

DISSOCIATION OF VALUE AND RESPONSE STRENGTH

WILLIAM VAUGHAN, JR.

HARVARD UNIVERSITY

Four pigeons were exposed to multiple schedules and later to concurrent-chains schedules, with terminal links that had previously been multiple-schedule components. For 2 birds, the terminal-link schedules arranged an inverse relationship between response rate and reinforcement rate; for the other 2 birds a direct corresponding relationship was arranged. Those response rates were further modified by differentially reinforcing either longer or shorter interresponse times, relative to the current means. Although the birds' initial-link responses indicated preferences for terminal links with higher rates of reinforcement, in half the cases the birds responded during the terminal links in such a way as to produce lower rates of reinforcement, rates their initial-link behavior indicated they did not prefer. That outcome is inconsistent with maximization theory, but consistent with a strengthening analysis of behavior on single-key schedules.

Key words: maximization, multiple schedules, concurrent chains, interresponse times, key peck, pigeon

Given single-key variable-interval (VI) and variable-ratio (VR) schedules, each providing the same rate of reinforcement, it is well established that responding is usually faster on the VR than on the VI schedule (Catania, Matthews, Silverman, & Yohalem, 1977; Ferster & Skinner, 1957; Zuriff, 1970). Two distinct accounts have been offered for this difference between response rates (Baum, 1973; Morse, 1966). One, which may be termed a *local* account, is based on the fact that on a VI schedule longer interresponse times (IRTs) have a higher probability of being reinforced than do shorter IRTs, whereas on a VR schedule all IRTs have the same probability of being reinforced.

Figure 1 (top left) shows the distance advanced by a VI 60-s schedule as a function of IRT, the time (in seconds) from the previous to the current peck. The convention followed here is that the average distance a schedule must advance for reinforcement to be made available is one unit. Thus, on a VI 60-s schedule an IRT of 1 s will accompany a schedule advance of 1/60, as shown (neglecting the fact that conventionally, schedule timing is suspended whenever a reinforcer is set up). Longer or shorter IRTs advance it longer or shorter

distances, respectively. On a VR schedule, the distance advanced at each response is independent of the time preceding that response. Figure 1 (top right) shows a VR 60. Each response, regardless of the time between it and the previous response, advances the schedule a distance of 1/60.

The second account of the difference between VI and VR response rates may be termed a *global* account. On a VI schedule, at very low rates of responding an increase in response rate produces a substantial increase in reinforcement rate, whereas at higher rates of responding the same increase produces a smaller increase in reinforcement rate (Figure 1, bottom left). On a VR schedule, on the other hand, regardless of the current rate of responding, an increase in response rate will produce the same increase in reinforcement rate (Figure 1, bottom right). Maximization theorists have formalized the global account further by pointing out that indifference curves, as shown in Figure 2 (left), can account qualitatively for the VI versus VR rate difference (see Green, Kagel, & Battalio, 1987, for a recent discussion). Superimposed on the VI and VR schedules shown in Figure 2 (left) are hypothetical indifference curves of trade-off between response rate and reinforcement rate. They are called curves of constant value, and may be thought of as analogous to contour curves on a topographical map. Each curve represents combinations of rate of responding and rate of reinforcement that will be equally valued. The

I thank Jeffrey Davis for participating in a number of pilot studies, David Coe for discussions as well as for running animals, and M. Davison and R. J. Herrnstein for comments. Requests for reprints should be sent to William Vaughan, Jr., Department of Psychology, Harvard University, Cambridge, Massachusetts 02138.

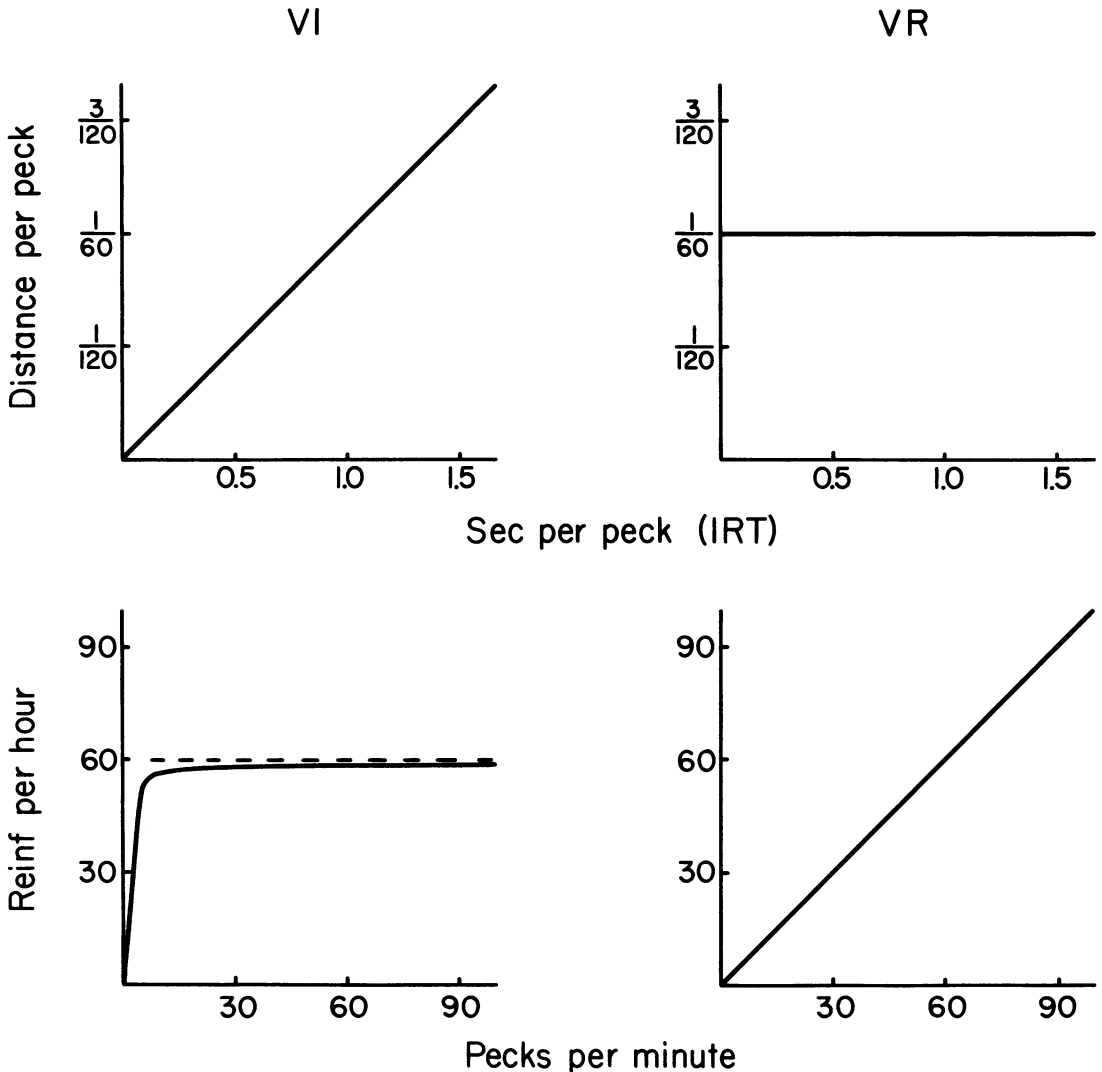


Fig. 1. Two representations of VI schedules (left panels) and VR schedules (right panels). Top: Distance advanced by the schedule, as a function of time per peck, or IRT. Bottom: Rate of reinforcement as a function of rate of responding.

numbers next to each curve show their respective hypothetical values. The hypothetical data point on the VI function (labeled "a") is the place at which that schedule is tangent to an indifference curve, one having a value of 30 units. Any other rate of responding on the VI schedule will produce less value (its corresponding point lying on a different indifference curve, lower on the graph), and so a maximization analysis predicts that behavior will fall at the point shown. The data point on the VR schedule (labeled "b") is likewise the place at which that schedule is tangent to an indif-

ference curve, one having a value of 20 units. Once again any other rate of responding would result in less value, and so should not be chosen.

