

*INTERRESPONSE-TIME SENSITIVITY DURING
DISCRETE-TRIAL AND FREE-OPERANT
CONCURRENT VARIABLE-INTERVAL SCHEDULES*

J. M. CLEAVELAND

DUKE UNIVERSITY

Two experiments investigated the sensitivity of pigeons' choice to elapsed time since the last response (i.e., to interresponse time [IRT]) during concurrent variable-interval variable-interval schedules. Experiment 1 used a two-key discrete-trial procedure with variable intertrial intervals. Experiment 2 employed a three-key free-operant procedure. In both experiments choice was found to be a function of the active-schedule IRT, defined as the time since the most recent response. Monte Carlo simulations show how this finding permits the joining of several seemingly incompatible data sets held to both support and contradict a kind of choice strategy, termed *momentary maximizing*, which attempts to maximize momentary reinforcement probabilities. The studies suggest that only two variables are needed to describe the static molecular structure of concurrent variable-interval choice: active-schedule IRTs and "response states" consisting of the last one or two schedule choices.

Key words: interresponse time, momentary maximizing, Markov chains, timing, variable-interval schedules, pigeons

Molar performance of animals under concurrent schedules of reinforcement is accurately described by the matching law and its derivatives (Baum, 1974, 1979; Herrnstein, 1961; Staddon, 1977). Dynamic models of choice, although certainly constrained by the findings that underpin the matching law, require a more molecular analysis of the variables that influence choice. For example, on a molar level an animal might allocate its responding equally between two alternatives that provide equal rates of reinforcement. How such a result is constructed on a moment-to-moment basis, however, may vary. A subject could alternate responses to the two alternatives, respond in long, though equal, runs at each alternative, or produce a mixture of these two response patterns. A molar analysis, such as the matching law, will not distinguish among these different response patterns nor among the mechanisms that might produce them. Dynamic models of concurrent choice thus necessitate a precise characterization of molecular variables.

The molecular structure of choice is not purely stochastic. Using concurrent variable-interval (VI) VI schedules of reinforcement,

experimenters have described sequential dependencies in an animal's moment-to-moment responding, both in actual responses (Heyman, 1979; Shimp, 1966, 1969; Silberberg, Hamilton, Ziriaux, & Casey, 1978) and in the allocation of temporal blocks (Hinson & Staddon, 1983a, 1983b; Silberberg & Ziriaux, 1985; Staddon, Hinson, & Kram, 1981). In light of these findings a molecular theory of choice, termed *momentary maximizing* (Shimp, 1966, 1969), hill climbing (Hinson & Staddon, 1983a), and molecular maximizing (Silberberg & Ziriaux, 1985), has been formulated. For the sake of clarity I will refer to all of these by the common name *momentary maximizing*.

Momentary maximization hypothesizes that choice is determined by the single variable of reinforcement probability. An animal said to be following such a strategy in a concurrent choice experiment would simply select the choice with the highest momentary probability of reinforcement. Under constant probability VI schedules (e.g., those produced by Fleshler & Hoffman, 1962) such probabilities are given by the equation

$$p_i = 1 - e^{-\lambda_i t_i}, \quad (1)$$

where p_i , the probability of reinforcement at choice i , depends upon the time, t_i , since choosing choice i and the average reinforcement rate, λ_i , assigned to choice i . In order to apply Equation 1 to concurrent choice exper-

I thank Armando Machado and John Staddon for suggestions on earlier drafts. The work was supported by grants from the NSF and NIMH to J. E. R. Staddon.

Correspondence should be addressed to John Mark Cleaveland, Department of Psychology: Experimental, Duke University, Durham, North Carolina 27708-0086 (E-mail: mark@psych.duke.edu).

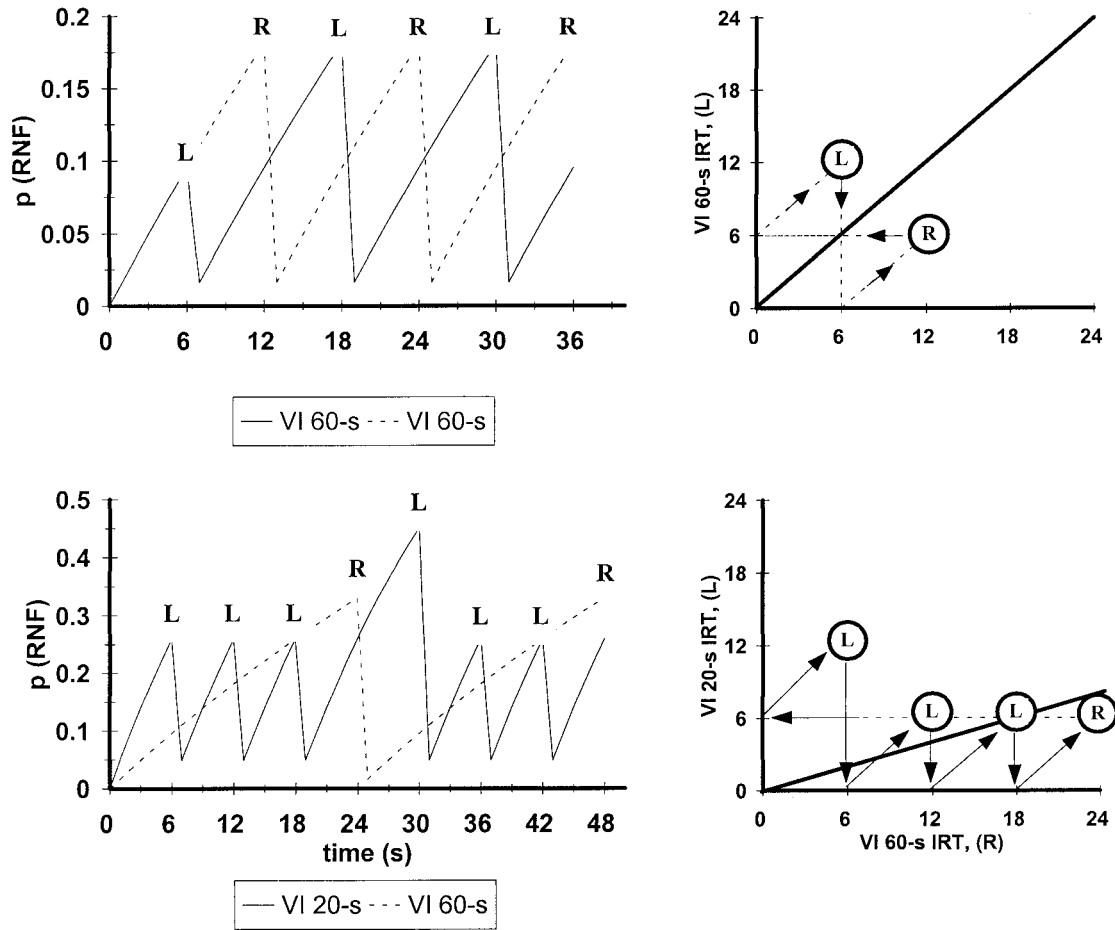


Fig. 1. The left half plots the reinforcement probabilities for each choice in a hypothetical concurrent VI 60-s VI 60-s and concurrent VI 20-s VI 60-s procedure. Reinforcement probabilities were calculated using Equation 1. The plots assume that a response occurs every 6 s as in, for example, a discrete-trial procedure and that the animal selects the choice with the momentarily highest probability of reinforcement. The right half illustrates how such response streams may be represented in a clock space. First, an indifference line serves to divide the clock space into regions of relative reinforcement value. Above this line the probability of reinforcement [$p(\text{rnf})$] is higher for Choice L. Below this line the probability of reinforcement is higher for Choice R. The temporal trajectory of each response is shown by the arrows. A given response serves to reset the IRT of the responded-to schedule to zero, and then a subsequent 6-s IRT “moves” the response diagonally. Assuming regular responding, the clock spaces show the emergent, stable response pattern for an animal following a momentary maximizing strategy. Only when relative IRT values move a response across the indifference line should the animal switch to the alternate schedule of reinforcement.

iments, one simply iterates Equation 1 for each choice. Figure 1, for instance, plots the probability of reinforcement for each choice in a concurrent VI 60-s VI 60-s schedule and a concurrent VI 60-s VI 20-s schedule. For simplicity the plots assume a response every 6 s as in, say, a discrete-trial procedure. Figure 1 shows that as a subject continues to respond at one alternative, which I will term the *active schedule*, the probability of reinforcement

grows at the nonselected or *background schedule*.¹

The differences in reinforcement probabil-

¹ The terms *active schedule* and *background schedule* are perhaps best thought of in relation to a Findley procedure (Findley, 1958). In such a procedure only a single schedule is ever available for responding. In the terminology of this paper, the available schedule is the active schedule, and the nonavailable schedule is the background schedule.

ity associated with the active versus the background schedule predict, according to a momentary maximizing strategy, specific molecular choice structures. These response patterns may be easily visualized if we define an indifference line at which the probabilities of reinforcement for the two choices are equal (Staddon et al., 1981). Setting p_1 equal to p_2 and simplifying yields

$$t_1 = t_2 \left(\frac{\lambda_2}{\lambda_1} \right), \quad (2)$$

where t_i is the time since the last response to Choice 1 or Choice 2 (i.e., the interresponse time or IRT), and λ_i is the programmed reinforcement rate for Choice 1 or 2. Note that the slope of the indifference line in the resultant *clock space* (a plot with time on both axes, in this case IRTs) is given by the ratio of the programmed schedule values, λ_2/λ_1 . Assuming a constant tempo of responding, this ratio directly gives the response pattern predicted by momentary maximizing. Figure 1 illustrates this fact using a concurrent VI 60-s VI 60-s and a concurrent VI 20-s VI 60-s schedule of reinforcement. Under a concurrent VI 60-s VI 60-s schedule of reinforcement, a subject should alternate its responding between the two alternatives ($\lambda_2/\lambda_1 = 1$). Conversely, under a concurrent VI 20-s VI 60-s schedule of reinforcement, the subject should respond three times at the rich (VI 20-s) alternative, followed by a single response to the lean (VI 60-s) alternative ($\lambda_2/\lambda_1 = 1/3$).

Momentary maximizing, however, predicts regularly repeated response patterns only if the subject responds at a constant tempo. Shimp (1966) and Silberberg et al. (1978) observed, for example, that although the momentary maximizing sequence was most probable, pigeons emitted a variety of response sequences under concurrent VI schedules of reinforcement. One might conclude from these findings that the behavior of pigeons is only imperfectly sensitive to reinforcement probabilities (e.g., Shimp, 1969, Experiment 3) or that another molecular process, perseveration, competes with a sensitivity to reinforcement probabilities (Silberberg et al., 1978). A simpler explanation is that pigeons rarely respond at a constant tempo.

