MAXIMIZING AND MATCHING ON CONCURRENT RATIO SCHEDULES¹

R. J. HERRNSTEIN AND DONALD H. LOVELAND

HARVARD UNIVERSITY

Pigeons on concurrent variable-ratio variable-ratio schedules usually, though not always, maximize reinforcements per response. When the ratios are equal, maximization implies no particular distribution of responses to the two alternatives. When the ratios are unequal, maximization calls for exclusive preference for the smaller ratio. Responding conformed to these requirements for maximizing, which are further shown to be consistent with the conception of reinforcement implicit in the matching law governing relative responding in concurrent interval schedules.

Probability matching refers to a phenomenon sometimes observed (Bitterman, 1965; Brunswik, 1939) when subjects choose between alternatives that reinforce probabilistically. Suppose alternative A pays off one time in five and alternative B pays off one time in 10. Probability matching implies a 2:1 ratio of choices for A and B respectively, since $\frac{1}{5} : \frac{1}{10} :: 2 : 1$. Sometimes the probabilities are set so that one choice or the other is correct on any trial, which means that the probabilities add to 1.0. Or, the probabilities may add to 1.0 without the further constraint that one choice or the other is correct on every trial. In any case, the phenomenon of probability matching is a psychological oddity, for it may be grossly at odds with the rational course of action, to an extent rarely seen in nature. If the subject places value on the pay-off itself, rather than on the sheer matter of guessing correctly, he is ill advised to match probabilities. Instead, he should always choose the higher probability alternative, a strategy that maximizes the winnings per response.

We make no effort here to review the complex and extensive literature on probability matching (see Estes, 1964; Hilgard and Bower, 1966). Rather, our purpose is to establish an empirical bridge between that literature and the study of the matching law for

concurrent schedules (mainly concurrent *interval* schedules). There are significant connections between the two that have so far been left undeveloped. Procedurally, a short step takes us from one sort of situation to the other. The matching law, as noted above, usually describes concurrent variable-interval schedules. Probability matching, when it occurs, turns up in variants of concurrent *ratio* schedules. The present experiment examined concurrent variable-ratio schedules.

METHOD

Subjects

Five adult male White Carneaux pigeons, all with experience in various experiments, were maintained at about 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 (0.15 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 extra-experimental sounds.

¹This research was supported by Grant MH-15494 from the NIMH to Harvard University. Reprints may be obtained from R. J. Herrnstein, Department of Psychology, William James Hall, Harvard University, Cambridge, Massachusetts 02138.

Procedure

The experiment was composed of three series of concurrent variable-ratio variable-ratio schedules, with nominal values as listed in Table 1, along with the numbers of sessions at each pair, in order of exposure. The ratios for left and right key are shown by left and right numbers, respectively. The left key was always transilluminated with red light; the right key, with green, except during reinforcements, when both keylights were extinguished. Reinforcement was 2- or 3-sec access to food, as explained below. In series 1 and 2 (see Table 1), reinforcement could not occur within 1.5 sec of a switch in responding from one key to the other. Responses executed during the changeover delay (COD) contributed to the ratio requirement, but could not complete it. It took at least one peck following any COD to earn reinforcement. In series 3, there was no COD.

Note, in Table 1, that the sum of the two ratios in series 1 was always 60, 61, or 62; in the second and third series, the sum was always 120. Henceforth, we occasionally refer to the first series as "Sum 60", the second series as "Sum 120", and the third series as "Sum 120, No COD". In all cases, the distribution of ratio values was constructed according to

Table 1

Numbers of sessions at successive concurrent VR VR for five subjects.

		Subjects				
		414*	456	475	300	468
SERIES 1						
VR 30	VR 30	54	54	54	54	54
VR 21	VR 41	56	56	56	56	56
VR 50	VR 11	85	85	85	85	85
VR 25	VR 35	101	101	101	101	101
VR 30	VR 30		67	67	67	67
SERIES 2						
VR 60	VR 60		93	93	93	93
VR 40	VR 80		68	68	68	68
VR 100	VR 20		26	26	26	26
VR 50	VR 70		40	40	40	40
VR 60	VR 60		74	72	74	74
SERIES 3						
VR 60	VR 60		60	60	60	60
VR 40	VR 80		45	44	45	45
VR 100	VR 20		20	20	20	20
VR 50	VR 70		65	65	64	63

*Subject died before full exposure to the fifth experimental condition.

the Fleshler and Hoffman (1962) progression with $N = 14$ and with time translated into responses. Reinforcement duration was 3 sec for the first series and 2 sec for the second and third. Throughout all series, sessions terminated with the sixtieth reinforcement.