Suppose, by means of the local schedules, it were possible to generate two distinct rates of responding on one global schedule (e.g., during two components of a multiple schedule), as shown by the two data points (labeled "c" and "d") in Figure 2 (middle). Because both points would represent combinations that were possible under *both* of the local schedules, a maximization analysis would then imply they were

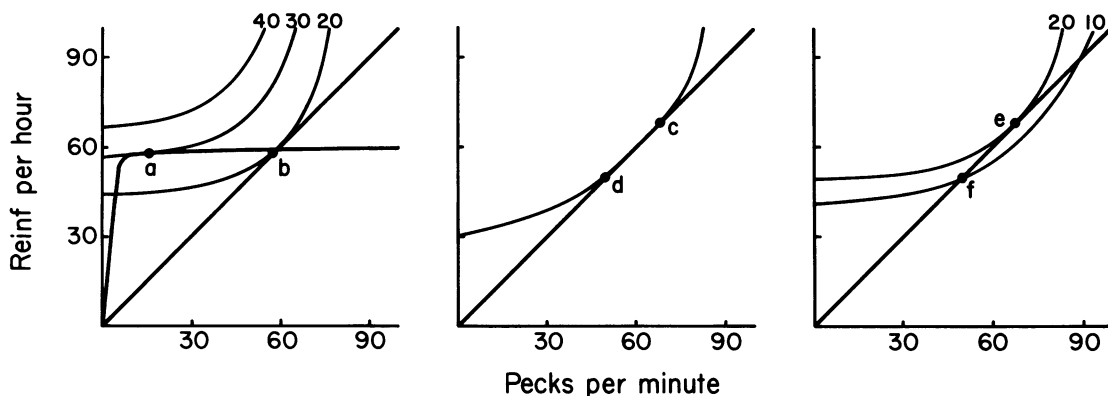


Fig. 2. Left: Hypothetical indifference curves, showing that with equal obtained rates of reinforcement, response rate on a VI schedule (a) should be less than on a VR schedule (b). Middle: Given two distinct rates of responding and reinforcement on a VR schedule (c and d), a maximization analysis requires they fall on the same indifference curve. Right: Hypothetical example in which point e lies on a higher indifference curve than point f, an outcome precluded by maximization. The numbers identifying successively higher indifference curves have only ordinal significance.

both on the same indifference curve, as shown in the figure. If they were on two indifference curves having different values, as shown in Figure 2 (right, labeled "e" and "f"), there is then no account, in terms of maximization, for why point "f" on the lower curve would not migrate to the indifference curve where "e" is located, because that curve represents a higher valued combination of response rate and reinforcement rate.

In the present experiment, an attempt was made to generate two rates of responding on one global schedule, and then to show, using a concurrent-chains procedure, that one of those rates was preferred over the other. Such an outcome would correspond to Figure 2 (right) and so constitute an inconsistency for maximization, because the unpreferred response rate, with its associated rate of reinforcement, could not be accounted for by the theory.

METHOD

Subjects. The subjects were four White Carneau pigeons, all of whom had prior experimental histories. They were all maintained at approximately 80% of their free-feeding weights.

Apparatus. A standard pigeon chamber, 30 cm wide, 33 cm high, and 30 cm long, was used. On the front panel were mounted three response keys, of which only the two outer keys were used. Those two were 2 cm in diameter

and 12.7 cm apart, center-to-center. Directly below the center key was a standard hopper for delivery of mixed grain; reinforcement consisted of 2.5-s access to the grain. The experiment was controlled by means of a Digital PDP-8® computer in conjunction with SuperSKED® software, with some standard relays as part of the interface.

Procedure for Conditions 1 and 2. There were seven conditions altogether, the first two involving multiple schedules and the remainder involving concurrent chains, with the components of the multiple schedule during Condition 2 serving as terminal links during the subsequent conditions. During Conditions 1 and 2, a four-ply multiple schedule was in effect. Two of the components (accompanied by green and amber) were on the left key, and two (accompanied by red and blue) were on the right key. The duration of each component was 2 min, and there was a 30-s period between components during which the keys were darkened and the houselight remained on. The order of components was random, with the constraint that each consecutive set of four components (e.g., 1 through 4, 5 through 8, and so on) contained one of each of the four colors. It was thus possible for a component to be repeated if it occurred at the end of a set of four components. When 40 reinforcers had been received within a session, the session was terminated when the current set of four components ended. The locations of the schedules in the final set of each session were stored and

