

Animal Learning & Behavior
1981, 9 (4), 476-482

Preference reversal and delayed reinforcement

GEORGE AINSLIE

Massachusetts Mental Health Center, Boston, Massachusetts 02115

and

R. J. HERRNSTEIN

Harvard University, Cambridge, Massachusetts 02138

Preference reversal and delayed reinforcement

GEORGE AINSLIE

Massachusetts Mental Health Center, Boston, Massachusetts 02115

and

R. J. HERRNSTEIN

Harvard University, Cambridge, Massachusetts 02138

Pigeons chose, in a two-key discrete-trial procedure, between 2- and 4-sec access to grain, with the larger amount always presented 4 sec later than the smaller. As the delay between the choice and the availability of the smaller reinforcement was varied from .01 to 12 sec, all subjects reversed preference from the small-early to the large-late reinforcement. The values of delay at which preference reversed were approximately consistent with the matching law as adapted for delayed reinforcement.

The effect of reinforcement on performance declines steeply as a function of delay. Both exponential and hyperbolic curves have been proposed to describe this decline (Ainslie, 1975; Chung & Herrnstein, 1967; Renner, 1964). Exponential curves are simple percentage discount functions, by which a reinforcement loses a fixed proportion of its value per unit of delay time. Hyperbolic curves depict the value of a reinforcement as inversely proportional to its delay; they are more concave than exponential curves. The exact shape of the function is important, since a hyperbola frequently predicts a striking reversal of preference, given pairs of reinforcers varying in amount and in delay. If the parameters are in a certain range, the hyperbolic function says that as a choice draws near, preference may suddenly shift from the more remote and larger reinforcer to the smaller but earlier one (e.g., Figure 1). Such preference reversals may play an important role in human impulsiveness and "ego defense" (Ainslie, 1975; Brown & Herrnstein, 1975; Rachlin, 1970; Strotz, 1956; Ainslie, Note 1). Exponential or still less concave curves do not predict preference reversals without additional assumptions (Strotz, 1956).

A hyperbolic delay function is approximately consistent with previously published data (e.g., Chung & Herrnstein, 1967). But Williams and Fantino (1978) have challenged the hyperbolic relationship with a reanalysis of Chung and Herrnstein's data and with a new experiment using fixed-interval schedules in the terminal links of a concurrent-chains procedure. Given the vagaries of curve fitting and the multiplicity of procedural parameters, it would be

This research was conducted with support from NIMH, Grant MH-15494 to Harvard University. G. Ainslie is now at Veterans Administration Hospital, Coatesville, Pennsylvania 19320. Reprints may be obtained from either author.




Figure 1. Hyperbolic discount curves showing the effectiveness of two hypothetical reinforcements as a function of delay: the effectiveness, at any given moment on the time axis, t , of a reinforcement occurring at time T vs. a reinforcement twice as large occurring at time $T+3$. Left curve shows $y = 10 / (T - t)$; right curve, $y = 20 / [(T + 3) - t]$.

hard to draw a firm conclusion about delay functions from existing data. The occurrence of the preference reversal phenomenon would be a more decisive form of evidence that the delay function is more concave than an exponential curve, besides being of major importance in its own right. Using various procedures

and rat, pigeon, and human subjects, preference reversals as a function of temporal remoteness of alternatives have been suggested by the data in a number of studies (Ainslie, 1974; Deluty, 1978; Fantino, 1966; Logan & Spanier, 1970; Navarick & Fantino, 1976; Rachlin & Green, 1972; Solnick, Kannenberg, Eckerman, & Walker, 1980). However, each of these studies has had limitations. The Ainslie experiment used responding vs. not responding rather than a symmetrical choice as a measure of preference. Solnick et al. found that differences in the instructions their subjects received determined whether or not preference reversal occurred. The other authors used response/reinforcement ratios much greater than one, so that the actual delay from most responses to the next reinforcer was variable and considerably larger than the nominal delay. Although nondiscrete-trial choices with variable interval or high fixed-ratio links before the delay undoubtedly add sensitivity to choice experiments, they cannot be assumed to motivate subjects in the same way as their corresponding discrete trial, FR1 schedules would. The present paper reports a simple, symmetrical test of preference reversal in pigeons, and, hence, of whether these subjects' discount curves for delayed reinforcement are more concave than exponential.