RESULTS

Figure 1 shows the distribution of pecks on the two keys during all three series. Each point plots an average for one subject over the final 10 sessions before a change in conditions. The ordinate is the proportion of pecks on the left key; since pecking on the right key is the complementary proportion, it has not been plotted. The abscissa is a proportion calculated by dividing the probability of reinforcement on the left key by the sum of reinforcement probabilities on both keys. When the probabilities are, for example, equal, the abscissa value is 0.5, and so on. Abscissa values in Figure 1 are based on actual, not nominal, ratio schedules. Sampling errors and the COD (for series 1 and 2) result in minor differences among the subjects, which is why the points do not quite line up over each other.

The diagonal line in Figure 1 traces probability matching. It clearly does not describe any subject's performance in any of the three series, although a few points perch on or near the line here and there.

The step function at 0.5 shows maximizing. To the left of the step, the subject maximizes reinforcements-per-response by avoiding the left key altogether. To the right of the step, maximization calls for all pecks on the left key. Although the subjects fail to take the step sharply, they seem to approximate it. Pecking tends to be exclusively at the left or right key as the abscissa departs from 0.5 in either direction.

Figure 2 shows the approach to exclusive preference directly. The data from Figure 1 are here combined into three curves, one for the average performance in each series. The ordinate is the proportion of responses to the left key minus 0.5 and expressed as an absolute value. Thus, if the left-key proportion is either 0.9 or 0.1, the ordinate would be 0.4. An ordinate value of zero therefore represents no preference, while an ordinate of 0.5 represents exclusive preference for one key or the other. The abscissa measures the difference

Fig. 1. Proportion of responses to the left key as a function of relative probability of reinforcement, for all subjects in the three series. Data points with slashes through them signify second determinations. Where the points cluster, it may be hard to see all the symbols. The diagonal shows probability matching; the step function, maximization of reinforcements per response.

between the relative probability (as defined for Figure 1) and 0.5, again expressed as an absolute number. An abscissa value of zero means that the two probabilities of reinforcement were equal; a value of 0.5 means that one or the other probability (but not both) was zero.

When the ratios summed to 60 (or 61 or 62), exclusive preference was attained with a smaller relative difference between the two ratios than when the sum was 120. A relative difference of 0.15 or thereabouts sufficed, on the average, for the smaller ratios, while a relative difference about twice as great barely sufficed for the larger ratios. But the larger ratios, it should be recalled, had a 2-sec magazine cycle instead of the 3-sec magazine cycle for the smaller ratios, so that there could easily be some confounding of effects here. In any event, given the small number of subjects and the substantial within-series variability (see Figure 1), these differences would doubtless fall well short of statistical significance.

Although the subjects were working on variable-ratio schedules, it is nevertheless possible to treat the results in the customary manner for concurrent variable-interval schedules. Figure 3 (for the three series) plots the proportion of responses on the left key as a function of the proportion of reinforcements

Fig. 2. The absolute value of the deviation from indifference in responding to the two keys as a function of the absolute value of the deviation from equal probabilities of reinforcement. Both variables are expressed as proportions. The subjects in the three experimental series have been averaged for the three curves.

Fig. 3. Each coordinate shows the relative frequency of responding to the left key as a function of its relative frequency of reinforcement, for one of the concurrent VR VR schedules. All three series are shown. The straight diagonals show matching of relative response to relative reinforcement; the curved functions show the loci of possible distributions of responding (see text for explanation). Data points with slashes through them signify second determinations. Where points cluster, it may be hard to see all the symbols.

actually obtained from pecking on the left. Separate coordinates have been drafted for each experimental condition in all three series. Once again, the points show the 10-session averages.

Each coordinate contains the diagonal for matching—not probability matching, but matching in the sense usually applied to interval schedules. In addition, there is a second curve except when the two ratios were equal. This second curve is the locus of all possible relations between the proportion of responses and the proportion of reinforcements, given the variable-ratio schedules on each key. For example, if the variable ratios are 40 and 80 on the left and right key, respectively, then each value of the ordinate implies one and only one value of the abscissa. If the ordinate is 0.5, then the proportion of reinforcements on the left must be 0.667; if it is 0.2, then the proportion of reinforcements on the left must be 0.333. The proportion of reinforcements on the left can be 0.5 only when the proportion of responses on the left had been 0.333.