Given variance in IRTs, no single response

sequence may be termed optimal. Thus, the analytical emphasis must shift from sequences to the molecular structure of IRT distributions (Shimp, 1969, 1973). A clock space is one method of representing a distribution of IRTs, and a clock space with the addition of an indifference line permits a qualitative estimate of whether a subject's behavior is correlated with momentary reinforcement probabilities. Figure 1 suggests, and indeed it can be mathematically proven (Hinson & Staddon, 1983a), that in a clock space that plots concurrent VI VI IRTs, the choice represented by the nearest axis, without crossing the indifference line, has the higher probability of reinforcement. Thus in Figure 1, the y-axis schedule (e.g., VI 20 s) will have the higher reinforcement probability at all points to the left of the indifference line, whereas the x-axis schedule will have the higher reinforcement probability at all points to the right of the indifference line.

A clock space can provide a visual estimate of whether behavior conforms to a momentary maximizing strategy. However, a more quantitative estimate of the degree to which a subject's IRT distributions conform to a momentary maximizing strategy is captured by the single statistic, M (Hinson & Staddon, 1983a). This statistic is given by

$$M = \frac{\sum \text{correct}}{\sum \text{correct} + \sum \text{incorrect}}, \quad (3)$$

where correct = $|p_i - p_j|$ if choice i when $p_i > p_j$ or choice j when $p_j > p_i$, and incorrect = $|p_i - p_j|$ if choice i when $p_j > p_i$ or choice j when $p_i > p_j$. M is in essence a proportion of reinforcement probability differences obtained at each choice. At every response the momentary probabilities of reinforcement for all choices are calculated and subtracted from one another. If the just-made choice was to the alternative with the highest probability of reinforcement, then the absolute value of this difference is added to the numerator and the denominator. Otherwise the difference is summed solely to the denominator. If a subject's behavior tracks momentary reinforcement probabilities, M will approach 1.0. Behavior that is indifferent to momentary reinforcement probabilities will produce M values of .5, and suboptimal behavior will produce M values less than .5. Hinson and Stad-

don (1983a, 1983b) found that pigeons in free-operant concurrent VI VI situations tended to produce M values of .5 early in training and values between .7 and .9 late in training, supporting the momentary maximizing hypothesis.

However, the data in support of momentary maximizing are seriously challenged by a single, well-established result. Namely, researchers typically find no correlation between switch probabilities and previous responding (de Villiers, 1977; Heyman, 1979; Nevin, 1969, 1979). That is, the probability of switching out of one schedule in a concurrent VI VI experiment appears to be independent of the number of pecks made to that schedule. To understand this result more thoroughly, let us refer to the bottom portion of Figure 1. As responses accrue to the active schedule (VI 20 s), the reinforcement probability at the background schedule (VI 60 s) grows. Momentary maximizing predicts that as the background IRT increases, the animal should eventually switch to that schedule. Allowing for choice "errors," one might not expect an immediate switch, but a positive correlation between background IRTs and switch probabilities should still be evident. No experimental support for this prediction has been published.

As was noted earlier, though, the crucial variable for momentary maximizing is not responses (e.g., run length) but is rather the distribution of interresponse times. Neither Nevin (1969) nor Heyman (1979) distinguish between these two variables. The experiments reported in this paper attempt to answer the simple question of whether the behavior of pigeons is indeed sensitive to IRTs during concurrent VI schedules. The first experiment utilized a discrete-trial procedure in which the intertrial intervals, and hence the IRTs, were variable, but were under experimental control. The second experiment used a free-operant procedure that permitted the pigeons to generate their own IRT distributions.

EXPERIMENT 1

Nevin (1969), using a discrete-trial procedure and concurrent VI VI schedules, reported that the probability of a switch remained flat as run length increased. This

result, however, does not address the issue of whether responding is sensitive to IRT. First, as noted above, run length, the dependent variable analyzed by Nevin, is a response-based variable. However, even in a discrete-trial procedure, which attempts to hold IRT constant, response times are somewhat variable because subjects must be given time to respond. For example, after a 6-s intertrial interval (ITI) Nevin's birds were given 2 s to respond. The first response in this response window initiated the next ITI (or reinforcement delivery). Thus, in reality, IRTs in Nevin's experiment ranged from 6 s to 8 s, assuming a response within each trial.

Another difficulty in using Nevin's (1969) results to rule out a momentary maximizing strategy concerns the experimental control exerted over IRTs. In a two-choice experiment each alternative has its own IRT (i.e., t_i from Equation 1). That is, there is the IRT associated with the active (i.e., most recently responded to) schedule and the IRT associated with the background schedule. In a typical discrete-trial procedure, then, the experimenter fixes the active-schedule IRT, and the animal is permitted to vary the background IRT (by varying run lengths). In Figure 1, for example, run lengths of one and three responses produce IRTs at the background schedule of 6 s and 18 s, respectively, but the active-schedule IRT is always 6 s. At most, Nevin's result suggests that a pigeon's responding is not sensitive to the IRT of the background schedule. However, Nevin's result is silent regarding active-schedule IRTs.

Experiment 1, therefore, attempted to extend Nevin's (1969) experiment by varying the IRTs emitted to both the background and active schedules. The experiment utilized a modified discrete-trial procedure in which ITIs were allowed to vary according to a Gaussian distribution. That is, most ITIs were 6 s, as in Nevin's experiment, but occasionally the subjects experienced shorter or longer ITIs.

METHOD

Subjects

Six White Carneau pigeons were used. None of these birds had previously experienced VI schedules, although all had been recently used in transitive inference tasks. All

subjects were maintained at 85% of their free-feeding weights during the experiments. The pigeons were fed a mixture of mixed grain and pellets (Purina ProGrain® and Purina Nutriblend® Green and Gold).

Apparatus

All daily sessions were conducted in a single two-key operant chamber. The chamber measured 37 by 31 cm at the base and 31 cm high. The floor was wire mesh. All other sides (except that holding the keys and food hopper) were made of Plexiglas. The front, metal wall held two keys. An opening (4 cm by 5 cm) at the wall's midline and 8 cm from the chamber floor allowed access to a food hopper that contained mixed grain. When the hopper was activated, a 7.5-W bulb illuminated it. The two keys were centered on the front panel 21 cm from the base of the chamber and 3 cm from each other. A force of 0.2 N on either key operated a microswitch. Behind each key was mounted a 6.3-V Grason-Stadler stimulus projector. These were used to project colors onto the keys. Another 7.5-W bulb (the houselight) served to illuminate the chamber during experimental sessions. The experiments were controlled and the data recorded by a 286 Compaq® computer.

Procedure

Training. The birds were placed individually in the experimental chamber with the food hopper up and the houselight on until they readily ate the mixed grain from the hopper. During all remaining sessions, the houselight was always on except during reinforcer deliveries.

Each bird was next given two 40-reinforcer sessions of forced alternation. During this training both keys were lit red and green (across sessions each key had a single color). A peck to the key different from the preceding choice (i.e., alternation) turned off the keylights and produced 4-s access to the hopper. After the hopper timed out, the keys were reilluminated.

After these two sessions, each bird was given three sessions of discrete-trial forced alternation. A trial began with the illumination of both schedule keys. If no peck occurred within 2 s or if the bird repeated its last response, then the keys darkened for a period of 6 s (the ITI). Pecks during the ITI reset the ITI

timer. A peck to a lit key that was different from the location of the preceding trial response (i.e., alternation) turned off the keylights and produced access to the hopper. Following a reinforcer delivery an ITI of 6 s ensued as described above. The three discrete-trial forced alternation sessions lasted for 40, 80, and 80 reinforcers, respectively.

Test. After training the birds were divided into two groups. Group 1 (Subjects 1348, 1371, and 4600) experienced a discrete-trial concurrent VI 20-s (red key) VI 60-s (green key) schedule. Group 2 (Subjects 1351, 1374, and 6900) experienced a discrete-trial concurrent VI 180-s (red key) VI 60-s (green key) schedule. Each group experienced these contingencies for at least 50 sessions followed by 50 more sessions of a concurrent VI 60-s (blue) VI 60-s (yellow) schedule.

Two separate clocks controlled the schedules, and the IRTs were given by an approximation of an exponential distribution (Fleshler & Hoffman, 1962). These clocks always ran until a reinforcer was assigned at one of the schedule keys, whereupon only the clock controlling that schedule stopped. As in training, a trial lasted until the sooner of 2 s or a peck. The ITI durations were given by a Gaussian distribution with a mean of 6 s and a variance of 1. Pecking during an ITI reset the ITI timer. A reinforcer consisted of 2-s access to mixed grain at the food hopper, and each session lasted for 80 reinforcers.

RESULTS

Figure 2 provides the conditional probabilities of a switch, or changeover, given a background IRT of increasing values. Stated loosely, Figure 2 asks the question, "What is the probability of returning to schedule x after remaining on schedule y for t seconds after a switch?" To be more specific, note that the background IRT is given by the summed ITIs separating active-schedule responding plus the ITIs preceding and ending the run. For example, let L denote a response to the VI 20-s schedule and R denote a response to the VI 60-s schedule. Given the sequence L, ITI₁, R, ITI₂, R, ITI₃, R, ITI₄, L, we would say that the VI 20-s (L) schedule has a background IRT equal to ITI₁ + ITI₂ + ITI₃ + ITI₄.

Figure 2 shows that the probability of a switch remained relatively constant as the IRT at the background schedule increased. This

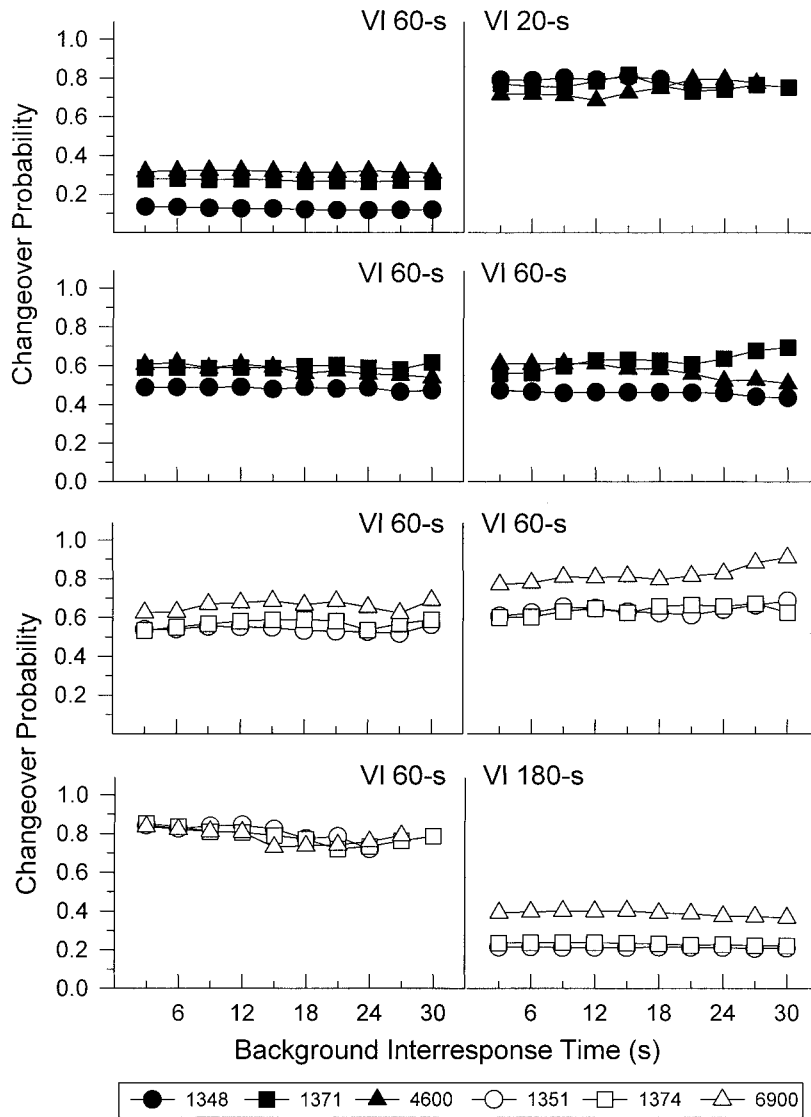


Fig. 2. The probability of a changeover to the background schedule given a background-schedule IRT. The background schedule is given above each graph. For example, as time away from the VI 20-s schedule accumulated, the probability that Subject 1348 switched to that schedule was high and flat. IRTs were divided into 3-s bins from 3 to 30 s, and relevant responses were collected over the last 20 sessions. Points were discarded for sample values of $n < 20$ (e.g., Subject 6900 at 30 s). Closed symbols refer to subjects 1348, 1371, and 4600. Open symbols refer to Subjects 1351, 1374, and 6900.