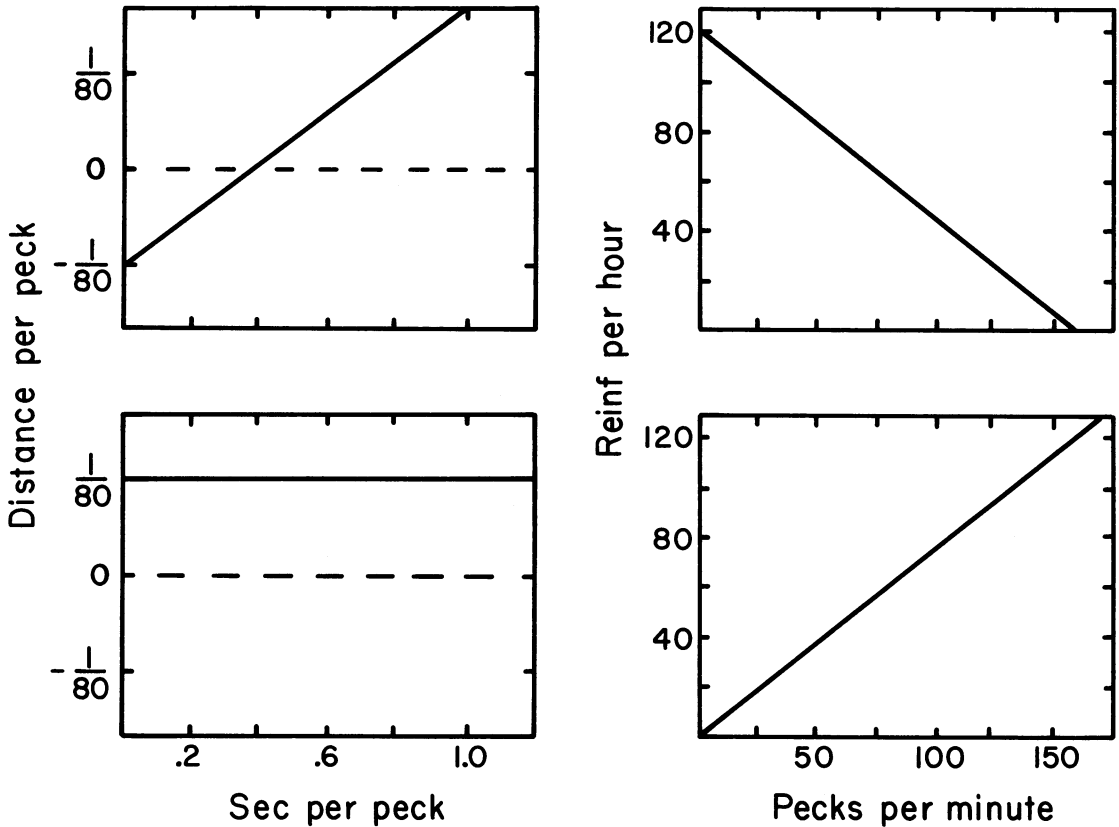


Fig. 3. Schematic diagrams showing characteristics of schedules in effect during Condition 1. Left: Distance that a schedule advanced as a function of IRT length (top: Birds 1 and 2; bottom: Birds 3 and 4). Right: Corresponding functions showing reinforcement rate as a function of response rate (top: Birds 1 and 2; bottom: Birds 3 and 4).

used to start the next session. On a multiple schedule, the reinforcement rate during one component can affect the response rate during another component, a phenomenon termed contrast (Reynolds, 1961). The random ordering, the 30-s period between components, and the long components were employed so as to minimize such effects. Williams (1981) has shown these effects to depend on factors such as short component durations that should have had little effect given the current procedure.

The schedules in effect during Condition 1 are shown in Figure 3. Birds 1 and 2 were exposed to the schedule shown (in two different ways) in the top two panels. The function in the top left panel shows how far the schedule advanced toward reinforcement as a function of each IRT. As mentioned above, the average distance to each reinforcer was one unit. There were 16 such distances, and their distribution was that specified by Fleshler and Hoffman

(1962). As can be seen in the panel, if the bird emitted a response, say, 0.2 s following the previous response, the schedule would move backward about $1/160$ th (actually 0.0058) of a unit. An IRT of 0.375 s would have no effect on the schedule, and longer IRTs would move it forward by a distance that was a linear function of the IRT duration. The function shown in the top left panel of Figure 3 is linear and has a slope of 0.033 (that is, $0.0125/0.375$) and an intercept of -0.0125 (or $-1/80$). IRTs were measured with a resolution of 0.1 s.

At each response, the distance the schedule should advance (as determined by the function in the top left panel of Figure 3) was compared with the remaining requirement for reinforcement. If the former was negative, or positive but less than the latter, the schedule was adjusted by the appropriate amount, with no other effect. If the distance to be moved was positive and equal to or greater than the current re-

quirement, reinforcement was delivered. Any remaining distance, up to the requirement for the next reinforcer, was also effective in advancing the schedule. For example, suppose the current requirement was $1/80$ th of a unit, and an IRT of 0.6 s was emitted. The top left panel of Figure 3 shows that the distance to be advanced is approximately $1/160$ (actually 0.0075), so the requirement for reinforcement would be reduced by that amount. Suppose the current requirement was $1/80$, and an IRT of 1 s was emitted. The distance to be advanced is almost $2/80$, so reinforcement would be delivered and up to almost $1/80$ would be subtracted from the next requirement. It was thus possible for a single response to produce reinforcement and to set up conditions for the next reinforcement, but nothing more.

The function shown in the top right panel of Figure 3 shows rate of reinforcement as a function of rate of responding, given the function in the top left panel. Increasing rates of responding should result in linearly decreasing rates of reinforcement, until, at 160 pecks per minute (equivalent to an average IRT of 0.375 s), reinforcement rate should fall to zero (i.e., on average the schedule would not advance). As can be inferred from the top left panel, still higher rates of responding would move the schedule backwards. At very low rates of responding, rather than the function precisely intersecting the ordinate at 120 reinforcements per hour (as shown in the figure, due to the scale) there was a maximum around two pecks per minute, and lower rates of reinforcement would result from yet lower rates of responding. This latter aspect follows logically from the fact that each reinforcer is response produced, and hence if response rate is very low, reinforcement rate must likewise be very low. In subsequent discussion here, the kind of function shown in the top left panel will be called an *IRT function*, and the kind of function in the top right panel a *rate function*. An IRT function may apply to individual IRTs, to average IRTs, or to both. A rate function will always be applied to average rates of responding and reinforcement.

Birds 3 and 4 were run on the schedule shown in the bottom two panels of Figure 3. The bottom left panel shows how far the schedule would be advanced by each response, which in this case is $1/80$ th of a unit regardless of the IRT. In other words, the schedule in effect

was a VR 80. The function in the lower right panel shows rate of reinforcement as a function of rate of responding on this schedule.

To summarize, during Condition 1, Birds 1 and 2 were exposed to a 4-ply multiple schedule. In each component, IRTs longer or shorter than 0.375 s decremented or incremented, respectively, the distance to reinforcement. Birds 3 and 4 were likewise on a 4-ply multiple schedule, with each component being a VR 80, but with colors changing from component to component (green or amber on the left key, red or blue on the right key). Each bird was shifted to Condition 2 when its response rates appeared stable from day to day.

During Condition 1, the distance moved for each IRT was the same as the average distance moved for an average IRT of the same duration. For Birds 1 and 2, for example, an IRT 0.2 s in duration would move the schedule backwards about $1/160$ th of a unit, and if a number of IRTs averaged 0.2 s the average distance moved would be backwards the same distance. Consistent with what was said earlier, that first aspect of a schedule (distance per IRT) is called its local property, and the second aspect (average distance per average IRT, or reinforcement rate as a function of response rate in the case of a rate function) is called its global property. During Condition 2, each global schedule of Condition 1 remained in effect, but the local schedules were modified, as shown in Figure 4. For Birds 1 and 2, the local function shown in the top left panel of Figure 4 was superimposed on one of the original global schedules (accompanied by green or amber) on the left key and on one on the right (accompanied by red or blue). Similarly, the local function shown in the top middle panel was superimposed on the other original global schedule on the left key, as well as on the right.