The possible relations between relative responding and relative reinforcement *must* obey the equation:

$$\frac{\frac{P_L}{(VR)_L}}{\frac{P_L}{(VR)_L} + \frac{P_R}{(VR)_R}} = \frac{R_L}{R_L + R_R} \quad (1)$$

P_L and P_R stand for the numbers of responses to the left and right key; R_L and R_R stand for the corresponding numbers of reinforcements. The two ratio values are given as $(VR)_L$ and $(VR)_R$. Either the relative number of responses or the relative number of reinforcements therefore fixes the other quantity. Here and there, a data point strays from the theoretical curve, but only because the actual ratio values occasionally differed from the nominal values used for Figure 3. No second curve is drawn when the two ratios are equal (*i.e.*, 30, 30 or 60,60) simply because it falls exactly on the matching line itself (see equation 1). To put it another way, when the ratios are equal, any distribution of responses conforms to matching.

The more the two ratios differ, the more the locus of possible values bows out from the matching line. However, no matter how different the ratios are, the curves are tied

Fig. 4. Deviations from matching of relative response to relative reinforcement, for the three experimental series and averaged across subjects.

together at the extremes—at exclusive preference for one key or the other. Thus, a subject can always match (in the variable-interval sense) by confining all its responses to one key. When the ratios are equal or nearly equal, any distribution of responses approximates matching.

Performance in relation to matching is shown in Figure 4, which averages across subjects, but keeps the experimental conditions separate. The average deviation from matching exceeded 5% only once—for VR 50 VR 70 without the COD. Figure 4 shows the left key; the right key would be the complement and is therefore omitted. In all three series, the largest deviation from matching was obtained when the ratios differed, but differed least—i.e., VR 25 VR 35 or VR 50 VR 70. Making the ratios either equal or more different reduces the deviation from matching. There is some suggestion in Figure 4 that the COD holds down the deviations from matching, but the effect is minor at most.

The final bit of results also averages across subjects but keeps the experimental conditions separate. Figure 5 shows the average session durations, minus the time the feeder operated, for each condition. Sessions always terminated at the sixtieth reinforcement, so this duration measures the time taken to earn

a fixed number of reinforcements. The independent variable in Figure 5 is the smaller of the two ratios, since the subjects tended to pick that one when the two were appreciably different. The general trend is clear, although there are a couple of reversals. The smaller the available ratio, the briefer the session. There may be some interaction with the larger ratio, for the sum of the two seems to exert an effect of its own. In general, the smaller sum favored slower responding, as shown by the upward displacement of the curve for the Sum 60 con-

Fig. 5. Session durations (minus reinforcement time) as a function of the smaller variable-ratio schedule in each pair, for the three experimental series and averaged across subjects.

dition. A lesser upward displacement may be attributed to the COD.

DISCUSSION

The pigeons in this experiment failed to match probabilistically, unlike those in experiments reported by Bitterman (1965). However, there are enough procedural differences so that we dare not be too positive about the cause of the discrepancy. It may be noteworthy that probability matching by animals (and to a considerable degree, by human subjects) has been most often observed with probabilities in the vicinity of 0.5 or at 1.0, or in experiments using a correction procedure (Graf, Bullock, and Bitterman, 1964; Sutherland and Mackintosh, 1971). In the vicinity of 0.5, our pigeons too hovered at nonexclusive levels of preference, which, with a large enough N , might have averaged near the ratio of the probabilities.

Maximization, not probability matching, was our main effect. Almost all the pigeons came at least close to exclusive preference for the smaller of a pair of variable-ratio schedules. If this finding surprises anyone, it should be only those who think probability matching is a basic principle of choice, if there are any. Nevertheless, though our finding will be intuitively obvious to most workers, and even though it has been anticipated by some fragmentary earlier data (Herrnstein, 1958), it would be premature to consider the issue closed. The prediction of maximization, and its relation to the matching law, is now more a matter of intuition than explicit theory. It seems useful to make the theory explicit, for it has logical consequences that are not ordinarily pursued empirically.