was true on average for all birds and all schedule values. Subjects 6900 and 1371, for example, both showed a slight increase in changeover probabilities given long background IRTs to one of the VI 60-s VI 60-s alternatives. However, Subject 4600 showed a slight decrease in changeover probabilities given the same schedule values.

Across schedule values the range of changeover probabilities varied. Given a 3:1 or 1:3 schedule ratio, and when the rich schedule was the background schedule, switch probabilities ranged from .69 to .88. Conversely, when the lean schedule was the background schedule, switch probabilities ranged from .11 to .4. Given a schedule ratio

of 1:1, switch probabilities ranged from .48 to .7, Subject 6900 being the sole exception to this latter finding.

Figure 3 plots the conditional probability of a switch to the background schedule given an IRT at the active schedule of t seconds, where t ranged from 0.0 to greater than 8.0 s. Recall that an active IRT is simply defined as the ITI immediately following a single choice. Given that subjects will respond substantially more to a richer alternative and given the Gaussian distribution that defined the ITI values, one might expect difficulty in collecting an adequate sample size for cases in which the lean schedule served as the active schedule. Therefore, the conditional probabilities shown were calculated over the last 20 sessions of each condition. Sample sizes ranged from $n = 50$ for the shortest and longest IRT values to $n = 2,414$ for the median IRT values.

Figure 3 shows that switch probabilities did vary as a function of the active-schedule IRT. All birds were more likely to switch from the richer of two VI schedules after a short IRT than after a long one. The probability of a changeover then decreased as the IRT increased. Conversely, from the lean schedule the probability of a switch increased as a function of IRT duration. The changeover function is essentially flat for all birds in the concurrent VI 60-s VI 60-s condition. Finally, although not explicitly shown here, the functions in Figure 3 were found to be approximately the same regardless of the sequential position of the peck that initiated the IRT. That is, changeover functions derived after a single peck to the VI 20-s schedule or after five pecks in sequence to the VI 20-s schedule approximated one another.

Figure 4 provides the relative IRT distributions generated by each subject on the last session of the concurrent VI 20-s VI 60-s and VI 180-s VI 60-s schedules. The IRT distributions have been separated according to the schedule choice (i.e., a left or right key peck). The resultant clock spaces, which also contain an indifference line, should be read as follows. First, in a discrete-trial procedure, as pecks accrue to, say, Schedule 1, the IRT associated with that schedule remains relatively fixed, whereas the IRT associated with Schedule 2 will increment. Such behavior will produce a line of points extending along the

Schedule 2 axis. Second, recall that the indifference line provides information as to the relative reinforcement probability at each schedule. That is, the y -axis schedule (here, the schedule assigned to the left key) has the higher probability of reinforcement at all points to the left of the indifference line. Conversely, the x -axis schedule has the higher probability of reinforcement at all points to the right of the indifference line. If a bird were following a momentary maximizing strategy, then whenever its relative IRTs placed the response to the left of the indifference line it should choose the left schedule key (i.e., the VI 20-s and VI 180-s schedules). Similarly, whenever a subject's relative IRTs placed it to the right of the indifference line it should choose the right schedule key (i.e., the VI 60-s schedule).

Figure 4 shows that pecks to the rich schedule key do not appear to conform to a momentary maximizing strategy. Rather, the subjects are seen to persevere on the rich schedule, producing points (i.e., pecks) with extraordinarily large lean-schedule IRT values. In terms of reinforcement probabilities, a peck to the VI 60-s schedule defined by $IRT_{VI\ 180} = 200$ s and $IRT_{VI\ 60} = 6$ s corresponds to $p(\text{rnf}_{VI\ 180}) = .67$ and $p(\text{rnf}_{VI\ 60}) = .26$, where rnf refers to reinforcement. Such a response is decidedly not optimal. In contrast, pecks to the lean schedule do support a momentary maximizing hypothesis, remaining on the appropriate side of the indifference line. Also, the figure shows that all of the subjects typically made only a single peck to the lean schedule key.

Finally, Table 1 provides the momentary maximizing statistic, M , for the first three and last three sessions of each condition for each bird. With the exception of Subject 6900, no consistent increase in M values were observed across these sessions. Whereas all other subjects typically produced M values in the range of .4 to .6, indicating an indifference to momentary reinforcement probabilities, the performance of Subject 6900 consistently produced M values greater than .7 by the end of training.

DISCUSSION

The birds used in Experiment 1 cannot be said to have followed a momentary maximizing strategy. First, an index of momentary

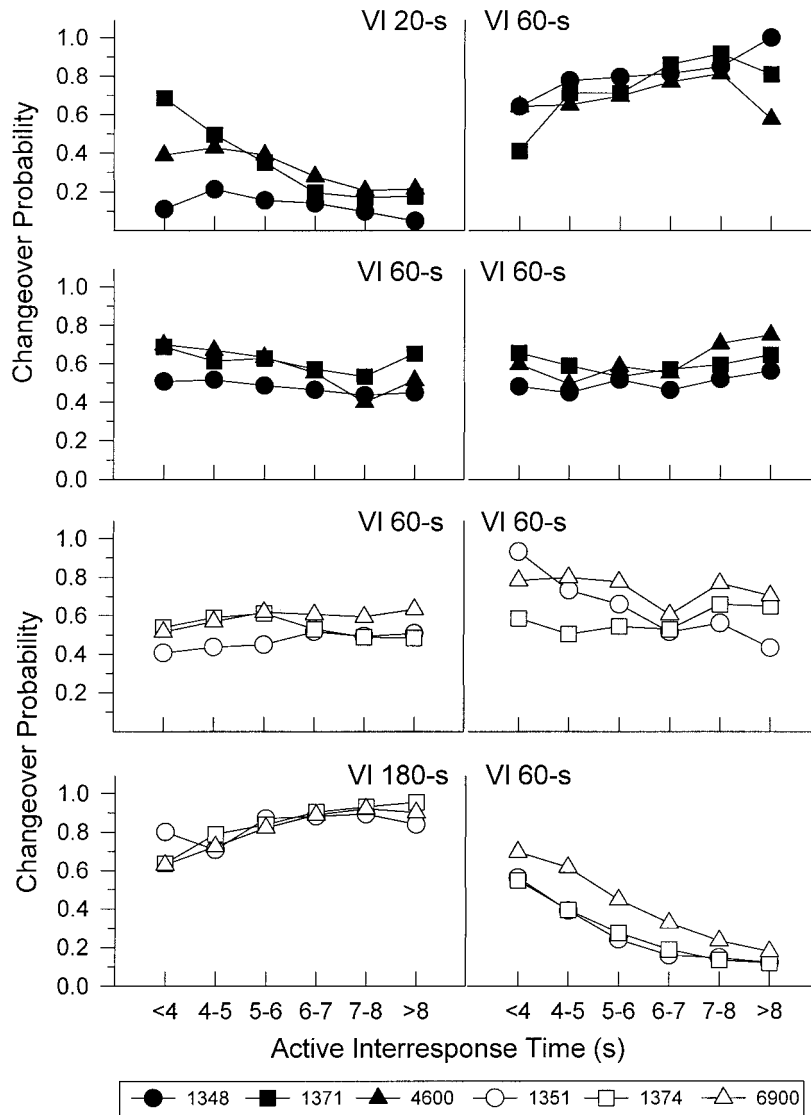


Fig. 3. Changeover probabilities to the background schedule given active-schedule IRT. Note that this IRT is defined as the programmed ITI occurring after the most recent schedule choice. IRTs were divided into two bins of less than 4 s and greater than 8 s. The remaining IRTs were placed into 1-s bins from 4 to 8 s. Relevant responses were collected over the last 20 sessions. Closed symbols represent Subjects 1348, 1371, and 4600. Open symbols represent Subjects 1351, 1374, and 6900.

maximizing, M , yielded only intermediate values that remained roughly constant over the course of the experiment (Table 1). Second, choice was found to be insensitive to IRTs accumulating to the background schedule (Figures 2 and 4). However, in contrast with these two results, Experiment 1 showed that changeover probabilities did vary as a function of the IRT that followed the most recent peck (Figure 3). Experiment 1, then, both

contradicts and supports momentary maximizing, and it simultaneously confirms and extends the results reported by Nevin (1969, 1979) and Silberberg et al. (1978, Experiment 1).