To see how this was done, consider the top right panel in Figure 4. The global function (from Figure 3, top left) is the solid line with a positive slope, while the local function (with a steep negative slope) is from the middle panel. (This global function is an IRT function that applies to average IRTs only, whereas this local function is an IRT function that applies to individual IRTs only.) At each response, two quantities were calculated: the IRT of that response (represented by C for current), and the mean IRT of the last 10 responses, in-

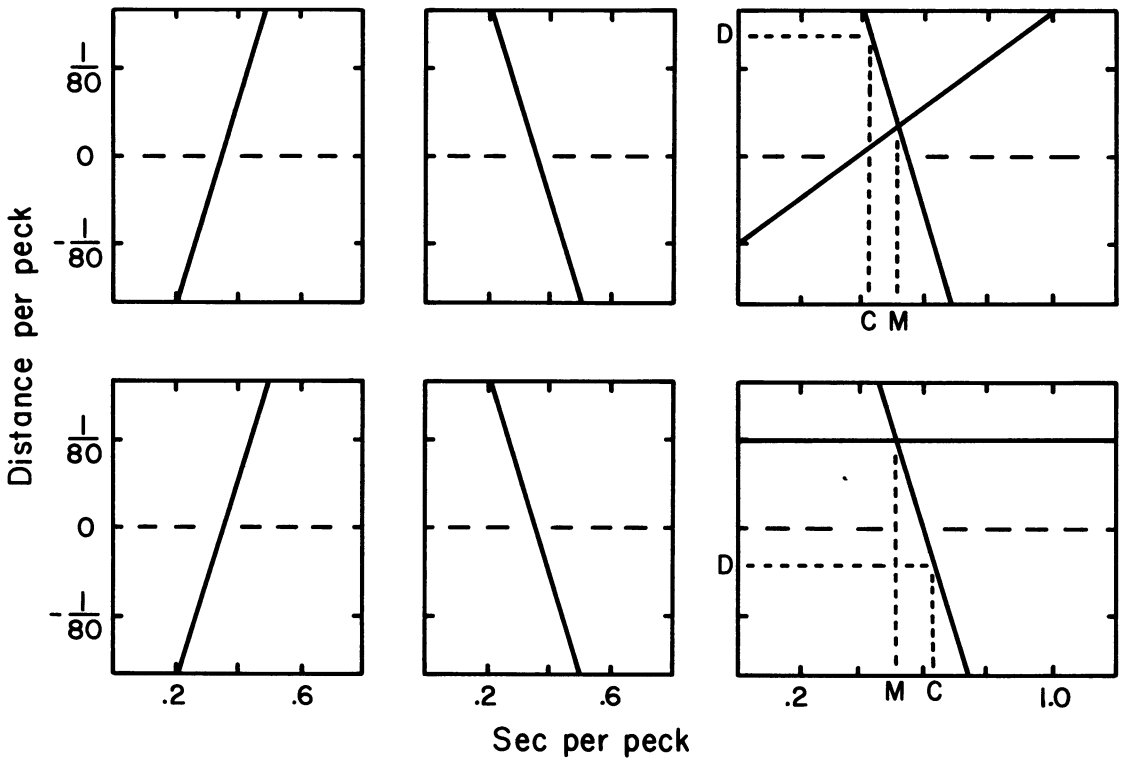


Fig. 4. Local schedules in effect during Conditions 2 through 7. The three top panels are for Birds 1 and 2, and the bottom panels are for Birds 3 and 4. Left panels: Local function reinforcing longer IRTs more than shorter IRTs. Middle panels: Local function reinforcing shorter IRTs more than longer IRTs. Right panels: Representation of superposition of local schedules shown in middle panels on global IRT schedules in effect. See text for details.

cluding the last (represented by M for mean). The local function was then (algebraically) placed so as to intersect the global function just above M , the mean of the last 10 IRTs. Then the distance the schedule should advance, D , was determined by C , the duration of the current IRT, in conjunction with the superimposed local function.

Suppose responding were approximately stable when averaged over 10 responses. Shorter IRTs would then have a higher probability of producing reinforcement than longer IRTs (i.e., they would drive the schedule farther than longer IRTs would), as specified by the local schedule. In addition, because on average the deviation of IRTs about their mean is zero, on average the global schedule would be satisfied: In terms of the rate function, the higher the average rate of responding the lower the average rate of reinforcement (Figure 3, top right panel). In other words, there are now two schedules in effect. Shorter IRTs have a higher probability of reinforcement than longer

IRTs (the local property), and higher rates of responding, on average, produce lower rates of reinforcement (the global property). Were the average rate of responding to change, shorter IRTs would still drive the schedule farther than longer IRTs, but the respective distances would change. They would both increase if the average response rate decreased and decrease if the average response rate increased.

For Birds 3 and 4, the same local schedules (Figure 4, bottom left and middle panels) were superimposed on the original global schedule, the VR 80 (Figure 4, bottom right panel). In the case shown in the figure, rather than every response driving the schedule the same distance (as was the case previously), IRTs longer than the mean of the last 10 IRTs drove the schedule a shorter distance than $1/80$, and IRTs shorter than the mean drove it a longer distance. On average, however, every response continued to drive the schedule a distance of $1/80$, and so on average the global VR 80 was

in effect. Overall rate of reinforcement remained proportional to response rate, but IRTs shorter than the current mean advanced the schedule farther than did IRTs longer than the current mean. On a traditional VR schedule (e.g., Condition 1, Birds 3 and 4) overall rate of reinforcement is proportional to response rate, and all IRTs advance the schedule the same distance.

To implement these schedules, it was necessary for the computer program to perform a number of calculations following each response. Although somewhat cumbersome within SuperSKED®, the use of a FOR . . . NEXT loop allowed the shifting of the last 10 IRTs within a buffer, as well as the calculation of their sum and a number of other operations, within a single clock tick (10 ms). The nominal time to finish all calculations lasted a maximum of 20 ms. Because the same number of calculations was performed for the two kinds of local schedules, there should have been no differences in reinforcement latencies between the two. During Condition 2, as well as during later conditions, both the location of the schedules and the last 10 IRTs of a session were stored and carried over to the next session. The colors that accompanied the two local schedules were counterbalanced across the two pairs of birds. Condition 2 was in effect for a minimum of 50 sessions, in case the effect of the superimposed local schedules was merely transitory.

Results for Conditions 1 and 2. Figure 5 shows the mean response rates and reinforcement rates for the last five sessions of Condition 1 (filled shapes) and of Condition 2 (open shapes), as well as the programmed rate functions. For Birds 1 and 2 (top two panels) the open triangles represent performances on the left-key and right-key schedules that reinforced longer IRTs more than shorter IRTs (the values of the individual schedules are shown in Table 1 below). The filled triangles represent performances on the same schedules during Condition 1. In these cases, then, when longer IRTs (relative to Condition 1) had higher probabilities of reinforcement than shorter ones, response rate fell and reinforcement rate increased. Performances on the other two schedules are represented by circles (filled for Condition 1, open for Condition 2). In these cases, when shorter IRTs came to have higher probabilities of reinforcement than longer ones,

response rate increased and reinforcement rate fell. For Birds 3 and 4, the open triangles represent performances on the left- and right-key schedules that reinforced shorter IRTs more than longer IRTs during Condition 2, whereas the filled triangles represent performances on the same schedules during Condition 1. The circles represent performances on the other two schedules. None of the data points deviates substantially from the programmed rate function.

Procedure for Conditions 3 through 7. During Conditions 3 through 7, concurrent-chains schedules were in effect. The initial link consisted of interdependent (Stubbs & Pliskoff, 1969) conc VI 30-s VI 30-s schedules with a 1-s changeover delay (COD). The left and right keys were white, and the schedules only advanced within 2 s of a response on either key. If the schedules had stopped, only a response at least 1 s after they restarted could produce a transition to the terminal link. Interdependent schedules were used to prevent unequal numbers of entries to the terminal links, an outcome that could in turn affect preference. The COD was employed to ensure separation between the two concurrent schedules.