METHOD

Subjects

Six male White Carneaux pigeons, experienced in a number of other operant procedures, were maintained at 80% of their free-feeding weights.

Apparatus

A standard, two-key chamber for pigeons was used. The two keys were on one wall, side by side 5 in. (13 cm) apart, at a height of about 9 in. (23 cm). A force of about 15 g (.15 N) was required to operate each key. Each peck on the keys also operated a feedback relay in the chamber. The keys could be transilluminated with lights of different colors. Beneath the keys was an opening into the standard feeding magazine, which was lit with white light when food was available. The chamber was continuously illuminated with a pair of 7-W white bulbs, and continuous white noise masked extraexperimental sounds. The opening to the feeder contained a swinging gate, hinged at the top and connected to a microswitch, which permitted the duration of reinforcement to be timed from the moment the pigeon's head entered the opening. The procedure was controlled by a PDP-8 computer, which also collected the data.

Procedure

Each subject was given 40-45 trials daily, each of which lasted 60 sec regardless of the subject's behavior. At the beginning of each trial, the left key was green and the right key, white. A peck on the left (green) key darkened both keys and raised the food hopper D sec later. The hopper remained raised for 2 sec after the subject's head entered the feeder opening. The keys then remained dark until the beginning of the next trial, 60 sec from the trial's onset. A peck on the right (white) key darkened both keys and raised the food hopper D + 4 sec later. The hopper remained raised for 4 sec after the subject's head entered the feeder opening, whether or not the bird kept his head continuously in the opening. The key then remained dark until the beginning of the next trial.

If no peck occurred within 10 sec of the beginning of a trial, the keys darkened and remained dark until the beginning of the next trial. The experimental session ended when the subject had pecked on 40 trials or when 45 trials had taken place, whichever happened first.

The value of D was initially set at .01 sec for all subjects, making the smaller reinforcer available virtually immediately after the appropriate peck. For Subjects 1, 2, and 3, D was increased from .01 sec to 2, 4, 6, 8, and 12 sec, then returned to .01 sec; for Subjects 4, 5, and 6, D went from .01 to 12 sec and was then decreased to 8, 6, 4, 2, and .01 sec. Although stable preferences usually formed within the first few sessions of a new contingency, D was not changed until no subject was pecking outside of the range of rates it had emitted on previous trials with the same value of D, and no subject seemed to be moving toward the edge of that range. The experiment ran for 320 sessions extended over a period of 11 months and was divided roughly equally among the seven conditions.

RESULTS

Data from the last 10 sessions in each condition were analyzed.

All subjects initially preferred the smaller reinforcer at .01 sec rather than the larger reinforcer at 4 sec (Figure 2). As the delay, D, was increased, all subjects reversed preference, choosing the large-late

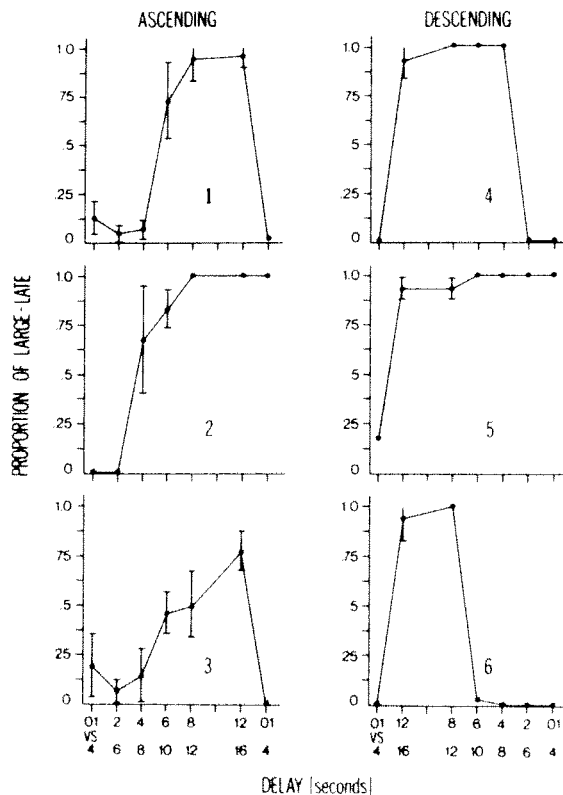


Figure 2. The ordinate shows the ratio of large-late choices to total choices; the abscissa shows the two reinforcement delays in each choice. Each graph is for a single subject, using the mean of the final 10 sessions of each experimental condition. Vertical lines indicate ± 1 standard deviation around the means and are omitted when the standard deviation was almost or exactly zero.