In terms of momentary reinforcement probabilities, an animal that experiences concurrent VI schedules of reinforcement should be more likely to switch given increasing IRTs at the background schedule. Figure 2 shows

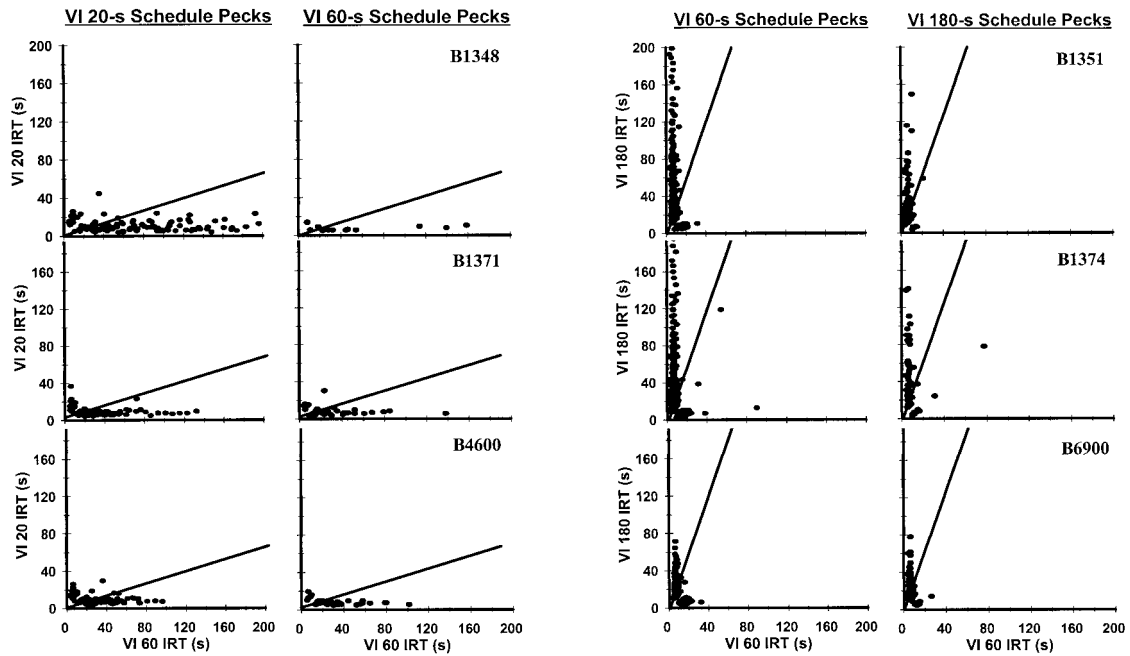


Fig. 4. The clock space distributions generated by each subject on the last session of the concurrent VI 20-s VI 60-s and concurrent VI 180-s VI 60-s schedules. The indifference line indicates the points at which the probability of reinforcement is equal for both alternatives. Pecks have been divided by the choice schedule to which they occurred, and every point defines a peck in terms of the IRT values. A point with axis coordinates of (120, 6) would indicate a peck in which the last VI 20 peck was 6 s ago, and the last VI 60 peck was 120 s ago.

that this was not the case. The birds were more likely to switch out of a lean versus a rich schedule, but the probability of a switch was roughly constant, despite the ever-increasing reinforcement probability at the al-

ternative. This result confirms the findings of Nevin (1969, 1979) and Silberberg et al. (1978). These two studies also employed discrete-trial procedures and found that the probability of a changeover remained constant as the number of pecks made on a single schedule increased.

Figure 3, however, shows that switch probabilities following IRTs to the active schedule roughly conformed to a momentary maximizing strategy. During a concurrent VI VI experiment, the alternate schedule always has the higher probability of reinforcement immediately following a choice (see Figure 1). This is true regardless of whether the just-emitted choice was to the rich or lean schedule because responding resets the IRT of the selected choice, thus temporarily setting the reinforcement probability of that alternative to zero. As the active-schedule IRT increases, momentary maximizing predicts two results. First, if the active schedule is the richer of the two schedules, with longer subsequent IRTs the higher rate of gain associated with its reinforcement probability will eventually make

Table 1

The momentary maximizing statistic, *M*, calculated for Experiment 1.

Session	Subject					
	1351	1374	6900	1348	1371	4600
	VI 180 s VI 60 s			VI 20 s VI 60 s		
1	0.48	0.52	0.52	0.63	0.47	0.58
2	0.49	0.49	0.61	0.37	0.38	0.49
3	0.47	0.50	0.55	0.51	0.41	0.50
Last 2	0.40	0.29	0.71	0.39	0.54	0.61
Last 1	0.44	0.50	0.74	0.27	0.58	0.50
Last	0.59	0.41	0.78	0.18	0.62	0.61
	VI 60 s VI 60 s					
1	0.33	0.54	0.62	0.25	0.44	0.38
2	0.46	0.5	0.64	0.46	0.53	0.43
3	0.36	0.61	0.51	0.36	0.47	0.4
Last 2	0.47	0.57	0.61	0.33	0.57	0.53
Last 1	0.52	0.57	0.63	0.27	0.54	0.59
Last	0.51	0.52	0.59	0.40	0.60	0.57

it, once again, the more likely to be reinforced. That is, following a rich-schedule choice, the longer the IRT, the more likely it is that the subject should repeat its choice instead of switching. However, if the schedule values are the same or if the just-emitted response was to the leaner of two schedules, then a switch will always produce a higher probability of reinforcement, and this probability will only increase with time. In such situations momentary maximizing predicts flat or increasing changeover probabilities with longer IRTs. Figure 3 shows each of these predicted results.

Finally, the results reported here illustrate the importance of the IRT distribution that the experimenter allows the organism. For example, let us for the moment assume the generality of the results found in Experiment 1. Using a discrete-trial concurrent VI VI choice procedure, if one makes the ITI progressively longer, then one would expect more and more perseveration on the rich schedule and, hence, overmatching. Conversely, birds might be made to approach perfect alternation, or undermatching, given an IRT/ITI distribution that yields long IRTs after a lean-schedule choice and short IRTs after a rich-schedule choice. To my knowledge, these experiments have not been performed.

EXPERIMENT 2

Experiment 1 suggested that pigeons on a concurrent VI VI schedule base their choice on the time since the most recent peck (i.e., the active-schedule IRT). However, there is reason to believe that a free-operant procedure, with its relatively high response rates, might produce different results, namely a sensitivity to both active and background IRTs. To understand this prediction, let us consider the response sequence L, IRT₁, R, IRT₂, R, IRT₃, ??, where ?? is an unknown choice. Here the active-schedule IRT = IRT₃, and the background-schedule IRT = IRT₁ + IRT₂ + IRT₃. That is, background and active IRTs are defined in reference to a specific choice, L versus R. Now, if we assume that the memory for each of these choices is a decreasing function of time, then the longer IRT₁, IRT₂, and IRT₃, the weaker will be the memory for L, and hence, the weaker will be any sensitivity to the background-schedule

IRT. In contrast, if IRT₁, IRT₂, and IRT₃ are very short, then one would expect a stronger memory for (i.e., control by) L, and hence a greater sensitivity to the background-schedule IRT.

Indeed, perhaps the strongest evidence in favor of momentary maximizing, and hence a sensitivity to both active and background IRTs, does come from free-operant procedures. Not only do the response sequences emitted by pigeons under such conditions approximate the ratio of the schedule values (i.e., the maximizing sequence; Shimp, 1969; Silberberg et al., 1978) but, more important, given variable response rates, Hinson and Staddon (1983a, 1983b) showed that pigeons in concurrent VI VI schedules allocate their relative IRTs in a manner that correlates with reinforcement probabilities. Using clock space representations (see Figure 1), they found that when a subject's relative IRTs placed it on one side of the indifference line, the subject tended to respond to the schedule with the higher probability of reinforcement.

However, Heyman (1979) showed that in terms of run length, the behavior of pigeons in a free-operant concurrent VI VI procedure did not appear to be sensitive to the time spent away from the background schedule. His birds, like Nevin's under a discrete-trial procedure (Nevin, 1969), were no more likely to switch schedules after one peck than after 10 pecks to a single alternative. That is, the probability of a switch did not increase as a function of run length (however, see Todorov, Souza, & Bori, 1993). Of course, as repeatedly stressed in this paper, run length is not an accurate measure of reinforcement probability, especially in a free-operant procedure that permits variable rates of responding.

Experiment 2, then, explicitly took IRTs as its independent variable and asked if, during a free-operant concurrent VI VI procedure, pigeons' responding would show a sensitivity to both active and background IRTs. The experiment used a procedure similar to that of Findley (1958), in which changeover keys determined which of two schedules was available at a central key. This procedure served to clarify the distinction between active-schedule and background-schedule IRTs.

METHOD

Subjects

Five White Carneau pigeons were used (Subjects 2087, 2195, 7323, 7351, and 7380). None of these birds had previously experienced VI schedules, although all had recently been used in a relative numerosity pilot study. All subjects were maintained at 85% of their free-feeding weights during the experiments. The pigeons were fed a mixture of mixed grain and pellets (Purina ProGrain® and Purina Nutriblend® Green and Gold) when not run or to maintain target weight.

Apparatus

All daily sessions were conducted in the apparatus described for Experiment 1 with the following changes. A third key (termed here the center key) was added to the front panel above the two other keys (termed here side keys) such that the three keys formed an equilateral triangle with sides equal to 3 cm. A 6.3-V Grason-Stadler stimulus projector was mounted behind each of the keys.

Procedure

Training. All subjects were given 1 day of Phase 1 training and 2 days of Phase 2 training at the start of each concurrent VI VI condition. Phase 1 trained forced alternation. A side key was lit, say green. A peck to this key illuminated the center key with the same color, green, and turned off the side key. Ten pecks to the center key then operated the food hopper (a fixed-ratio [FR] 10 schedule). Upon reinforcement the opposite side key became lit, say red, and the center key was turned off. A single peck to this key illuminated the center key, red, and turned off the side key. Again, 10 pecks to the center key produced access to grain. Such alternations occurred until 40 reinforcers had been collected.

During Phase 2 training, the side keys always remained lit, say red and green. A peck to either side key determined the color of the center key (e.g., a peck to the red side key lit the center key with a red hue). Reinforcements were alternately assigned to the center key colors using an FR 10 schedule. All birds quickly learned to use the side keys to control the center key color and to alternate peck runs to these colors.

Test. After training, the pigeons experienced concurrent schedules in the three-key operant chamber. Respective key colors were green and red for Condition 1 (concurrent VI 180-s VI 60-s schedule) and yellow and blue for Condition 2 (concurrent VI 20-s VI 60-s schedule). Each condition lasted for 25 sessions of 80 reinforcers each.

Side keys acted as changeover keys, and the center key served as the schedule key. For example, in Condition 1 the left side key (red), if pecked, changed the center key from green to red. The right side key (green), if pecked, changed the center key to green. The pigeon received mixed grain after pecks to the center key as determined by the active reinforcement schedule. Two separate clocks controlled the reinforcement schedules, which were given by an approximation of the exponential distribution (Fleshler & Hoffman, 1962).

RESULTS

Figure 5 plots the conditional probability of a switch given an IRT at the background schedule of 0.0 s to 6.0 s over the last three sessions of Condition 1 (concurrent VI 180 s VI 60 s) and Condition 2 (concurrent VI 20 s VI 60 s). In Condition 1, only Subjects 7351 and 2087 displayed an increase in changeover probabilities with time away from the leaner schedule (VI 180 s), but these increases were slight. Most often, the probability of a switch, given an increasing background-schedule IRT for either schedule, remained flat for all birds. Changeover probabilities ranged from .61 to 1.0 when the background schedule was the VI 60-s schedule. When the background schedule was the VI 180-s schedule, changeover probabilities ranged from .13 to .53.