When one terminal link was produced, the other key was darkened and responses on it had no effect. The two terminal links consisted of one of the two schedules that had been arranged on the left key during Condition 2 and one of the two schedules that had been on the right during Condition 2. Terminal links remained in effect for 2 min (excluding any reinforcer time), at which time the initial links were again reinstated. Fantino and Herrnstein (1968) have shown that duration of a terminal link, independent of rate of reinforcement there, can affect preference for that terminal link, so a constant duration was used to help ensure that only rate of responding and rate of reinforcement would affect preference. When 40 reinforcers had been received, the session was terminated as soon as the end of the current terminal link was reached.

During Condition 3, both the left and right terminal links were those schedules with lower rates of reinforcement (the schedules represented by the open circles in Figure 5). In Condition 4, the right terminal link was unchanged, but the left terminal link was the left schedule with the higher rate of reinforcement.

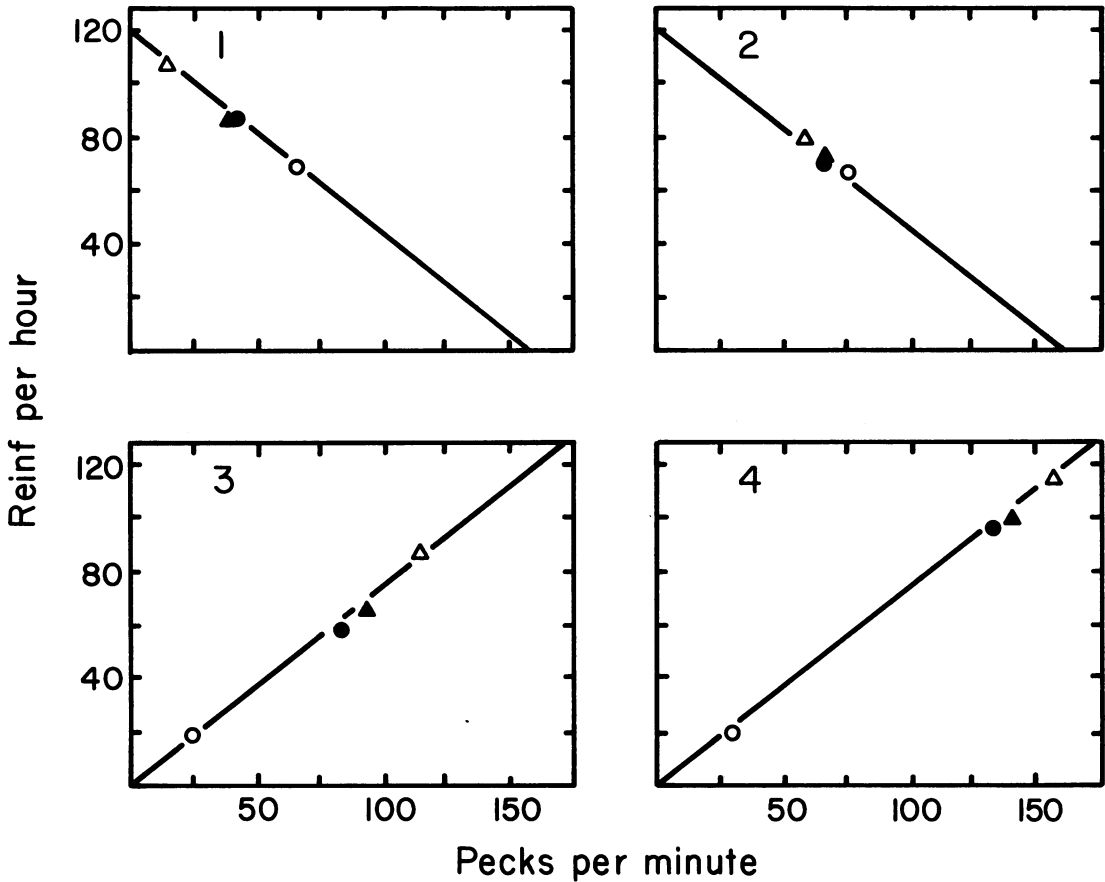


Fig. 5. Summary of results from Condition 1 (filled shapes) and Condition 2 (open shapes). The solid-line functions are the programmed rate schedules in effect. The lower rates of responding (open triangles for Birds 1 and 2, open circles for Birds 3 and 4) correspond to schedules on which the probabilities of reinforcement were higher for longer IRTs than for shorter IRTs. The higher rates of responding (open circles for Birds 1 and 2, open triangles for Birds 3 and 4) correspond to schedules with the converse relationship: The probabilities of reinforcement were higher for shorter IRTs than for longer IRTs.

In Condition 5, both terminal links were the schedules with higher rates of reinforcement (open triangles in Figure 5), in Condition 6 the left terminal link was the schedule with the lower rate of reinforcement (and the right terminal link was unchanged), and in Condition 7 both terminal links were the schedules with the lower rates of reinforcement (replication of Condition 3). Each bird was shifted to the next condition when its initial-link choices appeared stable. The minimum number of sessions for any bird was 18.

RESULTS

Table 1 shows, for all seven conditions, response rates on the four keys (exclusive of the

choice keys), reinforcement rates, colors of the keys, and the local schedule in effect. For the latter, "0" indicates no superimposed local schedule, "+" indicates a local schedule that should result in a higher rate of reinforcement, and "-" indicates a local schedule that should result in a lower rate of reinforcement.

During Conditions 3 through 7, components that had been part of the multiple schedule during Condition 2 became terminal links of the concurrent-chains schedules. In spite of that change, for the most part there were few substantial changes in average rate of responding for any particular schedule, suggesting the context of those schedules did not significantly affect behavior during them. Bird 2 did show a decrease in response rate within amber be-

Table 1
Color of keys, local schedules, response rates and reinforcement rates for all conditions.

Condition	Left key				Right key			
	Color	Schedule	Pecks per minute	Reinforcers per hour	Color	Schedule	Pecks per minute	Reinforcers per hour
Bird 1								
1	Green	0 ^a	46.3	82.5	Red	0	36.7	91.5
	Amber	0	46.8	84.0	Blue	0	35.2	90.0
2	Green	+ ^b	12.0	107.5	Red	+	16.4	103.0
	Amber	- ^c	69.6	64.0	Blue	-	62.1	76.5
3	Amber	-	76.9	60.5	Blue	-	73.0	64.2
4	Green	+	12.8	104.7	Blue	-	76.4	59.8
5	Green	+	12.6	96.1	Red	+	13.1	109.9
6	Amber	-	69.9	64.0	Red	+	13.2	114.4
7	Amber	-	66.2	65.3	Blue	-	66.7	80.0
Bird 2								
1	Green	0	60.9	82.8	Red	0	72.8	64.8
	Amber	0	57.5	72.0	Blue	0	75.4	68.4
2	Amber	+	60.5	78.1	Blue	+	58.3	80.4
	Green	-	72.8	70.4	Red	-	77.6	64.8
3	Green	-	71.4	66.5	Red	-	83.0	56.9
4	Amber	+	46.0	86.3	Red	-	72.6	62.1
5	Amber	+	39.5	88.6	Blue	+	59.4	75.7
6	Green	-	70.7	73.5	Blue	+	35.5	82.0
7	Green	-	70.4	78.8	Red	-	74.8	66.1
Bird 3								
1	Green	0	82.2	51.4	Red	0	85.0	63.6
	Amber	0	85.1	54.4	Blue	0	103.1	78.8
2	Amber	+	104.2	80.6	Blue	+	127.0	96.4
	Green	-	27.8	22.8	Red	-	24.4	16.0
3	Green	-	20.7	16.1	Red	-	19.5	14.8
4	Amber	+	123.1	91.0	Red	-	19.5	13.9
5	Amber	+	114.4	87.5	Blue	+	128.1	94.6
6	Green	-	17.6	14.2	Blue	+	113.1	88.1
7	Green	-	16.6	11.0	Red	-	20.1	15.6
Bird 4								
1	Green	0	135.4	92.0	Red	0	148.6	108.0
	Amber	0	131.6	92.0	Blue	0	139.4	104.5
2	Green	+	141.9	100.8	Red	+	174.4	128.4
	Amber	-	28.0	19.2	Blue	-	30.0	19.2
3	Amber	-	18.0	13.4	Blue	-	23.6	17.7
4	Green	+	139.6	105.3	Blue	-	21.8	15.2
5	Green	+	132.3	100.4	Red	+	178.4	136.1
6	Amber	-	16.4	7.5	Red	+	174.9	137.0
7	Amber	-	20.2	17.1	Blue	-	14.0	16.7

