MATCHING, MAXIMIZING, AND HILL-CLIMBING

JOHN M. HINSON AND J. E. R. STADDON

DUKE UNIVERSITY

In simple situations, animals consistently choose the better of two alternatives. On concurrent variable-interval variable-interval and variable-interval variable-ratio schedules, they approximately match aggregate choice and reinforcement ratios. The matching law attempts to explain the latter result but does not address the former. Hill-climbing rules such as momentary maximizing can account for both. We show that momentary maximizing constrains molar choice to approximate matching; that molar choice covaries with pigeons' momentary-maximizing estimate; and that the "generalized matching law" follows from almost any hill-climbing rule.

Key words: optimal behavior, momentary maximizing, matching law, concurrent schedules, variable interval, variable ratio, key peck, pigeons

---

Early research has shown that a hungry ass regularly allowed to choose between two piles of hay, one large and one small, will consistently pick the large. This result conforms to the matching law (Herrnstein, 1961, 1970), but is not explained by it—consistent choice of the small pile would fit as well. Two other theories do explain this result: overall maximization (e.g., Rachlin, Green, Kagel, & Battalio, 1976; Staddon & Motheral, 1978) and hill-climbing, namely, picking the best alternative at a given time (e.g., Hinson & Staddon, 1983; Shimp, 1969). Overall maximization works because consistent choice of the large pile maximizes food intake. Hill-climbing works because always picking the large pile means always picking the best alternative offered.

More recent research has shown that on concurrent variable-interval variable-interval (concurrent VI VI), variable-interval, variable-ratio (concurrent VI VR), and some other reinforcement schedules, pigeons, rats, and people allocate their responding according to the formula

$$x/y = a[R(x)/R(y)]^b, \qquad (1)$$

where x and y are the average response rates to the two alternatives, $R(x)$ and $R(y)$ the reinforcement rates obtained, and a and b con-

stants. For concurrent VI VI, a and b are both approximately unity (i.e., the matching law), although some studies find b values less than one (undermatching); for concurrent VI VR, b is near unity but a may be less than one (biased matching); for concurrent chained schedules, a is unity, but b is generally greater than unity (overmatching: Bacotti, 1977; Baum, 1974; Davison, 1981; Lea, 1981).

Equation 1, the "generalized matching law," is flexible enough to accommodate almost any monotonic empirical relationship between choice and reinforcement ratios, given usual experimental variation. Hence, Equation 1 is significant mainly for the consistent relationship between parameter values and particular procedures. The standard matching law cannot account for systematic deviations from unity in the two parameters.

Behavior on concurrent schedules has properties in addition to aggregate choice ratios: distributions of interresponse times (IRTs) to each alternative, interswitch times (i.e., choice bout lengths), and temporal relationships between the sequence of responses to each alternative. Some of these properties are affected in striking ways by procedural features. For example, if each switch (changeover) starts a timer that prevents the delivery of food for t sec after the switch (changeover delay: COD), then interswitch times increase (e.g., Baum, 1982; Hunter & Davison, 1978). As another example, IRTs to the VR alternative on concurrent VI VR are shorter than IRTs to the VI. None of these properties of choice are explicable by the matching law.

The limited applicability of matching-law formulations has led in recent years to a search for more comprehensive alternatives that can explain simple matching, power-ratio matching (Equation 1), COD and schedule-type effects, as well as exclusive choice. For example, the effects of COD on interswitch time make perfect sense from the point of view of either overall maximizing or hill-climbing; IRT differences and biased matching on concurrent VI VR follow from a form of hill-climbing (Hinson & Staddon, 1983; Staddon, 1980), and as we have seen, exclusive choice in simple situations follows from almost any optimizing theory.

Molar measures, such as the aggregate response and reinforcement rates that enter into Equation 1, are made up of molecular elements: moment-by-moment choices of one alternative or the other. Both molar matching and overall reinforcement-rate maximization when they occur must depend upon rules that relate the animal's instantaneous choice to its past history. Such a rule can be called a *strategy* or *policy* without implying conscious deliberation by the animal. When molar principles fail, the obvious place to look for an explanation is the molecular strategies that underlie them.

We can probably rule out one possibility—what we have elsewhere called the *molar comparison strategy* (Hinson & Staddon, 1983)—that is, the idea that when animals maximize, they do so by varying choice proportions from day to day, remembering the average reinforcement rate associated with each, and choosing the proportion that gives them the highest payoff rate. Memory limitations mean that animals are much more likely to maximize (and match) by adopting a moment-by-moment strategy that does not demand comparisons over long time periods.

Hill-climbing, picking the best option available from moment to moment, is one class of strategies that makes limited demands on memory, yet often achieves molar outcomes close to the theoretical optimum (Minsky, 1961). This paper discusses one such hill-climbing strategy, *momentary maximizing* (Shimp, 1966). We have previously shown that pigeons conform reasonably well to momentary maximizing on both concurrent VI VI and concurrent VI VR schedules (Hinson & Staddon, 1981, 1983). Here we show that when pigeons show momentary maximizing, molar choice proportion approximately matches reinforcement proportion; that molar choice varies with the animal's momentary estimate of payoff probabilities for the two alternatives on these two schedules; and that this relation is not an artifact of our method of measuring the animal's estimate. Hence matching, a molar relation, is probably the outcome of moment-by-moment hill-climbing. We show in the Discussion that a class of hill-climbing strategies of which momentary maximizing is a special case implies molar results conforming to Equation 1. We conclude from these simulations that Equation 1 is implied by almost any hill-climbing rule.