In Condition 2, Subjects 7380, 7351, and 2087 produced results similar to those in Condition 1. That is, regardless of the background schedule, the probability of a switch remained essentially flat with increasing IRT and was either high or low depending on the background schedule value. Subjects 7323 and 2195, however, showed a large transient decrease in changeover probabilities when the richer schedule (VI 20 s) served as the background schedule. For these 2 birds, the probability of switching to the richer schedule decreased quite substantially with background IRTs of 0.5 s to 3.0 s. The probability

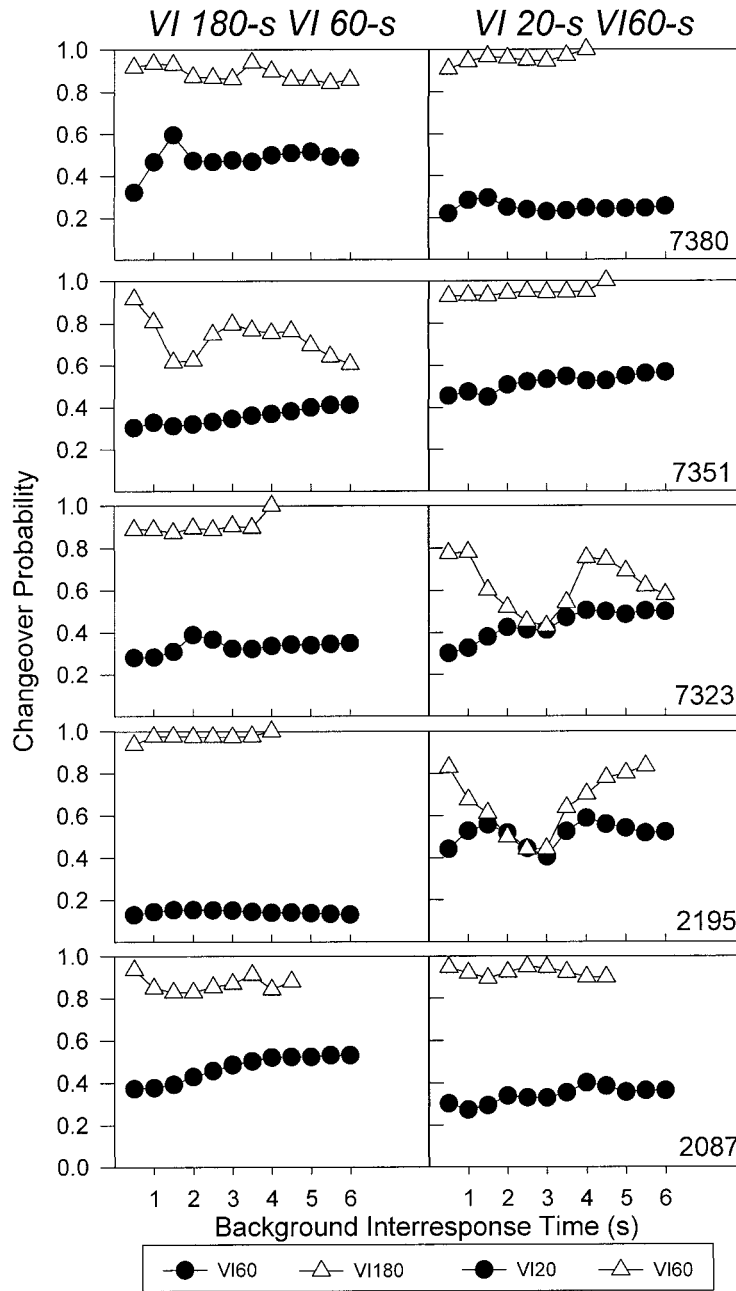


Fig. 5. The probability of a changeover given increasing background IRT for the rich and lean schedules during free-operant choice. IRTs were divided into 0.5-s bins and were taken from the last three sessions of each condition. Bins with fewer than 20 accumulated IRTs were not considered. The background schedule is given by the symbol type. For example, for Subject 7380 during the concurrent VI 180-s VI 60-s condition, as time away from the VI 180-s schedule (solid circle) increased, the probability of a switch to that schedule remained flat at approximately .5.

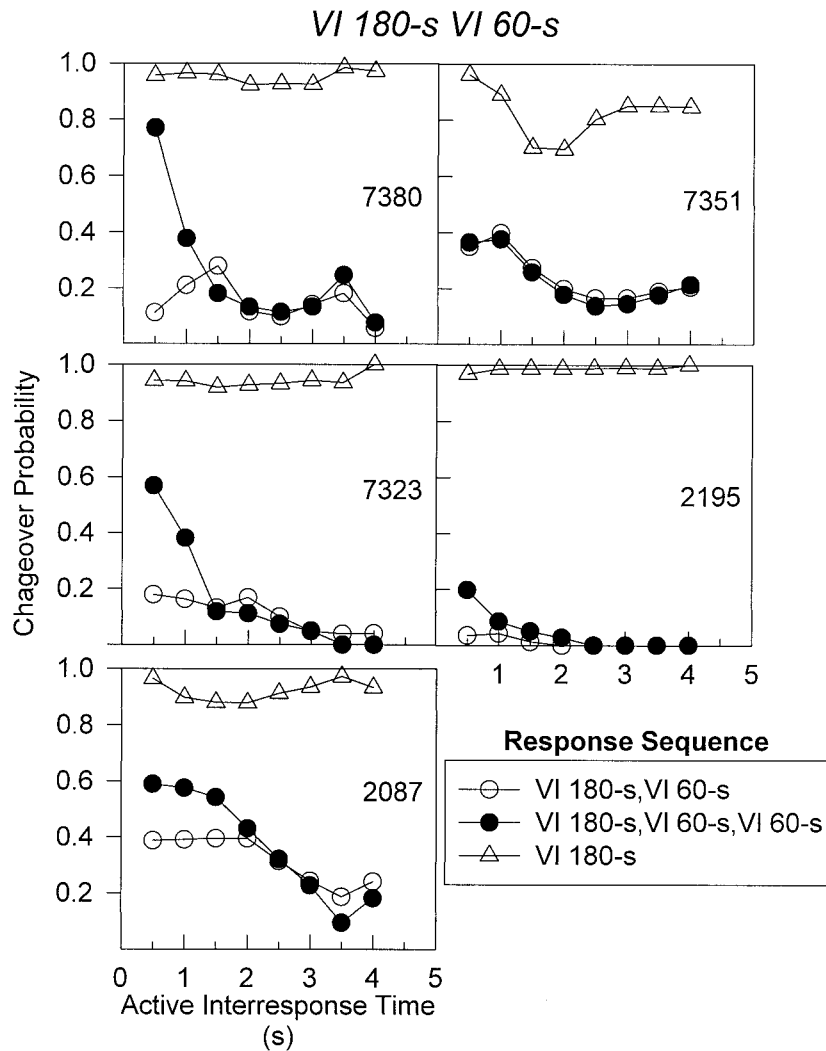


Fig. 6. The changeover probabilities given active-schedule IRT during the concurrent VI 180-s VI 60-s condition. IRTs were divided into 0.5-s bins during the last three sessions of the condition. Probabilities with a sample of $n < 20$ were not considered. The sequential position of the active-schedule response was found to influence the probability of a switch, but only when this response was to the richer schedule. Hence, plots are shown for two types of rich-schedule response: those that were preceded by a lean-schedule response (open circles) and those that were preceded by a rich-schedule response (closed circles).

of a switch to the rich schedule then tended to increase with longer background IRTs.

Figures 6 and 7 plot the conditional probability of a switch given active-schedule IRTs of 0.0 s to 3.5 s over the last three sessions of Condition 1 and Condition 2, respectively. For Condition 1 when the most recent choice consisted of the VI 180-s schedule, all birds showed a high, flat changeover probability, irrespective of the time since this choice (i.e.,

the active-schedule IRT). Further, the sequential context of the VI 180-s choice did not affect the subsequent changeover probabilities. Conversely, when the active schedule was the rich VI 60-s schedule, both IRT and sequential context affected the subsequent changeover probability. First, Figure 6 shows that changeover probabilities were an inverse function of active-schedule IRT when the rich schedule served as the active schedule. Sec-

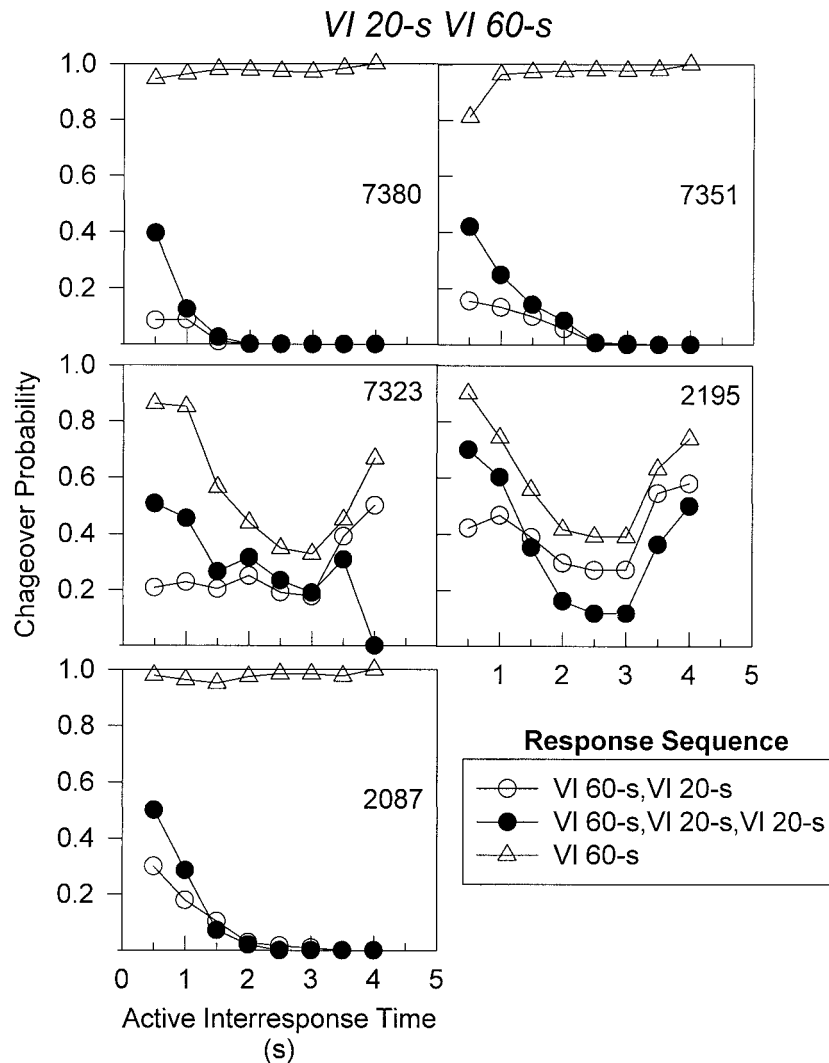


Fig. 7. The changeover probabilities given active-schedule IRT during the concurrent VI 20-s VI 60-s condition. See Figure 6 for explanation.

ond, Figure 6 also shows that the range of changeover probability decrement depended upon the sequential position of the prior rich-schedule choice. For example, given an IRT of 0.5 s, all subjects in Condition 1 were more likely to switch after the choice sequence VI 60 s VI 60 s than after VI 180 s VI 60 s.