reinforcer on most or all trials. When D was decreased to .01 sec again, four subjects again chose the small-early reinforcer on virtually all trials; two subjects (2 and 5) continued to choose the large-late reinforcer on all trials. Figure 2 summarizes the main findings for individual subjects. The ordinate gives the large-late reinforcers as a proportion of all reinforcers; the abscissa gives the two delays at each condition, which differ by 4 sec, plotted in the order in which they were imposed. Somewhere between .01 and 8 sec, all subjects reversed preference. Lines through the data points show ± 1 standard deviation (estimated population value) around the mean of the individual sessions. Where no variability lines are shown, the standard deviation was zero or virtually zero.

The sharp transitions in preference limit the value of an average across subjects. Nevertheless, Figure 3 captures some of the salient features of the individual curves. Figure 3 shows the median values of the points plotted in Figure 2. The median curve shows that, on the average, preference reversed at about 4-6 sec to the earlier reinforcer. At greater remoteness, subjects consistently chose the larger reinforcer; at lesser remoteness, they chose the smaller, but earlier, one.

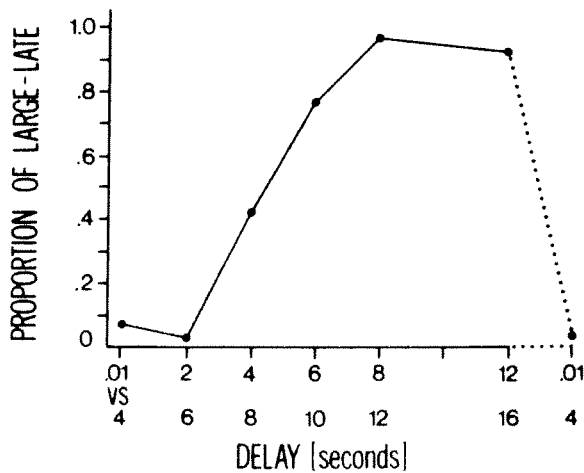


Figure 3. Median values of the points plotted in Figure 2 (not in temporal order for half the subjects).




Figure 4. Latencies to peck for those trials in which the subjects pecked, as a function of the reinforcement delays in each choice. The solid line shows latencies for trials in which the earlier alternative was chosen; the dashed line shows those in which the later alternative was chosen. Points are averages across subjects of the latencies in Table 1.

Sessions were terminated after the 40th reinforcement or the 45th trial, whichever came first. The average number of reinforcers per session ranged from 37.9 (over all conditions) for Subject 1 to 39.9 for Subject 5. Latencies to peck are summarized in Table 1. Latencies to peck tended to be longer with longer delays, in accordance with the notion that response strength is inversely related to delay of reinforcement. Figure 4 shows the arithmetic means of the large-late and small-early peck latencies in Table 1. Exclusive choice of one alternative meant no latency for the other choice, as shown by the blanks in Table 1. The averages in Figure 4 are based only on the entries in Table 1. They suggest slightly slower pecking to the large-late alternative, except at the

Table 1
Peck Latency (in Seconds)

Delay	Subject 1		Subject 2		Subject 3		Subject 4		Subject 5		Subject 6	
	E	L	E	L	E	L	E	L	E	L	E	L
.01 vs. 4	1.75	3.47	2.10		1.59	1.56	1.65	1.00	1.42	.82	1.11	1.20
2 vs. 6	1.47	3.21	1.98	2.15	1.94	2.47	2.60	1.46		2.10	.83	
4 vs. 8	2.85	2.63	3.04	3.03	1.72	3.11		4.16		2.43	1.12	
6 vs. 10	3.11	3.28	1.35	3.23	1.83	2.81		4.05		2.16	1.86	4.87
8 vs. 12	3.91	3.38		4.53	2.30	4.58		4.21	2.47	2.08	4.40	3.04
12 vs. 16	2.39	3.36	4.80	2.86	2.87	3.54	3.86	1.90	3.67	1.96	5.91	3.76
.01 vs. 4	1.05	.82		2.29	1.27		1.65	1.20		2.10	.75	

Note--E = earlier reinforcement; L = later reinforcement.

longest delays, though an analysis of variance failed to confirm the significance of this difference.