Maximization in Time

How does a pair of unequal variable-ratio schedules guide behavior to maximization? That is the obvious prime question. Confronted with a pair of unequal ratio schedules, a subject may be influenced by the differing rates of reinforcement the schedules yield. Figure 6 shows how the time to earn a fixed number of reinforcements depends on the relative frequency of responding to the left key, assuming that the rate of pecking remains constant. (Plausible violations of this assumption do not change the general point.) When the

subject shifts toward the more lucrative key, it earns reinforcement sooner; when it shifts the other way, the rate of pay declines. Figure 6 contains examples based on VR 100 VR 20 and VR 40 VR 80, when the rate of pecking is assumed to remain constant while the relative amount of pecking on the left or right key shifted. A constant change in the proportion of reinforcements produces a constant change in the time taken to earn a fixed number of reinforcements. The larger the difference between the ratios, the steeper the line; the steeper the line, the greater the savings gained by shifting toward the smaller ratio.

Fig. 6. Theoretical time to earn a fixed number of reinforcements as a function of the relative frequency of reinforcement on the left key for *conc* VR 100 VR 20 and *conc* VR 40 VR 80. A constant rate of responding has been assumed for this calculation. The ordinate plots the expression $[(VR)_L - (VR)_R]R_L + (VR)_RN$, as a function of $\frac{R_L}{R_L + R_R}$. The expression

equals the time taken to obtain a fixed number of reinforcements, N . N , which equals $R_L + R_R$, has been arbitrarily set at a value of 0.1, and the ordinate reads in the corresponding arbitrary units of time.

Maximization of reinforcements per unit time means sliding down the lines in Figure 6. Equal ratios make the line horizontal, so the behavior should show no tendency to slide one way or the other. Figure 3 showed that equal ratios produced little systematic preference, while maximally different ratios often produced exclusive preference. Figure 5 confirmed that the behavior did, in fact, adjust to

produce an increasing rate of reinforcement as the ratios became more different.

Minimization of Deviations from Matching

Although maximization of reinforcements in time handles much of the data, it seems to leave at least one serious gap. It fails to account for the tendency of the subjects to minimize deviations from matching. Figure 6 should help explain the problem. Suppose the subject is working on a pair of ratios different enough to produce a detectable slope in the line. Once behavior begins to slide toward maximization, there is nothing in the theory to explain why it does not go all the way to exclusive preference.² The pigeons' behavior, however, occasionally stopped short of total preference (see Figures 3 and 4), apparently within the limits set by the typical spread around matching.

In Figure 3, the data points crowd toward the matching line for every pair of variable ratios. When the two ratios differ most, the locus of possible distributions of responding and reinforcement arches sharply away from matching (see Figure 3). Yet, Figure 4 shows that the deviation from matching stayed within the 5 to 10% limits commonly observed on concurrent *interval* schedules. For a given pair of ratios, the deviation from matching depends on the distribution of responding. Figure 7 shows the theoretical deviations from matching for two concurrent ratio schedules (VR 100 VR 20; VR 40 VR 80). The more different the ratios, the longer the possible deviations from matching. However, all concurrent ratio schedules approach matching as preference becomes extreme. For equal or nearly equal ratios, approximate matching is achieved by any distribution of responding. Consequently, subjects may be expected to

vary markedly in how they scatter their responses, as they did (see Figure 3). In contrast, when the ratios pull apart, matching is approximated only at the extreme preferences. Consequently, the subjects then bunched together at one extreme or the other. The matching function seemed to hold responding within a band no thicker than about 15%.

Fig. 7. Theoretical deviations from matching of relative response to relative reinforcement for *conc* VR 100 VR 20 and *conc* VR 40 VR 80. The curves show the value of the expression, $\frac{(VR)_L R_L}{(VR)_L R_L + (VR)_R R_R} - \frac{R_L}{R_L + R_R}$, as a function of $\frac{R_L}{R_L + R_R}$. R_L and R_R are absolute frequencies of reinforcement; $(VR)_L$ and $(VR)_R$ are the values of the ratio schedules for the two alternatives. The curves plot the logically necessary deviations from matching given all distributions of reinforcement within each of two concurrent VR VR.

But matching by itself may not seem to lead uniquely to maximizing with concurrent ratio schedules, for the subject can match by choosing the larger ratio exclusively, instead of the smaller ratio. While Figure 3 shows the data points drawing close to the matching line, there is clear evidence that the pigeons favor the smaller ratio of unequal ratios. They approximate matching advantageously. That they do so raises the possibility that matching and maximizing may be dual aspects of a single process, which is the process of reinforcement itself. It is axiomatic that, given incompatible responses that differ in reinforcement and given that the difference has been detected, a subject will choose the more

