have been about 0.09/s. Part of the difference may have to do with Snyderman's use of different houselight colors as terminal-link cues. Because, for simplicity, I have chosen not to include a parameter in the model to reflect differences in the potency of different types of conditioned reinforcers, differences there will affect the value of q. Figure 1 shows that this treatment accommodates the data reasonably well, accounting for 93% of its variance, with the average deviation of data points from predictions being 4%.

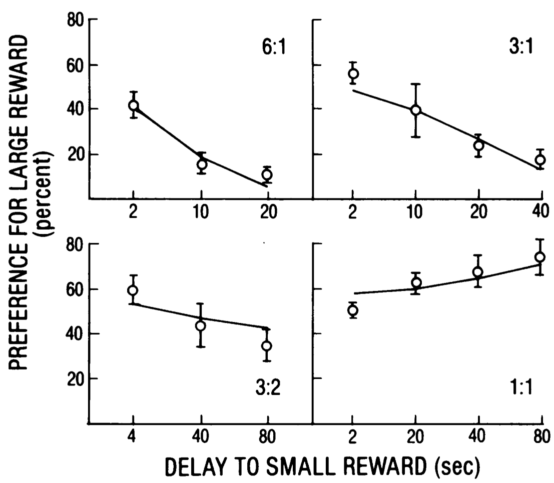

Fig. 1. Rate of responding for a 6-s reward relative to the rate for a 2-s reward as a function of the reward delay. The data are Snyderman's (1983). The delay to the large reward was greater than the delay to the small reward by the ratio noted in each panel. The circles mark the mean across subjects, and the vertical bars its standard error. The solid lines are the predictions of the model with values of 0.77/s for lambda (assuming a 1-s latency to begin eating) and 0.04/s for q. The difference of these slopes from zero rules out the description of these data by the "generalized matching law" (Baum, 1974), since that must treat all data within each panel as the same, by virtue of their same ratios of delays.

This treatment may be compared with Fantino's (1969) "distance to reward" model. When that is modified by using Equation 3 to permit the value of reward to be a nonlinear function of its duration, it accounts for 85% of the data variance. It accomplishes this with only one parameter, but shows some systematic error, as noted by Snyderman. It is interesting that the ratio of values of v required by Fantino's model, 1.4/1, is quite similar to that given by my model, .99/.77 = 1.3/1. By either account, 6 s of eating time is worth substantially less than twice as much as 2 s of eating time.

Reanalysis of the data from Green and Snyderman (1980) using Equation 3 also improves the accuracy of the basic model, increasing the percentage of variance accounted for from 86% to 89%, and decreasing the average deviation from 6.5% to 5%. The value of lambda required by these data is 0.29/s, substantially less than for the Snyderman data. However, in Snyderman's experiment the duration of reward was controlled by a photocell in the hopper, so that the values cited for the durations are good estimates of the time spent eating. That was not the case for the Green and Snyderman experiment. If we assume that it took the pigeons 1 s to enter the hopper and begin eating, the parameters converge to exactly the same values as those used for the Snyderman data: 0.04/s for q and 0.77/s for lambda.

It is unfortunate that hopper-approach latency must enter as a separate parameter. (I have argued elsewhere [Killeen, 1972, 1985] for specifying all independent variables as closely as possible to the values experienced by the subjects, both to increase accuracy in specifying the controlling variables and to decrease the complexity of the models.) This additional parameter is not necessary if eating is timed with a photocell. But the major effect of imprecisions in our specification of reward duration will be to change the optimal value of lambda (and, for very short terminal links, to change the value of q); goodness of fit will not be greatly affected. Therefore, in studies where access time (rather than actual eating time) is controlled and reported, I will subtract 0.5 s from the stated durations. This is a conservative correction, for I suspect that a value between 0.5 and 1.0 s is more typical. However, when scheduled durations are very short, small differences in latency of hopper approach will result in differing incentive values, which should in those situations shape rapid approach.