The contingency table (Table 2) tests the hypothesis that latency and preference were associated. The contingency table was constructed from Table 1, using only conditions that provided latencies for both alternatives. It is evident that in the resulting sample, the relative latencies to the alternatives were independent of their preferability.

Table 3 summarizes the latencies between the presentation of the feeder tray and the pigeon's head's entering the feeder opening. The reinforcement durations, it may be recalled, were timed from the moment of entry. These feeder approach latencies were mostly between .75 and 2.25 sec, did not differ for the two alternatives, and did not interact significantly with delay. The averages in Figure 5 are arithmetic means based on all the entries in Table 3.

DISCUSSION

Preference reversed from a small to a large reinforcer separated by 4 sec as D, the time before the earlier but smaller reinforcer, increased from .01 to 12 sec. This finding probably indicates that the pigeon's delay function for food reinforcement is more concave than an exponential curve. For an exponential curve to be compatible with preference reversal, it would have to be assumed that the constant which determines the rate of time discounting varies with amount of reinforcement within a single modality (food), an assumption that would lack both parsimony and, at present, empirical justification. In any event, change of preference as a function of D is here

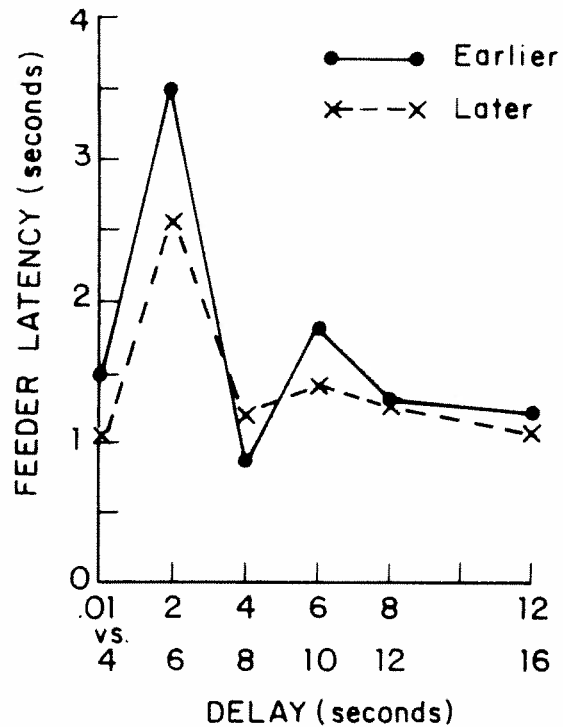


Figure 5. Latencies between the signaled availability of food and the opening of the door on the feeder. The solid line shows latencies for trials in which the earlier alternative was chosen; the dashed line shows those in which the later alternative was chosen. Points are averages across subjects of the latencies in Table 3.

observed directly, confirming the existence of an important problem for theories of motivational conflict (Ainslie, 1975).

The data not only argue against an exponential delay function, but also provide at least limited evidence to support hyperbolic time discounting, as implied by the matching law. The matching law states that an organism on a pair of concurrent variable-interval schedules, each with its own rate, amount, and delay of reinforcement, distributes responding in direct proportion to the rates and amounts and in inverse proportion to the delays of reinforcement

Table 2
Association Between Latency and Preference

Shorter Latency to:	Preference for:	
	Earlier	Later
Earlier	8	5
Later	7	8

Note— $\chi^2 = .6244$ (n.s.).

Table 3
Feeder Latency (in Seconds)

Delay	Subject 1		Subject 2		Subject 3		Subject 4		Subject 5		Subject 6	
	E	L	E	L	E	L	E	L	E	L	E	L
.01 vs. 4	2.51	.93	1.52		1.24	1.44	1.17	.87	1.08	1.22	.99	1.60
2 vs. 6	1.00	.73	1.23	1.15	1.12	1.38	2.06	5.78		1.43	2.11	
4 vs. 8	.87	.82	1.21	1.15	.86	1.23		1.06		1.63	.54	
6 vs. 10	4.89	1.11	.53	.89	1.29	1.88		1.07		1.21	.59	2.27
8 vs. 12	.78	.79		.89	1.87	2.20		1.54	1.16	1.20	1.40	.75
12 vs. 16	.60	.54	1.00	.54	1.92	1.64	1.14	1.16	1.32	1.22	1.30	1.27
.01 vs. 4	2.20	.60		.69	1.34		1.55	.90		1.34	.97	