In Condition 2, Subjects 7380, 7351, and 2087 produced results similar to those in Figure 6. The probability of a switch after a choice to the lean schedule (VI 60 s) was high and flat with increasing IRT. The probability of a switch after a choice to the rich schedule

depended upon the sequential position of the peck and was an inverse function of IRT. However, Subjects 7323 and 2195 showed anomalous results. Although the birds did show an effect of the sequential position of the rich-schedule choice, these 2 subjects displayed a large transient decrease in changeover probabilities as a function of active-schedule IRTs. This result was evident regardless of the schedule value.

Table 2 provides *M* values for the first three and last three sessions of Condition 1 (concurrent VI 180 s VI 60 s) and Condition 2 (concurrent VI 20 s VI 60 s). Recall that *M* is

Table 2

The momentary maximizing statistic, M , calculated for Experiment 2.

Session	Subject				
	2087	2195	7323	7351	7380
VI 180 s VI 60 s					
1	0.32	0.20	0.36	0.30	0.29
2	0.25	0.14	0.40	0.49	0.41
3	0.41	0.28	0.44	0.40	0.14
Last 2	0.81	0.19	0.75	0.71	0.89
Last 1	0.80	0.29	0.72	0.79	0.90
Last	0.80	0.31	0.74	0.73	0.84
VI 20 s VI 60 s					
1	0.12	0.15	0.17	0.24	0.13
2	0.56	0.38	0.40	0.56	0.48
3	0.65	0.41	0.55	0.62	0.58
Last 2	0.77	0.74	0.75	0.66	0.51
Last 1	0.75	0.74	0.78	0.59	0.59
Last	0.82	0.75	0.75	0.57	0.60

a ratio of reinforcement probability differences calculated across all pecks emitted during a session. In Condition 1, with the exception of Subject 2195, all subjects showed an increase in M values as the experiment progressed. That is, the birds came to allocate their pecks in such a manner that their behavior was positively correlated with momentary reinforcement probabilities. In contrast, Subject 2195's behavior was consistently sub-optimal, and showed no improvement over the course of the condition.

In Condition 2, only 3 of 5 subjects produced M values at the end of training that were consistently higher than those in the first three sessions. Interestingly, one of these, Subject 2195, was the only subject to show no such improvement in Condition 1. However, it must be pointed out that only 25 sessions were conducted in each condition, and Condition 2 directly followed Condition 1, using the same experimental apparatus.

Finally, Figures 8 and 9 plot the IRT distributions from the last session of Conditions 1 and 2 for each subject. The resultant clock spaces are similar to those reported by Hinson and Staddon (1983a). For instance, all of the clock spaces show a band devoid of points along the diagonal. This space gives the minimum changeover time imposed by the apparatus. That is, the diagonal of a clock space represents the points at which $IRT_{\text{Choice 1}} = IRT_{\text{Choice 2}}$, and as the changeover time gets

larger the points of a clock space will necessarily diverge from the diagonal by a greater and greater extent. Likewise, consecutive pecks to a single schedule also require a minimum amount of time, and the empty bands along the choice axes indicate this time.

In terms of momentary maximizing, Figures 8 and 9 make two points. First, responses to the lean schedule tended to occur in a single band along the lean-schedule axis. This observation is seen for Subjects 2195, 7323, and 7380 in Figure 8 (i.e., Condition 1) and for Subjects 2087, 2195, 7323, and 7380 in Figure 9 (i.e., Condition 2). As was noted in Experiment 1, such response structure indicates that the subjects tended to make only a single peck to the lean schedule before switching. Second, responding to the rich schedule, in terms of relative IRTs, tended to be quite variable. However, responding associated with high M values did tend to fall on the side of the indifference line closest to the rich-schedule axis.

DISCUSSION

The findings of Experiment 2, using a free-operand procedure, were similar to those of Experiment 1, which used a discrete-trial procedure. That is, the behavior of all of the subjects in Experiment 2 showed some sensitivity to IRT at the active schedule. If the most recent peck was to the rich schedule, then the probability of a switch on the subsequent choice decreased with increasing IRT. Conversely, if the most recent choice was to the lean schedule, then the probability of a switch remained high and flat with increasing IRT. Such a result, a sensitivity to active-schedule IRT, is predicted by a momentary maximizing strategy. A momentary maximizing hypothesis is further corroborated by the structure of responding shown in Figures 8 and 9. Momentary maximizing specifies that an animal should make only a single response on the lean schedule before switching from that schedule (Houston & McNamara, 1981). Such a result was found. Finally, the calculated M values in Table 2 showed an increase over the course of both conditions for the majority of the subjects.

However, although the data in Experiment 2 showed a sensitivity to active-schedule IRT, they showed no such sensitivity to background-schedule IRT. This latter finding is

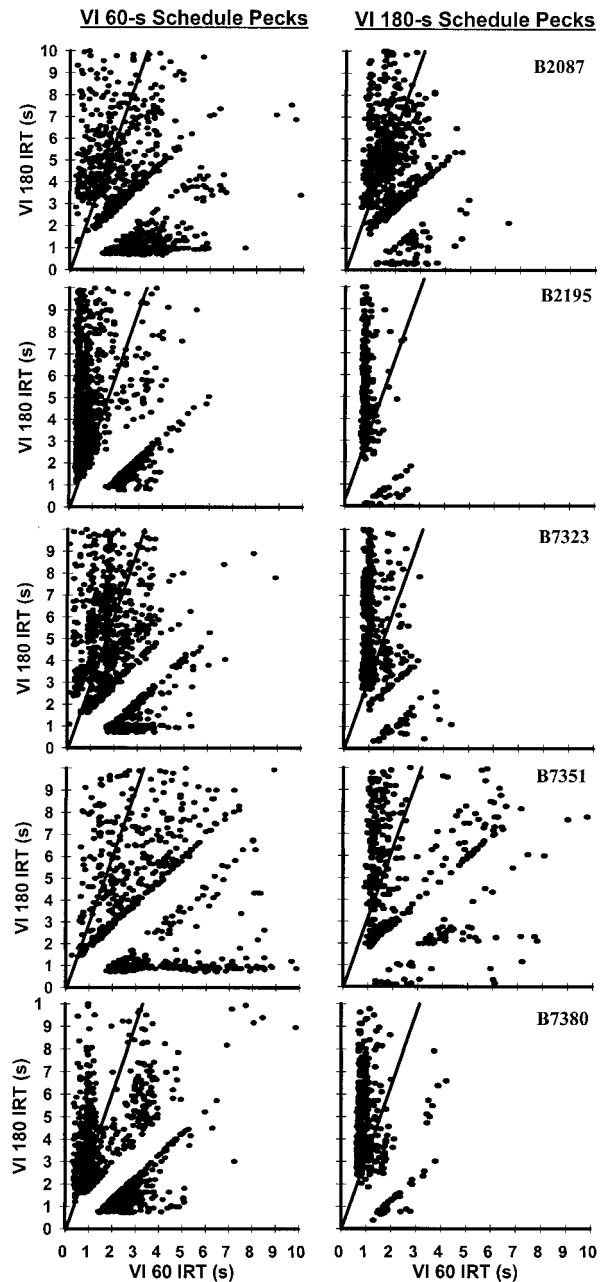


Fig. 8. Clock space plots for each bird during the last session of the concurrent VI 60-s VI 180-s condition. For simplicity pecks were divided into those made to the rich and lean schedules. Each point represents a peck defined in terms of the time since the last peck to both schedules so that a point with axis coordinates of (1, 5) represents a peck that occurred 1 s after the last VI 60-s choice and 5 s after the most recent VI 180-s choice. The region to the left of the indifference line produces higher reinforcement probabilities for the VI 180-s schedule. The region to the right of the indifference line produces higher reinforcement probabilities for the VI 60-s schedule.

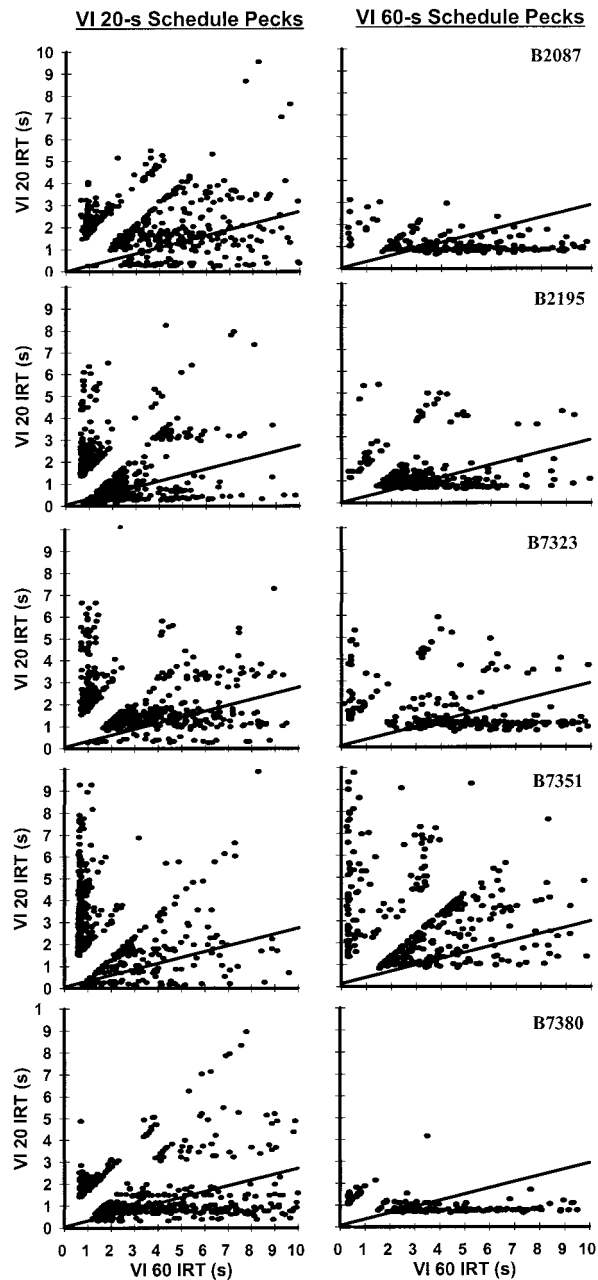


Fig. 9. Clock space plots for each bird during the last session of the concurrent VI 60-s VI 20-s condition. By convention, the *x* axis represents the schedule value associated with the left key of the apparatus; hence, the different slope of the indifference line when compared to Figure 8.

contrary to a momentary maximizing strategy. In Experiment 2, responses on the two side keys determined the reinforcement schedule at the center key. With this design the background schedule was synonymous with whichever schedule was not currently

present at the center key. As the background-schedule IRT increased, the probability of a changeover, or switch, remained flat, regardless of the relative schedule value. This result extends the findings of Heyman (1979), who also used a free-operant concurrent VI VI

procedure based on a Findley (1958) design. Heyman found that changeover probabilities remained flat as run length at a choice increased. However, as described earlier only an insensitivity to run time, not run length, contradicts momentary maximizing. The current study did indeed find an insensitivity to run time.