OTHER STUDIES

Neuringer (1967) gave pigeons 2-s access to grain in one terminal-link FI 5-s schedule, and equal or longer durations in the other terminal-link FI 5-s schedule. Delivery of food was probabilistic, and when it was not delivered the animals received a brief blackout and then were returned to the initial link. In the initial link each response on a key constituted a choice, and moved the animal into the terminal link. We may apply the model by taking R to be proportional to the number of reinforcers received from each terminal link. The predicted (and obtained) percentage of choices

Table 1

Obtained and Predicted Choice Proportions

Source	Amount Constant	Amount Variable	Choice Proportion Obtained	Choice Proportion Predicted
Neuringer (1967)	2.0	2.0	.49	.49
		2.25	.54	.57
		2.5	.63	.61
		3.0	.69	.68
		4.0	.74	.74
		6.0	.82	.82
		10.0	.86	.87
Menlove et al. (1979)	8 or 0	2	.33	.32
		3	.48	.51
		4	.61	.59
Young (1981)	10 or 0	1	.02	.00
		2	.11	.03
		3	.13	.12
		4	.24	.24
		5	.37	.35
		6	.37	.45
		7	.52	.53
		8	.58	.58
		9	.62	.63
		10	.72	.67

for various durations are given in Table 1. Those predictions required values of 0.10/s for q and 0.54/s for lambda. Here, as in other cases, the accuracy of predictions would fall substantially if we forced value to be proportional to duration (the average deviation would increase from 1% to 6%).

Menlove, Inden, and Madden (1979) gave pigeons 2-s, 3-s, or 4-s access to grain at the end of one FI 15-s terminal link, and 8-s or 0-s access to grain at the end of the other FI 15-s terminal link. This latter alternative provided the large amount half the time, and a brief hopper flash the other half of the time. If we assume the overall strength of this terminal link to be just half of that provided by one with a consistent 8-s access to grain, we may apply the model. With values of 0.08/s for q and 0.54/s for lambda, the predicted and obtained percentage preferences for the fixed amounts are given in Table 1.

Young (1981) provided the only evidence for a linear relation between the amount of reward and its incentive value. He gave pigeons access to 0 pellets or 10 pellets with equal probability at the end of one terminal link. At the end of the other terminal link he gave them a certain number of pellets on every trial. In the series with the most data, the certain pellets equaled 10 in the first condition, 9 in the second, and on down to 1 in the last condition. Each response in the initial link carried the animal into one of the terminal links, which were 4 s of blackout for both alternatives (during which time the pellets were being delivered into the hopper). Sixteen of every 17 trials were guidance trials, with only one alternative being offered; the 17th was a choice trial from which preferences were estimated. Because there were no differential terminal-link stimuli, the conditioned reinforcement term drops out of Equation 4. With values of 0.10/s for q and 0.03/s for lambda, the model accounts for 97% of the data variance. See Table 1 for the exact predictions. (Duration of incentive was estimated as 0.5 s per pellet, based on comments of Young and of Schneider [see discussion that follows].) With so small a value for lambda, the utility function is essentially linear. If we replace Equation 3 with $v = 0.75d$, the model continues to account for 97% of the data variance; in both cases the average deviation of data from predictions is 3%.

It is not clear why the relation between the magnitude of a reward and its value is linear in Young's (1981) study. There are a number of features of this experiment that make it quite different from the others considered here—the large number of forced trials, the use of a strictly descending series of values, the delay in blackout while the pellets were being delivered, and the use of number of pellets rather than duration of access to grain.

Independent of the nature of the utility function, Young's study demonstrates an important point about the interaction of magnitude and delay gradients. By itself, Equation 3 suggests decreasing (or at least, nonincreasing) marginal utility for larger magnitudes. This means that 10 units of reward can be worth no more than twice 5 units. But Young's data showed (and the model predicted) a substantial preference (63%) for a 50% chance of 10 units reward over a sure 5 units. How can this be? If the above options were delivered immediately, the model would indeed predict a preference for the small, consistent amount, or at least indifference. But remember that increasing the magnitude of reward flattens the delay-of-reinforcement gradient. As the delay to the options increases, the large reward gains an advantage. This crossover may help us understand the dynamics of gambling in general; a large but uncertain reward may be preferred to a small immediate one to an extent greater than any atemporal analyses, such as expected utility theory, can predict.