²We should note that this problem could not have been avoided by defining maximization reciprocally, as a minimization of time per reinforcement. The reciprocal of the ordinate in Figure 6 produces curves instead of lines, but the curves are bent just the wrong way to explain the data. If the subjects were minimizing time per reinforcement, it would be necessary to account for the peculiar fact that the approach to the minimum sometimes comes to rest just where the function has the highest slope. As before, the subjects' tendency to approximate matching short of exclusive preference remains unexplained.

highly reinforced alternative.³ No other axiom makes sense, for to "differ in reinforcement" means only that one response is stronger, and "stronger" means preferred, whether one thinks of preference as fundamental to response strength (Herrnstein, 1970) or as derivative (Nevin, 1974). In the present experiment, the choice of the smaller ratio is simultaneously an instance of both matching and maximizing, within the framework of the single process of reinforcement.

Momentary Maximizing

Some workers (Shimp, 1966, 1969; Silberberg and Williams, 1974) see a potential opposition between matching and maximizing, not a complementarity as suggested here. They look for, and, in some cases, find, response sequences at the level of individual responses that seem to be guided by optimization at each successive moment—"momentary maximizing", it is called. The matching on schedules like concurrent variable intervals is said to arise from optimizing reinforcement with each response. However, it is said that circumstances can be arranged in which responses individually optimized add up to nonmatching. Matching is thus said to be derivative, not fundamental. Some form of local maximizing is said to be fundamental.

We can give our objection to this theory by first developing a hypothetical example (see Table 2) of how momentary maximizing may lead to matching. Consider two concurrent response alternatives, A and B. Suppose that following a B, the probability of reinforcement for an A has been set at 1.0 and for a B, at 0. Following an A, the probabilities of reinforcement are reversed. Suppose the subject optimizes each response, as demanded by momentary maximizing. At the end of a session, there must be an equal number of As and Bs, since optimization means strict alternation. Because there must also be equal reinforcement from each alternative, the subject will have matched.

³For simplicity, this discussion assumes that the response alternatives are neutral or equivalent in inherent reinforcing value, in analogy with pigeons pecking at two or more standard keys. Clearly, responses in nature often have different inherent reinforcing value. Maximization would then involve both inherent and extrinsic variables. The calculation is more complex, but the principles are the same as far as we know.

Table 2
Probability of Reinforcement

	FOLLOWING B	FOLLOWING A
A	1.0	0
B	0	1.0

But now consider the contingencies in Table 3. Reinforcement probability for B following A is set at 0.5 instead of 1.0, otherwise the procedure is the same as for Table 2. Suppose the subject again optimizes each response. The result once again is strict alternation. However, the number of reinforcements can no longer be the same. At the end of a session, the subject will have received

Table 3
Probability of Reinforcement

	FOLLOWING B	FOLLOWING A
A	1.0	0
B	0	0.5

twice as many reinforcements from A as from B. By optimizing here, the subject *cannot* match responses and reinforcements over an entire session.

In the relevant experiments and analyses (Shimp, 1966, 1967, 1969; Silberberg and Williams, 1974), the response-reinforcement contingencies are not as simple as in Tables 2 and 3. However, the logic of the argument is preserved by our example. In short, the theory says that a subject matches when local contingencies dictate, and not otherwise.

To explain our uneasiness with this line of argument, let us add one further change to the hypothetical experiments. Suppose a B response turns on a green light that lasts until the next A response, and an A response turns on a white light that lasts until the next B response. At every moment, the light is then either green or white, depending on what the last response was. If it is green, then the probability of reinforcement for A is 1.0 and for B is 0; if it is white, then the probability for A is 0 and for B it is either 1.0 (Table 2) or 0.5 (Table 3). Adding the stimulus produces a multiple schedule with a response-dependent alternation of components. The components are concurrent ratio-extinction schedules. The subject, choosing between a ratio schedule and extinction, is optimizing when it chooses the reinforcement schedule. It is also matching within each component. Within each component, optimization means matching by exclusive preference because the alternatives are tantamount to unequal ratio schedules. No one expects responses summed over components to match reinforcements summed over components.

Adding a stimulus to the hypothetical procedures does not reduce the generality of our objection to momentary maximizing as an alternative to matching, for the objection centers on inhomogeneity in the schedule of reinforcement: if a subject's behavior tracks an inhomogeneity in reinforcement probability, then its behavior is under stimulus control, whether the stimulus is exteroceptive or interoceptive. Regular sequential shifts in response show that the procedure has in effect generated a multiple schedule with shifting contingencies of reinforcement. The proper test of matching is *within* each of the components of the effective multiple schedule. Matching describes the distribution of behav-

ior under effectively homogeneous conditions of reinforcement, not across differentiated conditions of reinforcement.