Note—E = earlier reinforcement; L = later reinforcement.

from the alternatives (Baum & Rachlin, 1969; Chung & Herrnstein, 1967):

$$\frac{B_1}{B_2} = \frac{R_1}{R_2} \cdot \frac{A_1}{A_2} \cdot \frac{D_2}{D_1} \quad (1)$$

where B_i is rate of emitting behavior to the alternatives, and R_i is the rate, A_i is the amount, and D_i is the delay of reinforcement.

When R_1 and R_2 are unequal, Equation 1 cannot be solved to predict a rate of preference under the contingencies used in this experiment; we will discuss this problem presently. However, at the point at which the subject is indifferent between the two alternatives, $B_1 = B_2$ and thus $R_1 = R_2$. Since $A_1/A_2 = .5$ in this experiment (2 sec of feeder time divided by 4 sec of feeder time), this indifference or crossover point should be reached when D_2/D_1 passes through 2.0.

Since D_2 is 4 sec greater than D_1 , we can readily solve for the critical value of D , as follows:

$$\frac{D_2}{D_1} = \frac{D+4}{D} = 2, \text{ so } D = 4. \quad (2)$$

Figures 2 and 3 show that preference did, in fact, reverse when the earlier reinforcer was delayed by approximately 4 sec. On the average, the departure from prediction was toward values of D a second or two longer. It could be argued that the usual convention of measuring delay (D) from a subject's peck to the elevation of the feeder should be abandoned in favor of the interval from the peck until actual eating takes place, or even until the food enters the subject's bloodstream. The "true" moment of reinforcement is not known, and might not even exist as a distinct instant; given the matching law, such a point could be empirically defined, but in the present case this procedure would be circular. However, since pecking-to-eating intervals can be calculated from our data by adding each feeder latency to the corresponding scheduled delay, we can at least calculate the consequences of using these intervals for D s.

Table 3 contains enough data cells to predict crossover points for the first four birds, using this alternative interpretation of D ; we assume that Bird 4's missing feeder latency for the earlier reinforcer at $D=4$ would have been at least .53 sec, a value far less than any other feeder latency obtained from that subject. In all four birds, this calculation predicts crossovers between the 2- vs. 6-sec and 4- vs. 8-sec conditions, instead of right at 4 vs. 8 sec, as the peck-to-feeder-elevation convention predicts. Birds 2 and 4 actually crossed between these two conditions; Birds 1 and 3 crossed or sharply started to cross in the next interval, between 4 vs. 8 sec and 6 vs. 10 sec. Thus, an alternative prediction of crossover points

taking the pecking-to-eating interval as D also approximates our data. Deviations from the precise predictions of either convention might be due to slight inaccuracies in the ratio of A_1/A_2 or to oversimplification in the formulation of the matching law, speculation about which would not be appropriate here.

One problem we should not avoid, however, is the assumption of equal frequencies of reinforcement, that is, $R_1/R_2 = 1.0$. The assumption reflects the equivalent reinforcement schedules for the two alternatives, namely continuous reinforcement. However, obtained reinforcement rates, as distinguished from programmed rates, depend on the subject's behavior as well as on the schedule. Consequently, Equation 1, as applied here, fails to take into account the feedback relationship between preference and reinforcement rate. Qualitatively, we would expect that preference would tend to shift rapidly, for the ratio R_1/R_2 in Equation 1 changes as B_1/B_2 changes. Significant preference for either alternative over the other on the basis of amount and delay of reinforcement, reflected in the ratio B_1/B_2 , would change the ratio R_1/R_2 , which would then change B_1/B_2 in the same direction, causing a further change in R_1/R_2 , and so forth, in a positive feedback system until the preferred side was chosen exclusively. Figure 2, indeed, shows sharp transitions from exclusive preference for one alternative to exclusive preference for the other in all subjects but No. 3. Although it is possible in principle to do so (see Prelec & Herrnstein, 1978), the data do not permit us to test any particular positive feedback system here.