Experiments by Ito and Asaki (1982) suggest that the relation between reward magnitude and strength may also be concave for rats, as it typically is for pigeons. These investigators kept the amounts of incentive fixed at either 3 pellets or 1 pellet, and studied the changes in rats' preference as they varied the delay in the terminal links of concurrent chained schedules. They also noted that their data could not be accounted for by a simple multiplicative interaction of amount and delay. We may apply Equation 4 to their data, specifically to the conditions of Experiment 2, where both terminal links were of the same duration: 5, 10, 20, or 40 s. They found average preferences for the large amount to increase with delay—66, 77, 79, and 87%. The model, as described above, predicts 69, 73, 80, and 86%, using parameter values of 0.062/s for q and 0.75/pellet for lambda. (Incentive theory has a problem with the data from the first experiment, in which the parameter values and the predictions are the same as for the second experiment, but where the data from the 5-5 condition were less extreme [60%] than either the prediction or the replicated data in Experiment 2. I believe that this is because in the first experiment, but not the second, the investigators had first established indifference to equal amounts of food, and the effects of this manipulation carried through to the second [5-5] condition. The other conditions are quite close to both the predictions and to the Experiment 2 replications.)

Davenport (1962) also found an interaction of incentive magnitude and delay when studying the runway performance of rats. Davenport used 4/2, 8/2, and 16/2 pellets and calculated the delays necessary to cause the animals to be indifferent in their choices. Given only three data points, the model has no trouble in accommodating them perfectly (predicting isopreference delays of 8, 20, and 26 s, for $q = 0.024$/s and lambda = 0.26/pellets). The important part is that the model predicts, and the investigator found, both a concave utility function and an interaction; this is consistent with the other studies and has implications for models of self-control.

Concurrent Schedules

The above studies involved an interaction of incentive magnitude with its delay. Does the model accommodate simpler situations in which the delay is minimal, such as basic concurrent schedules? It is possible to simplify the model for such situations: There are no differential terminal-link stimuli, and so that term may be removed from the equation for S_d. The assumed 0.5 s hopper entry time is taken as the value of the terminal link. It permits only a small range of values for the direct control by primary reinforcement, so that variations in the value of q have little effect on the goodness of fit. Therefore it may be fixed at a representative value, such as 0.09/s. The 0.5 s latency is also subtracted from the stipulated durations of the feedings.

Brownstein (1971) measured the amount of time pigeons spent in the presence of one of two different cue lights, each correlated with intermittent feedings programmed by a response-independent VI 90-s schedule. The pigeons could switch between cues by responses on a changeover key. The durations of feedings provided by each schedule, and the predicted (and obtained) percentage of time spent in the presence of each, are: 3/3: 52 (53); 2/4: 35 (33); 1.5/4.5: 27 (28). I used the obtained relative number of reinforcements on each key and a lambda of 0.53/s to derive these predictions.

Fantino, Squires, Delbrück, and Peterson (1972) kept reward durations fixed at 6 s and 1.5 s, and varied the length of the equal concurrent VI schedules. A value for lambda of 0.78/s permits the model to describe their four averaged data points quite accurately, with only 1% deviation between the obtained and predicted points.

Todorov (1973) correlated green, red, and yellow keylights with VI 36-s, VI 45-s, and VI 90-s schedules. In each condition of the experiment he provided one of three different hopper access durations (2 s, 4 s, and 8 s) for each of the schedules. Within each daily session, he compared preference between two of the schedules for one third of the session, between another two for a third of the session, and between the last two for the remainder of the session. With a value of 0.65/s for lambda, the model accounts for 83% of the data variance (average deviation = 5%).