^a 0 = no superimposed local schedule.

^b + = local schedule producing higher rate of reinforcement.

^c - = local schedule producing lower rate of reinforcement.

tween Condition 2 and the time amber appeared as a terminal link (Condition 4), but there had been no decrease upon moving from Condition 1 to 2 in amber, even though longer IRTs came to have higher probabilities of reinforcement than shorter IRTs, relative to Condition 1. It is thus possible that the continued exposure to the superimposed local schedule was responsible for the eventual shift

in response rate. This general lack of change in response rate between Condition 2 and subsequent conditions, in conjunction with the presumed minimization of contrast during Conditions 1 and 2, suggests that in terms of maximization (or any other comparable theory) the schedules may be construed as equivalent to single-key schedules.

As pointed out above, the existence of two

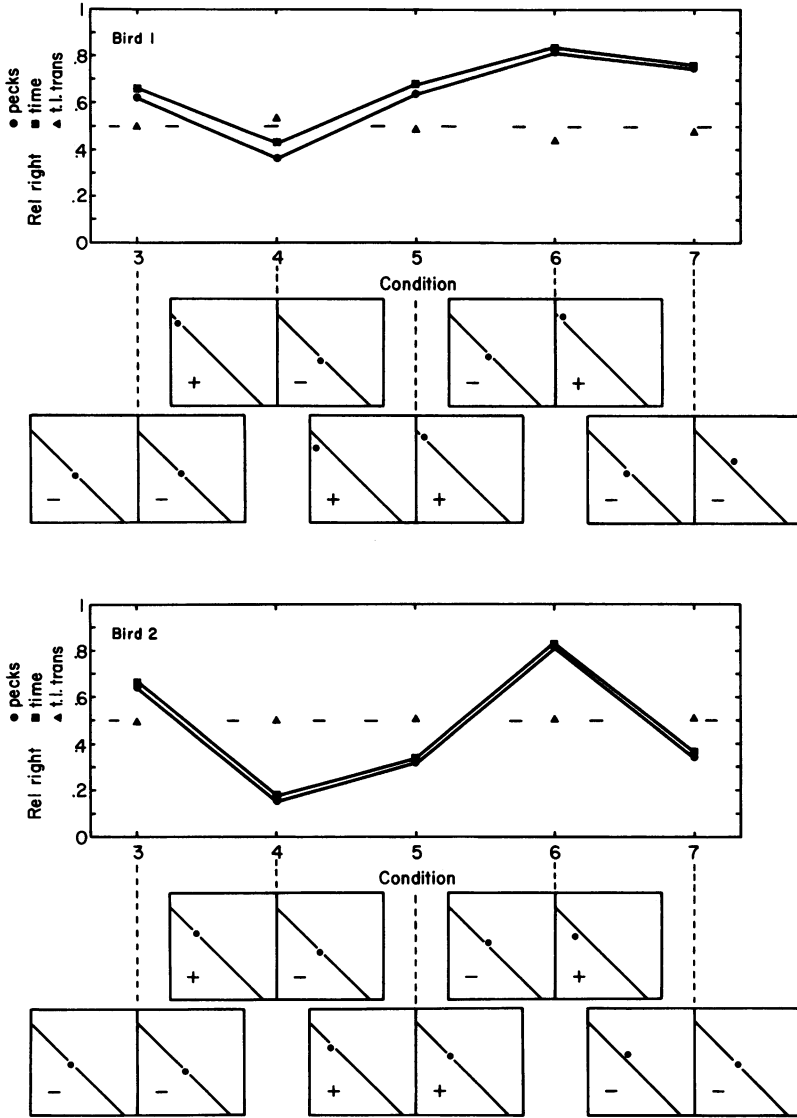


Fig. 6. Relative right pecks, time, and terminal-link transitions for Birds 1 (top) and 2 (bottom), as well as a representation of terminal-link response rates (abscissae) and reinforcement rates (ordinates), for Conditions 3 through 7.

response rates on one global schedule is not necessarily inconsistent with a maximization analysis, because it is logically possible that the two response rates were approximately equally preferred. In that case, the effect of the different local schedules would have been to shift behavior between points of equal value, possibly along lines of equal value. To derive an inconsistency with maximization, it is necessary to show that the birds preferred one of the two response rates over the other. Figure 6 shows, for Birds 1 and 2, relative pecks, times, and terminal-link entries during the ini-

tial links (top panel for each bird), as well as a schematic representation of the associated terminal-link rate functions (bottom five pairs of panels for each bird) for Conditions 3 through 7. The left and right terminal links are represented by the left and right part of each pair, respectively. As in Table 1, “+” indicates a schedule that should result in a higher rate of reinforcement, and “-” indicates a schedule that should result in a lower rate of reinforcement.

During Condition 3, Bird 1 responded at approximately the same rate during the two