Unlike Experiment 1, however, Experiment 2 produced sequential dependencies in the sensitivity to IRT at the active schedule. For example, if R was the rich schedule, then the IRT after the response pattern LR was treated differently than after the response pattern LRR. This supports the notion that "switch" and "stay" are separate response types, or states, for pigeons under concurrent VI VI schedules (e.g., Heyman, 1979; Silberberg et al., 1978). Further, an important difference between discrete-trial and free-operant procedures might be the salience of these two response types, because Experiment 1 found no such sequential dependencies.

Another difference between Experiment 1 and Experiment 2 concerns the concurrent VI 20-s VI 60-s choice schedules (e.g., Figures 6 and 7). In Experiment 2, for 2 of 5 birds the probability of a switch decreased and then increased as a function of background IRT duration, but only when the VI 20 s served as the background schedule. Such curves were not seen in Experiment 1. A possible explanation is that this result is a transitional finding. After all, the lean-schedule key of Condition 2 served as the rich-schedule key for Condition 1. Although the color of the key was changed between conditions, this may be a minor change when compared to the experimental chamber as a whole. Also, instead of a performance criterion for stability, conditions were studied for 25 sessions. It is possible that the switch function would have flattened out with additional training.

GENERAL DISCUSSION

The results described in this paper may be summarized quite simply. First, both Experiments 1 and 2 suggested that, of the two interresponse times operating in a two-choice concurrent VI VI procedure, only the active-schedule IRT influenced the subsequent choice. Second, Experiment 2 showed that in

a free-operant procedure, the sequential position of the rich-schedule choice also influenced the subsequent schedule choice. The remainder of the paper will show that these two findings permit the reconciliation of a wide range of seemingly disparate choice data. That is, I will show that data in support of momentary maximizing (e.g., Hinson & Staddon, 1983a, 1983b; Shimp, 1966, 1969; Silberberg et al., 1978) and evidence contrary to such a strategy (e.g., Heyman, 1979; Nevin, 1969, 1979) may both be derived from the results described in this paper. I will do so by formalizing the results of this paper as a continuous-time Markov chain.

A Continuous-Time Markov Description

A Markov chain describes a finite set of discrete states in which the probability of a switch from one state to any other is given solely by the resident state (see Metz, 1974, for analyses and ethological applications). State refers here to nothing more than a stimulus, or stimulus set, that controls the current behavior of an organism. For example, let us consider a two-choice procedure, and for the moment, discrete transition probabilities as shown in Figure 10. If a subject's immediate response history has no bearing on its current choice (a zero-order Markov chain), then we might say that the subject possesses a single weighted coin that it flips prior to making a choice. In this instance, molecular choice probabilities equal the subject's molar response probabilities. If the subject can remember (i.e., is controlled by) only its most immediate choice (a first-order Markov chain), then we might say that the animal possesses two coins and, hence, two states: (a) When the last peck was to Schedule 1, flip Coin 1, and (b) when the last peck was to Schedule 2, flip Coin 2. As a final example, if we believe that a subject's behavior is controlled by not only its last response, but also by whether that response was a switch or a stay (a second-order Markov chain), then we must use four states to describe the animal's behavior.

A Markov chain description of the discrete-trial and free-operant data reported in Experiments 1 and 2 requires chains of first and second order (see Figure 10). This finding corroborates Heyman (1979). Heyman showed that a Markov chain of first or second

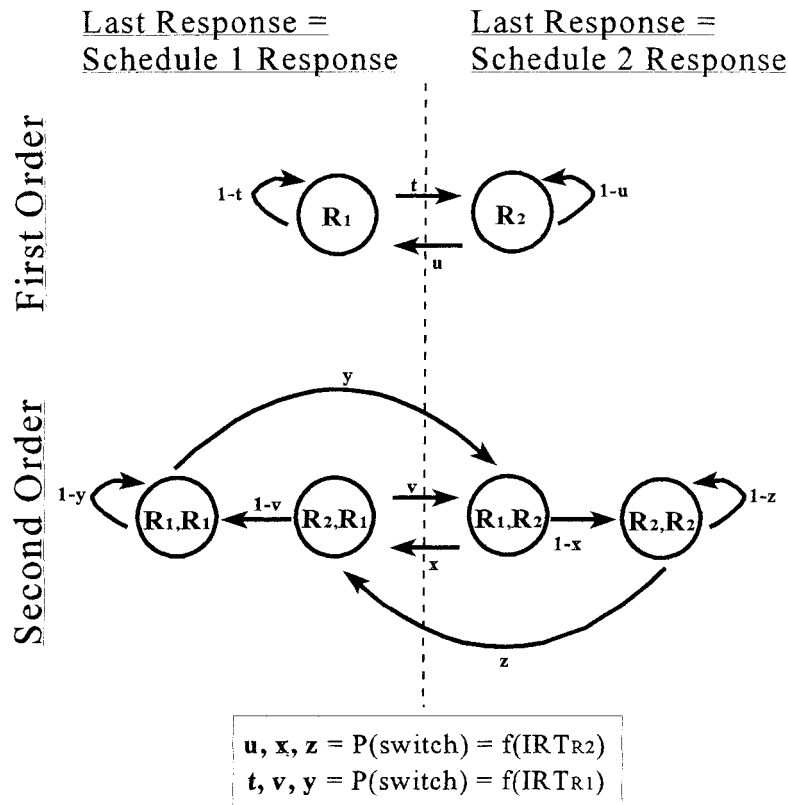


Fig. 10. Schematics of the Markov chains used to simulate Experiments 1 and 2. For the discrete-trial simulations, a first-order chain was used. Response states were defined by the schedule to which the most recent response occurred. To simulate the free-operant results, second-order Markov chains were used. Response states were defined by the two most recent schedule choices. The variables t through z represent the probability of exiting a given response state. These probabilities were not static, but rather were a function of the active-schedule IRT.

order could fit the molar response structure generated in free-operant variable-interval schedules. Such a conceptualization, independent of time, also supported the flat changeover probabilities associated with increases in run length that were observed by Nevin (1969, 1979) and Heyman (1979). However, it is unclear whether a stochastic process that is independent of time can generate the clock spaces observed by Hinson and Staddon (1983a, 1983b). Also, a purely peck-based Markov chain is contrary to the findings reported in this paper. The data reported in this paper, however, may be precisely formalized via a continuous-time Markov chain.

With a continuous-time Markov chain, the transition probabilities between states are not discrete. Rather, the probability of a

switch from a given state depends on the resident time in that state. Such a conceptualization allows for the findings, reported in this paper, that switch probabilities change as a function of time. A more precise characterization of this finding, however, must further specify that the temporal dependencies within a state reset after each event (i.e., pecks) that occur in that state (recall that the active-schedule IRT, by definition, is reset to zero with every response). An event that resets a state's clock in such a manner is termed a *renewal process*.

To summarize, then, the results reported in this paper may be formally characterized by continuous-time Markov chains. For the discrete-trial case (Experiment 1), a first-order Markov chain is sufficient to describe the data. Conversely, for the free-operant case

Table 3
The momentary maximizing statistic, M , calculated for simulated sessions.

Session	Experiment 1 Subject						Experiment 2 Subject					
	1351	1374	6900	1348	1371	4600	2087	2195	7323	7351	7380	
	VI 180 s VI 60 s			VI 20 s VI 60 s			VI 180 s VI 60 s					
Last 2	0.40	0.54	0.68	0.34	0.56	0.52	0.77	0.27	0.73	0.66	0.85	
Last 1	0.47	0.56	0.70	0.30	0.57	0.64	0.78	0.30	0.69	0.69	0.85	
Last	0.52	0.53	0.74	0.30	0.55	0.69	0.77	0.36	0.69	0.77	0.86	
							VI 20 s VI 60 s					
Last 2							0.74	0.73	0.71	0.65	0.52	
Last 1							0.71	0.76	0.69	0.53	0.65	
Last							0.71	0.78	0.69	0.57	0.59	

Note. The given M values were calculated from simulations of the last three sessions. All simulations used the active-schedule switch probabilities given in Figures 3 (discrete trial) and 6 and 7 (free operant) and randomized arrays of the actual IRT distributions recorded for each session.

(Experiment 2), a second-order Markov chain is required.

Simulations of Molecular Data

Having formalized the findings reported in Experiments 1 and 2, one derivative result becomes evident. By default, a continuous-time Markov chain of first or second order will show flat changeover probabilities with increasing run lengths (Heyman, 1979; Nevin, 1969, 1979). This result is expected because changeover probabilities will not be given by the run length itself, but by the last IRT of the run averaged over many, a great many, runs. For example, in a discrete-trial procedure with a 6-s ITI (e.g., Nevin, 1969), the last IRT at, say, the rich schedule will always approximate 6 s whether the prior run consisted of 3, 4, or 20 pecks. Therefore, the same response state would characterize each of these instances, and we would predict the same changeover probability for all run lengths. Similarly, in a free-operant procedure (e.g., Heyman, 1979), assuming that the same average rate of pecking applies to all run lengths, the average IRT after the last peck of any run would be roughly constant. Given that the bird would reside in the same state of the Markov chain, it would thus be obtaining the same transitional probability distribution after all run lengths.

First- and second-order continuous-time Markov chains will yield a data set held to be contrary to momentary maximizing. Will the same Markov chains account for data that seem to be indicative of a momentary maxi-

mizing process (e.g., Hinson & Staddon, 1983a; Shimp, 1966, 1969; Silberberg et al., 1978)? To answer this question, Monte Carlo simulations of the experiments reported in this paper were conducted run using the Markov chains depicted in Figure 10. First, the transition probabilities for each bird were taken from the observed active-schedule switch probabilities (i.e., Figures 3, 6, and 7). Second, an IRT distribution was created and sampled for each simulation. For discrete-trial simulations, the IRT distribution was given by the programmed Gaussian function with a mean of 6 s and variance of 1. Free-operant IRT distributions, on the other hand, were created by randomly drawing from the actual IRT distributions recorded for each subject in each session.