Schnieder (1973) varied from 2 to 30 the number of pellets that could be produced by responses on concurrent schedules whose mean value ranged from 12 to 180 s. Application of the model suggests a shallow curvature (lambda = 0.10/s, assuming 0.5 s to eat each pellet, as reported by Schneider, 1973). Accuracy of predictions was low, with an average deviation of 6% between expected and obtained observations. In some of Schneider's conditions, the amount of time consuming the pellets was twice as great as the average time in the initial link. We may generalize the model, as suggested in Killeen (1981) and Killeen and Smith (1984), and assume that this eating time be included in the calculation of reinforcement rate, and that it be weighted more heavily than other time included in that calculation. Such a modification, with a 3/1 weighting, reduces the average deviation to less than 5%.

Essock and Reese (1974) gave pigeons a choice between two FR 30 schedules, one leading to 5-s access to grain and the other to a variable-duration access to grain. The subjects responded more on the key that produced fixed-duration access. Again, this suggests a convex utility function, but here (unlike for Young, 1981) there is no interaction with terminal-link delay to negate that inference. My only hypothesis about this finding involves the experimenters' use of fixed-ratio schedules of reinforcement. Animals pause longer after a larger reward (Harzem & Harzem, 1981; Staddon, 1974; but see Powell, 1969) and response rate after the pause does not change reliably with quality (Lowe, Davey, & Harzem, 1974) or duration (Powell, 1969) of reward—perhaps because the ratio contingency has already moved response rate close to its ceiling. These peculiarities of fixed schedules, and in particular fixed-ratio schedules, somewhat weaken the force of Essock and Reese's findings as evidence against decreasing marginal utility for duration of a reward.

The final study to be analyzed, that of Davison and Hogsden (1984), provides more data than all of the others combined. They present data from six pigeons in 34 conditions, having conducted approximately 24 sessions for each condition. They varied both duration and rate of access to grain. Their study was unique in that it provided different amounts of grain, probabilistically, contingent upon left-key responses, while keeping the amount available from the right key fixed at 3-s access. In addressing the 204 data points that resulted, I deleted two that the authors considered suspect and averaged the remaining data across subjects. The model provides a poor fit to all of the conditions taken together. This is largely due to deviations in the last four conditions, in which the pigeons show a marked preference for the left key. (During the 200 sessions previous to those, only the frequency of reward and not its duration was varied; in the last four conditions, duration was varied, and the variation was no longer probabilistic, as in previous conditions.) If we set aside these last four conditions, the fit improves substantially to account for over 91% of the data variance in the first 30 conditions (average deviation = 3%). The theoretical parameters are: bias 0.76; lambda 0.27/s; and hopper latency of 0.1 s.

In summary, there have been numerous studies that have reported effects of varied duration and quality of reward on response rates on single reinforcement schedules. Almost all of these studies show a concave relation between magnitude of reward and response rate. There are multiple potential causes of such curvature—ceilings on response rates, competing reinforcement from background events and the concave utility function. Whatever its source, that concavity is consistent with the predictions of incentive theory. When both rate and magnitude of reinforcement are varied, the data are well fit by a hyperbolic function both of whose parameters are affected by the value of the magnitude (e.g., see McDowell & Wood, 1984); again, this is consistent with the predictions of incentive theory (Killeen, 1981, Equation 38).

DISCUSSION

Concave Utility Function

Equation 3 entails "decreasing marginal utility"—doubling the amount of a commodity less than doubles its value to the organism. The larger the parameter lambda, the greater this nonlinearity. This makes intuitive sense for the kinds of commodities that affluent humans deal with—the second car or steak dinner or million dollars may be good, but certainly not as good as the first. But in these examples, part of the decrease may be attributed to satiation (you may be able to fit only one car in your garage and one steak in your stomach; all basic needs have been satisfied by the first million, and any utility in the second must follow the cultivation of new appetites). Those intuitions do not apply to pigeons and rats maintained at 80% of their free-feeding weights and choosing between 2 s of eating versus 6 s of eating, given that it would take at least 2 min of eating to fill their bellies. One must seek other intuitions to reconcile both the data and their mathematical description with our tacit assumption that most behavior is in some sense optimal. The second 5 s of eating should be just as important as the first, and certainly will generate an identical increment in calories per hour, which is the type of utility scale that both economic behavioral models and foraging models seem to favor. Populations of animals that operate according to concave functions in a common feeding area ought to be easily displaced by more "linear" animals. What is wrong with these inferences?