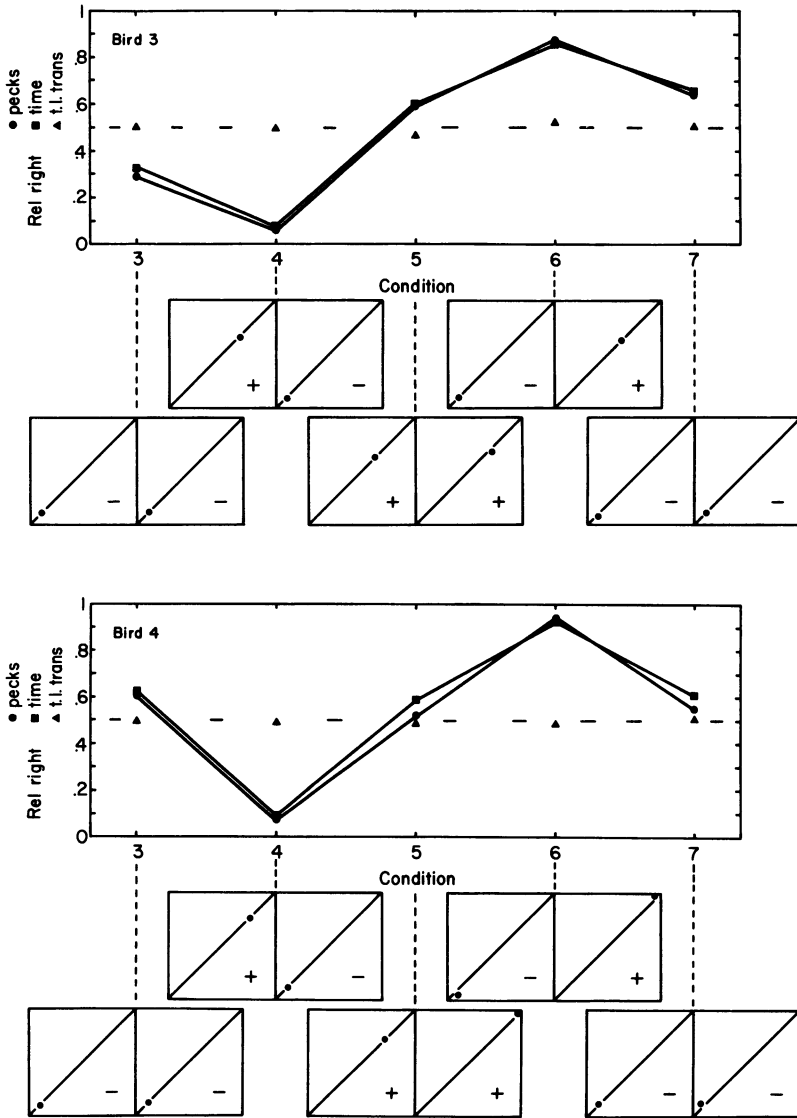


Fig. 7. Relative right pecks, time, and terminal-link transitions for Birds 3 (top) and 4 (bottom), as well as a representation of terminal-link response rates (abscissae) and reinforcement rates (ordinates), for Conditions 3 through 7.

terminal links and received similar rates of reinforcement, as shown in Figure 6 in the boxes directly below Condition 3. Initial-link behavior showed some preference for the right. During Condition 4, the left terminal-link response rate was lower, and that terminal-link reinforcement rate was higher; performance during the right terminal link remained fairly constant. The bird's initial-link choices showed a shift to the left, which indicates that the left terminal link during Condition 4 was more preferred than the left terminal link during Condition 3. That leaves unexplained, in terms

of maximization, why it responded at such a high rate during the left terminal link in Condition 3. In other words, during Condition 3 while on the left terminal link, the strength of responding was distinct from that which would have maximized value. It is not necessarily the case that value was maximized during Condition 4 on the left, only that it was greater during Condition 4 than during Condition 3.

The same general pattern occurs upon moving from Condition 4 to 5. In this case, reinforcement rate during the left terminal link remained approximately the same, but in-




Fig. 8. Log ratio (base e) of time allocated to the right over time allocated to the left during initial links, as a function of log ratio of reinforcement rate during the right terminal link over reinforcement rate during the left terminal link. Least-squares functions are shown, along with the proportions of variance accounted for.

creased on the right. In addition, preference for the right likewise increased. Again, there is no explanation in terms of maximization why Bird 1 responded at the higher rate on the right during both Conditions 3 and 4. The same general pattern holds for Bird 2, except in this case the differences in rates of reinforcement were not as large. For both birds, terminal-link entries remained approximately equal on each side, presumably as a result of the interdependent schedules used in the initial links.

Figure 7 shows the comparable data for Birds 3 and 4. For Bird 3, for example, the two terminal-link rates of reinforcement during Condition 3 were low, and there was some preference for the left. During Condition 4, reinforcement rate on the left was higher than previously, and the preference for the left increased. This shift shows the bird preferred the higher rate of responding with its correlated higher rate of reinforcement, so there is no explanation in terms of maximization why it responded at the low rate on the left during

Condition 3. Bird 4 showed both less deviation from indifference with nominally equal terminal links and the largest preference with unequal terminal links of all 4 birds. Again, terminal-link entries remained approximately equal.

Although Figures 6 and 7 show an increase in preference, during the initial link, for the terminal link with the higher rate of reinforcement, it is difficult to derive a quantitative functional relation directly from those figures. Figure 8 gives that information, showing the log ratio of times (to the base e) during the initial links as a function of the log ratio of obtained rates of reinforcement during the terminal links. In each case the slope of the least-squares function is positive, indicating that higher rates of reinforcement were preferred. For both birds with negative-slope terminal-link rate functions (i.e., higher reinforcement rates for lower response rates) the slopes of the functions shown in this figure are higher than for the other 2 birds. This may possibly reflect a greater increase in value as reinforcement rate increased if response rate decreased rather than increased. If so, this is one of the few demonstrations that on aperiodic schedules responding per se may affect value (see Baum & Nevin, 1981, for a discussion of some of the relevant literature).

Although Bird 2 showed little change in response rates between Condition 1 and later conditions, it showed the largest effect on preference during the initial link as terminal-link rates of reinforcement changed. Bird 4 exhibited the behavior closest to matching between initial-link distributions of time and terminal-link rates of reinforcement, with virtually no bias or undermatching.

Figure 8 is not intended to imply that on concurrent-chains schedules generalized matching occurs between initial-link responding and terminal-link reinforcement rates (see, e.g., Fantino, 1969), only that under the conditions studied here such a relation approximately held.

DISCUSSION

Figure 8 shows that all 4 birds preferred those combinations of responding and reinforcement that provided higher rates of reinforcement. Nevertheless, the local schedules were sufficient to induce them to respond in a manner that their initial-link behavior indi-

cated they did not prefer. Although the results are inconsistent with a maximization analysis of single-key schedules, they are qualitatively consistent with the assumption that reinforcement of different IRTs can affect the distribution of future IRTs (as revealed by overall rate of responding). The results are thus in accord with analyses of the VI VR response-rate difference that appeal to strengthening of IRTs (e.g., Peele, Casey, & Silberberg, 1984).

A maximization analysis would have predicted zero slopes for the functions shown in Figure 8, because in all cases the birds should have been indifferent between the two response rates with their correlated rates of reinforcement (see Figure 2 above). The use of equal and interdependent initial links would, if anything, have strengthened that tendency, because there was no advantage in spending more time on one side as opposed to the other.

The use of both negative- and positive-slope terminal-link rate functions in the present experiment serves to anticipate certain objections that might otherwise be brought against the interpretation given here. Suppose only positive-slope (i.e., VR) rate schedules had been used. It might then be argued that the preference for the higher rate of reinforcement derived from reinforcement being more immediate upon entry into a terminal link with higher rates of responding, because each reinforcer requires a response and responses occur more rapidly. A second possibility would have been that, through generalization, the higher terminal-link response rate would have induced a higher initial-link response rate on that side. The use of negative-slope rate schedules precludes those interpretations, because in that case preference was shown for the terminal link with the lower rate of responding. If only negative-slope rate schedules had been used, it would be unclear whether the birds preferred lower rates of responding or higher rates of reinforcement. The use of positive-slope terminal-link rate functions shows that high reinforcement rates in conjunction with high response rates were preferred over both rates being low.