Any form of momentary maximizing that shows up as a response strategy can be subsumed under matching by a suitable *post hoc* alignment of reinforcement contingencies. It is possible to decompose optimization into choices between higher and lower probabilities of reinforcement, which are tantamount to concurrent ratio schedules. Precisely the same logic that allows one to say that a response strategy is optimal also allows one to say that responding conforms to matching within each of those concurrent schedules. In real experiments, the probabilities may sometimes be obscure because of the latitude in analysis. There is latitude in how we divide up the reinforcement contingencies and in defining the response, for topographies can combine into sequences that are reinforced as a whole (Herrnstein, 1958). But a response strategy cannot be evidence against matching, for it is simultaneously optimal and matching on a concurrent ratio schedule. Experiments that show momentary maximizing are valuable, not as evidence against matching, but as data bearing on a subject's sensitivity to differential contingencies of reinforcement.

If the evidence suggested that matching always involved sequential response strategies, it would be another matter. While such a finding would not contradict or disprove matching, it might suggest that response-by-response optimizing was primary and matching was secondary. But the evidence shows the reverse. There is no consistent evidence of response-by-response optimizing in ordinary variable-interval schedules (Anger, 1956; Catania and Reynolds, 1968) and there is evidence *against* such molecular response processes in several experiments designed to study them (Nevin, 1969; Reynolds, 1963), though the matching law applies. In principle, matching on concurrent interval schedules might have been reduced to momentary maximizing, by showing that the behavior followed the shifting probabilities of reinforcement that these schedules produce. But, in fact, no one has succeeded in formulating such an analysis that explains the data. The overall matching on concurrent interval schedules remains unreduced, a molar phenomenon that is here replicated for concur-

rent ratio schedules. Whether the concurrent ratio schedule is explicitly arranged, as it is here, or inferred, as it has been in some analyses of concurrent interval schedules (Shimp, 1966; Silberberg and Williams, 1974), matching and maximization are just two words describing one outcome.

REFERENCES

- Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 1956, **44**, 145-161.
- Bitterman, M. E. Phyletic differences in learning. *American Psychologist*, 1965, **20**, 396-410.
- Brunswik, E. Probability as a determiner of rat behavior. *Journal of Experimental Psychology*, 1939, **25**, 175-197.
- Catania, A. C. and Reynolds, G. S. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 327-383.
- Estes, W. K. Probability learning. In A. W. Melton (Ed.), *Categories of human learning*. New York: Academic Press, 1964. Pp. 89-128.
- Fleshler, M. and Hoffman, H. S. A progression for generating variable interval schedules. *Journal of the Experimental Analysis of Behavior*, 1962, **5**, 529-530.
- Graf, V., Bullock, D. H., and Bitterman, M. E. Further experiments on probability-matching in the pigeon. *Journal of the Experimental Analysis of Behavior*, 1964, **7**, 151-157.
- Herrnstein, R. J. Some factors influencing behavior in a two-response situation. *Transactions of the New York Academy of Sciences*, 1958, **21**, 35-45.
- Herrnstein, R. J. On the law of effect. *Journal of the Experimental Analysis of Behavior*, 1970, **13**, 243-266.
- Hilgard, E. R. and Bower, G. H. *Theories of learning*. 3rd ed.; New York: Appleton-Century-Crofts, 1966.
- Nevin, J. A. Interval reinforcement of choice behavior in discrete trials. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 875-885.
- Nevin, J. A. On the form of the relation between response rates in a multiple schedule. *Journal of the Experimental Analysis of Behavior*, 1974, **21**, 237-248.
- Reynolds, G. S. On some determinants of choice in pigeons. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 53-59.
- Shimp, C. P. Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 443-455.
- Shimp, C. P. Reinforcement of least-frequent sequences of choices. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 57-65.
- Shimp, C. P. Optimal behavior in free-operant experiments. *Psychological Review*, 1969, **76**, 97-112.
- Silberberg, A. and Williams, D. R. Choice behavior on discrete trials: a demonstration of the occurrence of a response strategy. *Journal of the Experimental Analysis of Behavior*, 1974, **21**, 315-322.
- Sutherland, N. S. and Mackintosh, N. J. *Mechanisms of animal discrimination learning*. New York: Academic Press, 1971.

Received 29 October 1974.

(Final Acceptance 19 March 1975.)