They may be right. The concavity of the utility function may result from other causes, such as a slowing of eating within each bout. This is suggested by the more linear functions found when the independent variable is the number of food pellets given to pigeons, in which case they are guaranteed receipt of the amount scheduled. When hopper time is varied, animals may get proportionately less at longer durations. This explanation is supported by the research of Gunn (1983), who measured rats' rate of consumption of food by providing access to pellets on a continuous reinforcement schedule as rewards for responses on an intermittent schedule. Gunn found a rapid decrease in the rate of acquiring and consuming pellets throughout the reward episode.

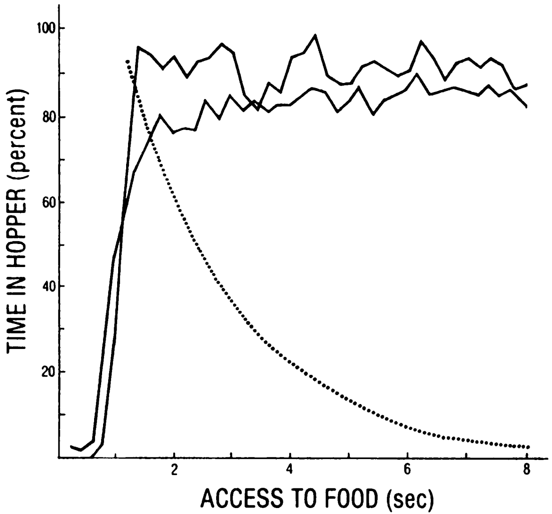

Fig. 2. The percentage of time that two pigeons had their heads in a food hopper, as a function of the time during the 8-s feeding episode. The curves rise to their maximum at about 1 s after the operation of the hopper, and remain at that level until the end of the interval. If a decrease in eating rate were the cause of the decrease in the marginal utility of food, the eating curves would have to follow a course such as the dotted curve, an exponential decay function with a value of 0.5/s for lambda.

But this does not address the question of why consumption rate should decrease. Perhaps that is a result of Equation 3, rather than a cause of it. Decreasing consumption rate is no kinder to our intuitions than is decreasing marginal utility. Furthermore, the data of Ito and Asaki (1982) and Davenport (1962) require concave utility functions, even though they are from rats choosing between fixed numbers of pellets.

Does consumption rate decrease for pigeons over the course of an extended hopper presentation? Apparently not: Figure 2 shows the percentage of time that two pigeons kept their heads in the hopper, as measured by a photocell, when given 8-s access to the hopper every 60 s. The data are based on 40 eating bouts. The figure shows that the animals regularly had their heads in the hopper about 1 s after its presentation, and that they kept their heads in the hopper (eating, as we observed) consistently thereafter. One pigeon showed a slight decrease in eating through the interval, the other a slight increase. In order to account for a concave utility function with a lambda of 0.5/s, a value typical for the data analyzed here, eating rate would have to follow the dotted curve, which shows the decrease in marginal utility of additional food.

But this is not the complete story, Epstein (1981) demonstrated that for one type of food hopper (the Gerbrands Model B feeder), the amount of food obtained over the course of an extended feeding was not proportional to the time the hopper was kept in operation. Apparently, the food depleted and was not replenished so long as the hopper was kept elevated. Subsequently, Epstein (1984) showed that this does not happen for a Lehigh Valley feeder such as the one used to generate the data in Figure 2.

The implication of Epstein's work is best understood in terms of Equation 3. The parameter lambda tells us the rate of decay of effectiveness of a reward as its duration is extended. Its reciprocal, the time constant of the system, tells us how long it will take for value to reach any stated level: In particular, in $1/\lambda$ s the marginal value will have decayed to 37%, in $2/\lambda$ s to 14%, and in $3/\lambda$ s to 5%. Thus, for a lambda of 0.5/s, the value of a 6-s reward is 95% of the maximum possible (i.e., it is 95% as good as a reward of indefinite duration). In this case, the effect of doubling a reward from 6-s to 12-s duration should be scarcely measurable. In studies that have varied the duration of punishing electric shock (e.g., Church, Raymond, & Beauchamp, 1967; Leander, 1973), the value for lambda is several times larger, indicating that the impact of the shock approaches to within 95% of its maximum at durations on the order of 1.5 s.