The hyperbolic function between behavior and delay implicit in Equation 1 predicted the average crossover point accurately in this experiment, but there was considerable variation among individual subjects. It may be possible to improve the fit of the matching law to individual subjects, or individual instances of behavior, by introducing an empirical constant, but, once again, the data at hand do not let us select among the many formal expressions for doing this. Nor can we say much about Equation 1's implication that a behavior reinforced with no actual delay whatever should completely replace its competitors, an extreme that is unlikely to occur in nature. However, truly immediate reinforcement may not occur, either. A formulation that attempts to deal with these complications has been presented in Herrnstein (1981). The simple formula given here describes only that range of delays which have been observed.

Whatever the exact shape of the delay function, the variations in the range of rates at which the subjects apparently discounted delayed reinforcement raises the possibility that there is no "best" rate, that is, none which is uniformly favored by natural selec-

tion. This possibility might seem nonoptimal, since the food intake of the subjects depended directly on their tendencies to wait for delayed reinforcement. It might therefore seem that natural selection should unequivocally favor low rates of time discounting, except for another consideration. The biological value of a low discount rate is limited by its requiring the organism to detect which one of all the events occurring over a preceding period of hours or days led to a particular reinforcer. As the discounting rate falls, the informational load increases. Without substantial discounting, a reinforcer would act with nearly full force not only on the behaviors that immediately preceded it, but also on those that had been emitted in past hours or days. The task of factoring out which behaviors had actually led to reinforcement could exceed the information processing capacity of a species. But even if the species' capacity is sufficient, low discounting rates have further risks. Since, in fact, behaviors tend to have decreasing control over increasingly remote events, an organism might learn to discard behaviors that preceded reinforcement by long periods of time, but such a high-order principle might take a fatally long time to learn. A high inborn discount rate would spare the organism the need to learn it by experience, although it would also reduce the likelihood that the organism could detect and be controlled by any remote consequences of its actions. Organisms with greater information-processing capacities or those living in more predictable environments could tolerate lower discount rates than others from the viewpoint of natural selection. In a diverse world, a number of different rates would have survived the evolutionary process. Inasmuch as Equation 1 provides no parameter for individual or species differences, it is apt to be incomplete.

Among steep discount curves, those concave enough to cross can be expected to offer a further selective advantage over those that are not. As D increases, the long tails of highly concave curves like the hyperbola implicit in Equation 1 approach proportionality to the sizes of the reinforcers on which they are based (see Figure 1). This can be shown formally, as follows. For a given constant difference between D_2 and D_1 , we can write $D_2/D_1 = (D + \Delta)/D$, in which Δ is the constant, D is the varying delay to the earlier reinforcer, and $D + \Delta$ is the varying delay to the later reinforcer. The ratio of frequencies, R_1/R_2 , is assumed to be 1.0, as earlier. As D increases,

$$\frac{B_1}{B_2} = \lim_{D \rightarrow \infty} \left(\frac{A_1}{A_2} \right) \left(\frac{R_1}{R_2} \right) \left(\frac{D + \Delta}{D} \right) = \left(\frac{A_1}{A_2} \right). \quad (3)$$

Less concave curves do not approach proportionality at increasing delays. Highly concave curves would

give an organism a mechanism for overcoming the motivational distortion caused by the steepness of the curve; less concave curves would not. If regularity of experience permits, the tails of curves which approach proportionality may motivate the learning of behaviors based on the objective amount of the reinforcement differential, and will do so without interfering with the organism's ability to learn rapidly from immediate reinforcement. In any recurring choice in which the alternatives are reinforced at different delays, the earlier reinforcer will continue to be disproportionately effective when their curves become steeply vertical (Figure 1); however, these segments will be preceded by low, but proportional, segments during which the "objective" differential may gradually have a compensating effect, motivating the learning of behaviors to forestall the later reversals of choice (Ainslie, 1974). Thus, two discontinuous time zones are defined by the operation of highly concave, continuous discount functions on an organism's preference. The persistence of learning in both time zones undoubtedly leads to the enduring patterns of motivational conflict that have been described elsewhere (Ainslie, 1975). Even counting the cost of this conflict, a combination of rapid learning favoring the earlier reinforcer and countervailing slow learning favoring the later reinforcer probably reaps a greater proportion of the available reinforcement than the learning that could be based on an exponential or other noncrossing curve. Any discounting according to noncrossing curves introduces a distortion in the relative effectiveness of nonsimultaneous reinforcers, and this distortion is not reduced by the perspective of distance.