Figure 3 shows that the separation between the high and low response rates was greater for the positive-slope rate functions (Birds 3 and 4) than for the negative-slope cases. In the latter case, as response rate fell reinforcement rate rose, possibly strengthening responding and hence retarding a decrease in responding

to some extent. As response rate increased reinforcement rate decreased, possibly retarding an increase in responding to some extent. In the case of positive-slope rate functions, the opposite global contingencies were in effect. These differences in separation provide additional evidence for a strengthening account.

All 4 birds showed some bias toward the right, possibly reflecting some color bias. DeLius and Emmerton (1979) mention a strong bias for blue shown by pigeons. Also, for Birds 2 and 3 in the present experiment, the results of Condition 7 failed to replicate those of Condition 1. For Bird 2 this did not prevent all data points from falling near a linear least-squares function, whereas for Bird 3 there was more spread about that function. In spite of these details, however, the main results are unambiguous.

To date there have been numerous experiments showing that matching, but not maximization, can account for behavior on concurrent schedules (see Heyman & Herrnstein, 1986, for a review of concurrent VI VR experiments). In contrast, although Vaughan and Miller (1984) and Davison and Charman (1987) provided evidence against a maximization analysis on single-key schedules, several recent experiments (Green et al., 1987; McDowell & Wixted, 1986) have presented arguments in favor of a maximization analysis of single-key schedules, and in favor of a global analysis, respectively.

Each of the latter two experiments, however, contains design flaws that limit the conclusions that can be drawn from them. McDowell and Wixted (1986) exposed human subjects to five VR schedules, as well as to five VI schedules, modified so that overall reinforcement rate was proportional to response rate. The latter schedules were designed to have the global properties of VR schedules and the local properties of VI schedules (as shown in Figure 1, above). Response rates on the two kinds of schedules were very similar and were higher than those found in an earlier experiment in which humans were exposed to simple VI schedules (McDowell & Wood, 1984).

If McDowell and Wixted (1986) had exposed their subjects to simple VI schedules as well, and found lower rates of responding for comparable rates of reinforcement, their conclusion would have been much stronger. In the absence of that control, they were forced to compare response rates across subjects, in dif-

ferent experiments, using different force requirements. The human subjects in the McDowell and Wixted experiment who responded faster had a force requirement of 120 N, whereas the subjects in the McDowell and Wood (1984) experiment had a requirement of 150 N (and in one case 152 N). Finally, even if it does turn out that humans are sensitive to the global properties of feedback schedules, the same may not hold for other organisms.

Green et al. (1987) exposed birds first to random-ratio (RR) schedules and then to RR schedules that provided half the food as previously, in conjunction with a variable-time (VT) schedule that made up the difference in total amount of food. On the combined schedules, the birds responded more slowly than on the simple RR schedules. Green et al. argued that the results were consistent with maximization and inconsistent with a response-strengthening account. Some of their arguments, however, are not entirely compelling. They say (p. 26) that a strengthening account might be able to account for the decrease in response rate with the combined schedules, but not for the simultaneous decrease in reinforcement rate. Surely that latter decrease follows from the fact that half their reinforcers were obtained on an RR schedule, and a decrease in response rate would have to produce a decrease in reinforcement rate. Perhaps more telling is the fact that earlier work directly contradicts their indifference curves (Vaughan, 1982; Vaughan & Miller, 1984), work they fail to confront.

SUMMARY

A maximization analysis of behavior on single-key schedules implicitly assumes that the local schedules do not affect behavior, and that animals choose that point on the global schedule that maximizes value. In the present experiment, it was possible to show not only that different local schedules were able to produce different rates of responding on one global schedule, but that one of those rates of responding, with its correlated rate of reinforcement, was preferred over the other. The experiment thus goes one step beyond that of Vaughan and Miller (1984), who showed that on negative-slope global schedules birds responded faster than necessary and obtained fewer reinforcements than necessary, but did

not actually show that the obtained combination was less valued than some other available combination.

The present experiment does not directly address the VI VR response-rate difference, but it does directly address the mechanisms that have been proposed to account for that difference. It is worth noting, however, that the schedules used here suggest that VI and VR schedules are only two points on a continuum rather than two distinct categories. The space from which the present schedules were drawn includes all linear functions, both with respect to the global schedule and independently with respect to the local schedule. It may well prove more illuminating to explore that space than to continue focusing on traditional VI and VR schedules of reinforcement.

REFERENCES

- Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, *20*, 137-153.
- Baum, W. M., & Nevin, J. A. (1981). Maximization theory: Some empirical problems. *Behavioral and Brain Sciences*, *4*, 389-390.
- Catania, A. C., Matthews, T. J., Silverman, P. J., & Yohalem, R. (1977). Yoked variable-ratio and variable-interval responding in pigeons. *Journal of the Experimental Analysis of Behavior*, *28*, 155-161.
- Davison, M., & Charman, L. (1987). Effects of response-allocation constraints on multiple-schedule performance. *Journal of the Experimental Analysis of Behavior*, *47*, 29-39.
- Delius, J. D., & Emmerton, J. (1979). Visual performance of pigeons. In A. M. Granda & J. H. Maxwell (Eds.), *Neural mechanisms of behavior in the pigeon* (pp. 51-70). New York: Plenum Press.
- Fantino, E. (1969). Choice and rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, *12*, 723-730.
- Fantino, E., & Herrnstein, R. J. (1968). Secondary reinforcement and number of primary reinforcements. *Journal of the Experimental Analysis of Behavior*, *11*, 9-14.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.
- Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, *5*, 529-530.
- Green, L., Kagel, J. H., & Battalio, R. C. (1987). Consumption-leisure tradeoffs in pigeons: Effects of changing marginal wage rates by varying amount of reinforcement. *Journal of the Experimental Analysis of Behavior*, *47*, 17-28.
- Heyman, G. M., & Herrnstein, R. J. (1986). More on concurrent interval-ratio schedules: A replication and review. *Journal of the Experimental Analysis of Behavior*, *46*, 331-351.
- McDowell, J. J., & Wixted, J. T. (1986). Variable-ratio schedules as variable-interval schedules with linear feedback loops. *Journal of the Experimental Analysis of Behavior*, *46*, 315-329.
- 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.
- Morse, W. H. (1966). Intermittent reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 52-108). New York: Appleton-Century-Crofts.
- Peele, D. B., Casey, J., & Silberberg, A. (1984). Primacy of interresponse-time reinforcement in accounting for rate differences under variable-ratio and variable-interval schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 149-167.
- Reynolds, G. S. (1961). Behavioral contrast. *Journal of the Experimental Analysis of Behavior*, *4*, 57-71.
- Stubbs, D. A., & Pliskoff, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, *12*, 887-895.
- Vaughan, W., Jr. (1982). Choice and the Rescorla-Wagner model. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 2, Matching and maximizing accounts* (pp. 263-279). Cambridge, MA: Ballinger.
- Vaughan, W., Jr., & Miller, H. L., Jr. (1984). Optimization versus response-strength accounts of behavior. *Journal of the Experimental Analysis of Behavior*, *42*, 337-348.
- Williams, B. A. (1981). The following schedule of reinforcement as a fundamental determinant of steady state contrast in multiple schedules. *Journal of the Experimental Analysis of Behavior*, *35*, 293-310.
- Zuriff, G. E. (1970). A comparison of variable-ratio and variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, *13*, 369-374.

Received February 23, 1987
Final acceptance July 17, 1987