To what extent, then, do the concave utility functions reflect apparatus design rather than evolutionary design? The "time constant" for the Gerbrands feeder, based on Epstein's (1981) data, is 4.6 s. This means that the hopper will be depleted of 95% of its available food in just 14 s. In studies using this type of hopper, we should not expect the value of lambda to be any smaller than the one found by Epstein, 0.22/s. However, most of the values for lambda that we found were well above that level, suggesting the operation of additonal "psychological" mechanisms involved in the observed curvature. Furthermore, the time constant found for Lehigh Valley feeders was 59 s, indicating that it would require some 3 min for the process of removing grain from the hopper to reach 95% of completion. This is on the same order as a satiation function for pigeons; depletion from feeders of this design is essentially linear over the range of hopper times of typical experimental interest and cannot account for the observed curvature.

Concave utility functions, such as represented by Equation 3, are inconsistent with assumptions that animals maximize the number of calories acquired in some unit of time. Because later instants of eating are "discounted" more than earlier instants, such concavity provides a way of weighting the magnitude of a reward less heavily than its frequency. It is possible to generate plausible foraging scenarios in which this is a good thing to do. But I turn now to a more molecular analysis. What behavioral mechanisms might underlie such concavity? The second, third, and *n*th seconds of access to food are not contiguous with the response that initiated the reward: They are separated from it by a delay—a delay filled with a most significant (eating) activity. That activity may block reinforcing effects of subsequent eating in that bout. If the probability of blocking is constant through the interval, then an exponential gradient will result, and the integral of reinforcement strength for the bout will be Equation 3.

It is also known that as hopper durations are increased to 4 s, pigeons' ability to report whether a response initiated the feeding decreases to chance levels (Killeen & Smith, 1984); that rate of decay of memory indicates a lambda of 0.5/s, which is within the range reported here.

I have not attempted to eliminate other possible utility functions, such as power functions. Equation 3 is recommended by its consistency with several theoretical perspectives (blocking, and exponentially weighted moving averaging of input; Killeen, 1981). Its predictions differ from those of power functions: For the latter, animals should show the same preference between two magitudes so long as their ratios are constant (e.g., 12 s should be preferred over 6 s by the same amount that 4 s is preferred over 2 s); for Equation 3, preference should decrease as absolute durations increase, even if the ratios are held constant. These details are less important than the fact that in all but two of the dozen studies inspected here, the utility function is nonlinear, and in most cases strongly so.

Interaction of Incentive Value and Delay

Increasing the magnitude of an incentive increases its ability to bridge delays, above and beyond biasing the organism in favor of that alternative. But why should greater incentive value flatten the delay gradient? The rationale for the exponential form of the gradient was based on the presumption of a constant probability of blocking of the primary effects of an incentive by intervening stimuli or responses. If this is the mechanism that produces the gradient, why should not increased value of reward increase blocking to the same extent that it increases the association with the target response, and therefore have no net effect? I earlier (Killeen, 1982b) suggested that the target response often has the advantage in association, because it is usually a "prepared," or "terminal," response—one that is likely to be elicited by the incentive; increased amounts of reward may merely increase this advantage. This hypothesis is lent support by the data of Reid and Dale (1983), who varied the amount of reward and measured its effect on schedule-induced drinking and hopper approach. They found that hopper approach increased and drinking decreased when large meals began or ended an interval.

If this chain of reasoning is correct, the interaction of value with gradient will be less extreme for responses that are not differentially associable with that incentive (i.e., that are "unprepared"). For responses that incur strong competition (i.e., are "contraprepared"), increasing the magnitude of a delayed reward may actually be counterproductive. The exact class of these responses depends, of course, on the nature of the incentive. These predictions should be easily tested and are important for our understanding of the nature of constraints on conditioning. This may, for instance, be the mechanism of "instinctive drift" (Breland & Breland, 1966; Timberlake, Wahl, & King, 1982), in which particular response topographies are favored by delayed rewards.