Existing evidence suggests that even pigeons can learn to wait for later but larger reinforcers, both with the aid of physical precommitting devices (Ainslie, 1974) and without this aid. Mazur and Logue (1978) showed that pigeons increase their tendency to choose the later and larger reinforcer after a "fading" procedure in which the value of Δ (see Equation 3) is, in effect, gradually increased. Two of four subjects exposed to the fading procedure ultimately chose 6 sec of food delayed by 6 sec over 2 sec of food delivered immediately on more than half the trials; four subjects without fading consistently chose the small-early reinforcer. In another experiment, one of us (Ainslie, Note 2) found that some, but not all, pigeons that were given choices between a small-early and a large-late reinforcer abruptly shifted preference to the large-late alternative after months of stable choice of the small-early one, in the absence of a fading procedure. The evidence does not tell us whether the pigeons that learned to forbear without physical help had learned some internal precommitting device or had actually changed the shape of their discount curves. In the present study, two pigeons (2 and 5) may have similarly learned forbearance (see

Figure 2), but the design of our procedure does not rule out a simple position bias which precluded the subject's sampling both alternatives after forming an exclusive preference for the later one.

The experimental analysis of preference reversal may contribute to the practical art of teaching impulse control, since it allows us to vary the controlling dimensions of reinforcement and delay and to observe the balance tipping between impulsiveness and forbearance.

REFERENCE NOTES

1. Ainslie, G. W. *Impulse and compulsion: A behavioral economic understanding of involuntary behavior*. Unpublished manuscript, 1980.
2. Ainslie, G. W. *Correlates of peck-withholding on a discrete trial specious reward schedule*. Unpublished manuscript, 1981.

REFERENCES

- AINSLIE, G. W. Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 485-489.
- AINSLIE, G. Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 1975, 82, 463-496.
- BAUM, W. M., & RACHLIN, H. C. Choice as time allocation. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 861-874.
- BROWN, R., & HERRNSTEIN, R. J. *Psychology*. Boston: Little, Brown, 1975.
- CHUNG, S.-H., & HERRNSTEIN, R. J. Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 67-74.
- DELUTY, M. Z. Self-control and impulsiveness involving aversive events. *Journal of Experimental Psychology: Animal Behavior Processes*, 1978, 4, 250-266.
- FANTINO, E. Immediate reward followed by extinction vs. later reward without extinction. *Psychonomic Science*, 1966, 6, 233-234.
- HERRNSTEIN, R. J. Self control as response strength. In C. M. Bradshaw, E. Szabadi, & C. F. Lowe (Eds.), *Quantification of steady-state operant behaviour*. Amsterdam: Elsevier/North-Holland Biomedical Press, 1981.
- LOGAN, R. A., & SPANIER, D. Relative effect of delay of food and water reward. *Journal of Comparative and Physiological Psychology*, 1970, 72, 102-104.
- MAZUR, J. E., & LOGUE, A. W. Choice in a "self-control" paradigm: Effects of a fading procedure. *Journal of the Experimental Analysis of Behavior*, 1978, 30, 11-17.
- NAVARICK, D. J., & FANTINO, E. Self-control and general models of choice. *Journal of Experimental Psychology: Animal Behavior Processes*, 1976, 2, 75-87.
- PRELEC, D., & HERRNSTEIN, R. J. Feedback functions for reinforcement: A paradigmatic experiment. *Animal Learning & Behavior*, 1978, 6, 181-186.
- RACHLIN, H. *Introduction to modern behaviorism*. San Francisco: Freeman, 1970.
- RACHLIN, H., & GREEN, L. Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior*, 1972, 17, 15-22.
- RENNER, K. E. Delay of reinforcement: A historical review. *Psychological Bulletin*, 1964, 61, 341-361.
- SOLNICK, J. V., KANNENBERG, C. H., ECKERMAN, D. A., & WALLER, M. B. An experimental analysis of impulsivity and impulse control in humans. *Learning and Motivation*, 1980, 11, 61-77.
- STROTZ, R. H. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 1956, 23, 166-180.
- WILLIAMS, B. A., & FANTINO, E. Effects on choice of reinforcement delay and conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*, 1978, 29, 77-86.

(Manuscript received August 18, 1980;
revision accepted for publication April 1, 1981.)