Table 3 provides M values calculated from simulations of the last three sessions of Experiment 1 (Condition 1) and Experiment 2 (Conditions 1 and 2). Given the crude level of the simulation, these values are remarkably similar to those given in Tables 1 and 2 (standard deviation of differences = 0.04). First, in both the real and simulated data, the free-operant procedure tended to produce higher M values than the discrete-trial procedure did. Second, if we were to rank subjects according to the M values shown in Tables 1 and 2, then Table 3 would produce a close match of this ranking. For instance, the highest M values recorded in Experiments 1 and 2 were produced by Subjects 6900 and 7380, respectively, under concurrent VI 180-s VI 60-s schedules. The simulated data also show

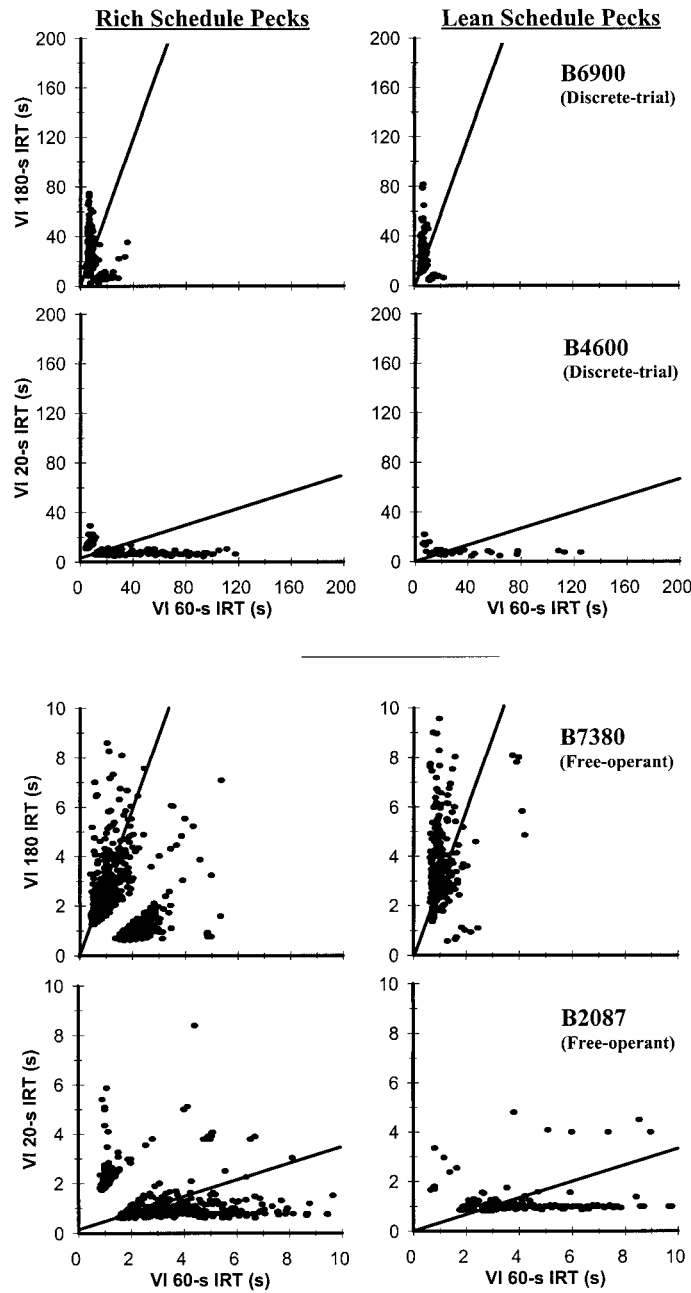


Fig. 11. Simulated clock spaces for Subjects 6900, 4600, 7380, and 2087. Simulations used the continuous-time Markov chains illustrated in Figure 10. IRT distributions for the discrete-trial procedure consisted of the Gaussian function used in Experiment 1. To simulate free-operant sessions, the sequence of real IRTs for each subject, independent of IRT type, was randomized and “fed” into the Markov chain.

these subjects as having the highest M values for their respective procedures. Similarly, the lowest M values recorded in Experiments 1 and 2 were produced by Subjects 1348 (concurrent VI 20 s VI 60 s) and 2195 (concurrent

VI 180 s VI 60 s), respectively, and again the simulated data correspond with these real results.

Finally, Figure 11 provides choice-by-choice relative IRT distributions (i.e., clock spaces)

generated from simulations of the last experimental session for Subjects 6900, 4600, 7380, and 2087. The performance of these subjects most closely approximated a momentary maximizing strategy for each of the four experimental conditions: concurrent VI 180 s VI 60 s (discrete trial), concurrent VI 20 s VI 60 s (discrete trial), concurrent VI 180 s VI 60 s (free operant), and concurrent VI 20 s VI 60 s (free operant). Given the stochastic nature of a Markov chain, we should not expect an exact match between simulated and actual data. After all, the probability that a thousand coin tosses (pecks) will produce the exact same sequence of heads and tails is infinitesimal. Nevertheless, in the case of discrete-trial simulations, the simulated clock spaces are nearly identical to the real data shown in Figure 4. Simulated free-operant clock spaces are also quite similar to the data shown in Figures 8 and 9.

Implications for Dynamic Models of Choice

The goal of a Markov chain analysis is to formally and precisely describe behavioral structure. To this end, the simulations presented in this paper suggest that choice during concurrent VI schedules is largely under the control of two molecular variables: (a) the *active interresponse time*, defined as the time since the most recent response, and (b) the current *response state*, defined by the most recent one or two choice allocations. Obviously, a dynamic mechanism must underlie each of these variables.

The response states illustrated in Figure 10 are defined by past choice allocations. Thus, an implicit assumption is that the sequential dependencies observed in Experiment 2 but not evident in Experiment 1 involve memory dynamics. Perhaps pigeons in a discrete-trial procedure can remember only the most recent response or IRT. Regarding memory for IRTs, Shimp (1981) reported evidence that pigeons can remember IRTs for several seconds after their emission. Regarding responses, Machado (1992, 1994) found that pigeons' choices under frequency-dependent schedules were controlled by the last one, two, or three responses. A similar result was reported using budgerigars and vocal responses (Manabe, Staddon, & Cleaveland, 1997). One might therefore expect that dis-

crete-trial concurrent VI VI procedures with a short ITI would show sequential dependencies similar to those of free-operant procedures. A dynamic process, however, will need to specify how the memory traces for individual IRTs and pecks decay or interfere with one another in real time.

Another issue that a dynamic model of choice must address is how IRTs become weighted to produce a switch response. That is, what sets the transition probabilities assigned to any given active IRT? Several studies have shown the importance of IRT reinforcement contingencies in determining response rate (Hawkes & Shimp, 1974; Platt, 1979; Shimp, 1973, 1974) during single VI schedules. Whether a similar analysis can be directed at concurrent VI VI schedules remains to be seen. Interestingly, the anomalous results of Subjects 7323 and 2195 from Experiment 2, Condition 2 (Figures 5 and 7), may be pertinent here. These birds showed a dip in the probability of a switch given intermediate lean-schedule active IRTs. These intermediate IRTs represent an area of weak selection pressure by differential reinforcement probabilities. That is, if birds are learning a relation between active-schedule IRT and the reinforcement probabilities at a set of alternatives, we would expect this learning to be most difficult when the differences between these probabilities are minimal (i.e., where the curves cross in Figure 1). Further experiments that examine the relation between IRT and switch probabilities during the transitions between a broad range of VI schedules should clarify this issue.

REFERENCES

- Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, 22, 231-242.
- Baum, W. M. (1979). Matching, undermatching and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, 32, 269-281.
- de Villiers, P. A. (1977). Choice in concurrent schedules and quantitative formulations of the law of effect. In W. K. Honig & J. Staddon (Eds.), *Operant behavior* (pp. 233-287). New York: Prentice Hall.
- Findley, J. D. (1958). Preference and switching under concurrent scheduling. *Journal of the Experimental Analysis of Behavior*, 1, 123-144.
- Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 5, 529-530.
- Hawkes, L., & Shimp, C. P. (1974). Choice between re-

- sponse rates. *Journal of the Experimental Analysis of Behavior*, 21, 109–115.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4, 267–272.
- Heyman, G. M. (1979). A Markov model description of changeover probabilities on concurrent variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 7, 133–140.
- Hinson, J. M., & Staddon, J. E. R. (1983a). Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior*, 39, 25–47.
- Hinson, J. M., & Staddon, J. E. R. (1983b). Matching, maximizing and hill-climbing. *Journal of the Experimental Analysis of Behavior*, 40, 321–331.
- Houston, A. I., & McNamara, J. (1981). How to maximize reward rate on two variable-interval paradigms. *Journal of the Experimental Analysis of Behavior*, 35, 367–396.
- Machado, A. (1992). Behavioral variability and frequency-dependent selection. *Journal of the Experimental Analysis of Behavior*, 58, 241–263.
- Machado, A. (1994). Polymorphic response patterns under frequency-dependent selection. *Animal Learning & Behavior*, 22, 53–71.
- Manabe, K., Staddon, J. E. R., & Cleaveland, J. M. (1997). Control of vocal repertoire by reward in budgerigars (*Melopsittacus undulatus*). *Journal of Comparative Psychology*, 111, 50–62.
- Metz, J. A. J. (1974). Stochastic models for the temporal fine structure of behaviour sequences. In D. J. McFarland (Ed.), *Motivational control systems analysis* (pp. 5–86). London: Academic Press.
- Nevin, J. A. (1969). Interval reinforcement of choice behavior in discrete trials. *Journal of the Experimental Analysis of Behavior*, 12, 875–885.
- Nevin, J. A. (1979). Overall matching versus momentary maximizing: Nevin (1969) revisited. *Journal of Experimental Psychology: Animal Behavior Processes*, 5, 300–306.
- Platt, J. R. (1979). Interresponse-time shaping by variable-interval-like interresponse-time reinforcement contingencies. *Journal of the Experimental Analysis of Behavior*, 31, 3–14.
- Shimp, C. P. (1966). Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, 9, 443–455.
- Shimp, C. P. (1969). Optimal behavior in free-operant experiments. *Psychological Review*, 76, 97–112.
- Shimp, C. P. (1973). Synthetic variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 19, 311–330.
- Shimp, C. P. (1974). Time allocation and response rate. *Journal of the Experimental Analysis of Behavior*, 21, 491–499.
- Shimp, C. P. (1981). The local organization of behavior: Discrimination of and memory for simple behavioral patterns. *Journal of the Experimental Analysis of Behavior*, 36, 303–315.
- Silberberg, A., Hamilton, B., Zirriax, J. M., & Casey, J. (1978). The structure of choice. *Journal of Experimental Psychology: Animal Behavior Processes*, 4, 368–398.
- Silberberg, A., & Zirriax, J. M. (1985). Molecular maximizing characterizes choice on Vaughan's (1981) procedure. *Journal of the Experimental Analysis of Behavior*, 43, 83–96.
- Staddon, J. E. R. (1977). On Herrnstein's equation and related forms. *Journal of the Experimental Analysis of Behavior*, 28, 163–170.
- Staddon, J. E. R., Hinson, J. M., & Kram, R. (1981). Optimal choice. *Journal of the Experimental Analysis of Behavior*, 35, 397–412.
- Todorov, J. C., Souza, D. G., & Bori, C. M. (1993). Momentary maximizing in concurrent schedules with a minimum interchangeover interval. *Journal of the Experimental Analysis of Behavior*, 60, 415–435.

Received July 15, 1997

Final acceptance July 28, 1999