The interaction of magnitude with delay is established, but its mechanism is uncertain. The above speculation—that increasing the magnitude of a delayed reward might faciltate "prepared" responses such as reaching for the reward, but undermine "contraprepared responses" such as withdrawal—has important implications for self-control. Some responses may be more easily emitted in efforts at self-control than others; it may be easier to say "No" than to leave a situation; people may clap a hand over their mouths to stifle a verbal response, and they may jam their hands in their pockets to prevent fidgeting (Skinner, 1953, p. 231). It is responses, not the organism, that are motivated and conditioned, and different response classes have different properties. Skinner (1953, p. 295) spoke of the "self" as "a functionally unified system of responses," but such unity is strained by rewards and punishers. We may wish one thing with our neutral ("unprepared") verbal responses, yet continually find ourselves doing otherwise with responses that are more tightly bound to their rewards.

REFERENCES

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior,* **22,** 231-242.

Breland, K., & Breland, M. (1966). *Animal behavior.* New York: Macmillan.

Brownstein, A. J. (1971). Concurrent schedules of response-independent reinforcement: Duration of a reinforcing stimulus. *Journal of the Experimental Analysis of Behavior,* **15**, 211-214.

Church, R. M., Raymond, G. A., & Beauchamp, R. D. (1967). Response suppression as a function of intensity and duration of a punishment. *Journal of Comparative and Physiological Psychology,* **63**, 39-44.

Davenport, J. W. (1962). The interaction of magnitude and delay of reinforcement in spatial discrimination. *Journal of Comparative and Physiological Psychology,* **55**, 267-273.

Davison, M., & Hogsden, I (1984). Concurrent variable-interval schedule performance: Fixed versus mixed reinforcer durations. *Journal of the Experimental Analysis of Behavior,* **41**, 169-182.

Epstein, R. (1981). Amount consumed as a function of magazine-cycle duration. *Behaviour Analysis Letters,* **1**, 63-66.

Epstein, R. (1984). *Amount consumed as a function of feeder design.* Unpublished manuscript.

Essock, S. M., & Reese, E. P. (1974). Preference for and effects of variable—as opposed to fixed—reinforcer duration. *Journal of the Experimental Analysis of Behavior,* **21**, 89-97.

Fantino, E. (1969). Choice and rate of reinforcement. *Journal of the Experimental Analysis of Behavior,* **12**, 723-730.

Fantino, E., Squires, N., Delbrück, N., & Peterson, C. (1972). Choice behavior and the accessibility of the reinforcer. *Journal of the Experimental Analysis of Behavior,* **18**, 35-43.

Green, D. R. (1969). *Preference in concurrent-chains as a function of reinforcement delay and amount.* Unpublished doctoral dissertation, Arizona State University.

Green, L., & Snyderman, M. (1980). Choice between rewards differing in amount and delay: Toward a choice model of self control. *Journal of the Experimental Analysis of Behavior,* **34**, 135-147.

Gunn, K. P. (1983). Quantification of rats' behavior during reinforcement periods. *Journal of the Experimental Analysis of Behavior,* **39**, 457-464.

Harzem, P., & Harzem, A. L. (1981). Discrimination, inhibition, and simultaneous association of stimulus properties: A theoretical analysis of reinforcement. In P. Harzem & M. D. Zeiler (Eds.), *Advances in analysis of behaviour: Vol. 2. Predictability, correlation, and contiguity* (pp. 81-124). Chichester, England: Wiley.

Ito, M., & Asaki, K. (1982). Choice behavior of rats in a concurrent-chains schedule: Amount and delay of reinforcement. *Journal of the Experimental Analysis of Behavior,* **37**, 383-392.

Killeen, P. R. (1972). The matching law. *Journal of the Experimental Analysis of Behavior,* **17**, 489-495.

Killeen, P. R. (1979). Arousal: Its genesis, modulation, and extinction. In M. D. Zeiler & P. Harzem (Eds.), *Advances in analysis of behaviour: Vol. 1. Reinforcement and the organization of behaviour* (pp. 31-78). Chichester, England: Wiley.

Killeen, P. R. (1981). Averaging theory. In C. M. Bradshaw, E. Szabadi, & C. F. Lowe (Eds.), *Quantification of steady-state operant behaviour* (pp. 21-34). Amsterdam: Elsevier/North Holland.

Killeen, P. R. (1982a). Incentive theory. In D. J. Bernstein (Ed.), *Nebraska symposium on motivation: Vol. 29. Response structure and organization* (pp. 169-216). Lincoln: University of Nebraska Press.

Killeen, P. R. (1982b). Incentive theory: II. Models for choice. *Journal of the Experimental Analysis of Behavior,* **38**, 217-232.

Killeen, P. R. (1985). *Incentive theory and the initial links of concurrent chain schedules.* Manuscript submitted for publication.

Killeen, P. R., & Smith, J. P. (1984). Perception of contingency in conditioning: Scalar timing, response bias, and erasure of memory by reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes,* **10**, 333-345.

Leander, J. D. (1973). Shock intensity and duration interactions on free-operant avoidance behavior. *Journal of the Experimental Analysis of Behavior,* **19**, 481-490.

Lowe, C. F., Davey, G. C. L., & Harzem, P. (1974). Effects of reinforcement magnitude on interval and ratio schedules. *Journal of the Experimental Analysis of Behavior,* **22**, 553-560.

McDowell, J. J, & Kessel, R. A. (1979). A multivariate rate equation for variable-interval performance. *Journal of the Experimental Analysis of Behavior,* **31**, 267-283.

McDowell, J. J, & Wood, H. M. (1984). Confirmation of linear system theory prediction: Changes in Herrnstein's k as a function of changes in reinforcer magnitude. *Journal of the Experimental Analysis of Behavior,* **41**, 183-192.

Menlove, R. L., Inden, H. M., & Madden, E. G. (1979). Preference for fixed over variable access to food. *Animal Learning & Behavior,* **7**, 499-503.

Navarick, D. J., & Fantino, E. (1976). Self-control and general models of choice. *Journal of Experimental Psychology: Animal Behavior Processes,* **2**, 75-87.

Neuringer, A. J. (1967). Effects of reinforcement magnitude on choice and rate of responding. *Journal of the Experimental Analysis of Behavior,* **10**, 417-424.

Powell, R. W. (1969). The effect of reinforcement magnitude upon responding under fixed-ratio schedules. *Journal of the Experimental Analysis of Behavior,* **12**, 605-608.

Reid, A. K., & Dale, R. I. (1983). Dynamic effects of food magnitude on interim-terminal interaction. *Journal of the Experimental Analysis of Behavior,* **39**, 135-148.

Schneider, J. W. (1973). Reinforcer effectiveness as a function of reinforcer rate and magnitude: A comparison of concurrent performances. *Journal of the Experimental Analysis of Behavior,* **20**, 461-471.

Skinner, B. F. (1953). *Science and human behavior.* New York: Macmillan.

Snyderman, M. (1983). Delay and amount of reward in a concurrent chain. *Journal of the Experimental Analysis of Behavior,* **39**, 437-447.

Staddon, J. E. R. (1974). Temporal control, attention, and memory. *Psychological Review,* **81**, 375-391.

Timberlake, W., Wahl, G., & King, D. (1982). Stimulus and response contingencies in the misbehavior of rats. *Journal of Experimental Psychology: Animal Behavior Processes,* **8**, 62-85.

Todorov, J. C. (1973). Interaction of frequency and magnitude of reinforcement on concurrent performances. *Journal of the Experimental Analysis of Behavior,* **19**, 451-458.

Williams, B. A. (1982). Blocking the response-reinforcer association. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Vol. 3. Acquisition* (pp. 427-447). Cambridge, MA: Ballinger.

Young, J. S. (1981). Discrete-trial choice in pigeons: Effects of reinforcer magnitude. *Journal of the Experimental Analysis of Behavior,* **35**, 23-29.

Received August 23, 1983
Final acceptance February 27, 1985