METHOD

This paper presents additional analyses of the experiment described in Hinson and Staddon (1983). The method is given in full in that paper, so we present only a brief summary here.

Subjects

Eight male, adult White Carneaux pigeons served.

Apparatus

All experiments were conducted in a standard aluminum and Plexiglas operant-conditioning chamber with two translucent pecking keys. The experimental contingencies and data recording were carried out by a microcomputer in an adjacent room. Data on the absolute time (to one msec) and identity of each experimental event were later transferred to a PDP 11 minicomputer for analysis.

Procedure

The pigeons received extensive training on concurrent VI VI and VR VI schedules, with random (Fleshler & Hoffman, 1962) interfood intervals and no changeover delay. The experimental conditions for all animals appear in Table 1. Sessions lasted 1 hour, excluding the time taken by food delivery.

RESULTS

Concurrent VI VI

Figure 1 shows molar matching: the relationship between the mean logarithms of choice ratios and the mean of the logarithms of ob-

Table 1
Conditions and Number of Sessions for Each Bird

| | Concurrent VI VI | | | |
| | Bird # | | | |
Condition	CO96	CO123	CO104	Sessions
1	VI 60 VI 60	VI 60 VI 60	VI 60 VI 60	90
2	VI 180 VI 60	VI 180 VI 60	VI 180 VI 60	90
3	VI 60 VI 180	VI 60 VI 180	VI 60 VI 180	60
4	VI 60 VI 60	VI 60 VI 60	VI 60 VI 60	30
	CD129	CD117	CD148	
1	VI 180 VI 60	VI 180 VI 60		30
2	VI 60 VI 180	VI 60 VI 180	VI 60 VI 180	30
3	VI 60 VI 60	VI 60 VI 60	VI 60 VI 60	17
4	VI 240 VI 60	VI 240 VI 60	VI 240 VI 60	15
5	VI 90 VI 180	VI 90 VI 180	VI 90 VI 180	16
6		VI 180 VI 30	VI 180 VI 30	15

| | Concurrent VR VI | | | |
| | Bird # | | Bird # | |
	CR129	Sessions	CR117	Sessions
1	VR 30 VI 60	30	VR 30 VI 60	30
2	VR 60 VI 60	30	VR 60 VI 60	30
3	VR 15 VI 60	30	VR 15 VI 60	30
	CR101		CR196	
1	VR 15 VI 60	40	VR 60 VI 60	20
2	VR 60 VI 60	40	VR 60 VI 180	17
3	VR 30 VI 60	15	VR 60 VI 120	18

tained reinforcement ratios of the last five sessions of each condition for each animal. Vertical bars for each point within the graph show the range of response ratios. Matching in this situation is within the normal range: Slopes of the fitted regression lines, shown for each animal, are generally less than one and the intercepts are close to zero.

There are two ways to demonstrate a relationship between matching and momentary maximizing, neither perfectly satisfactory. The first is to show that when animals conform well to the momentary maximizing rule, choice proportions are close to matching. (When animals maximize poorly, their choice proportions may or may not conform to matching, depending on what causes the poor performance. Consequently, molar choice proportion on days when maximizing is poor is less informative than choice performance on days when maximizing is good.) The second way to show a relationship between matching and momentary maximizing is to show that molar choice proportion reliably covaries with the animal's estimate of the momentary maximizing switching line. We will discuss each of these methods in turn.

Previous papers (Hinson & Staddon, 1981,

1983; Staddon, Hinson, & Kram, 1981) have shown that the relevant variables for momentary maximizing are t_1 and t_2, the times since the last responses to alternatives one and two. Momentary maximizing is defined by the switching line,

$$t_1 = \lambda_2 t_2 / \lambda_1,$$

where λ_i is the scheduled VI reinforcement rate for alternative i. The animal is performing perfectly when all instances of response one lie between the switching line and the t_1 axis in t_1/t_2 space (the clock space) and all instances of response two lie between the line and the t_2 axis.

Table 2
Procedure for obtaining the momentary maximizing estimate m.

|		Reinforcement Probability	
		$p(R\|t_1) > p(R\|t_2)$	$p(R\|t_1) < p(R\|t_2)$
Response	1	$\Sigma\|p_1 - p_2\| = a$	$\Sigma\|p_1 - p_2\| = b$
	2	$\Sigma\|p_1 - p_2\| = c$	$\Sigma\|p_1 - p_2\| = d$
		$m = (a + d)/(a + b + c + d)$	

Fig. 1. Log-log coordinate space for ratios of obtained response and reinforcement rates for each bird in concurrent VI VI. Each square is the average of the last five sessions in each condition. The vertical bars indicate the range of ratios within each average. Best-fitting regression lines are drawn through the data points. The regression equation is supplied for each panel. The ordinate represents the logarithm of responses to Schedule 1 divided by responses to Schedule 2. The abscissa represents the logarithm of reinforcers obtained from Schedule 1 divided by the reinforcers obtained from Schedule 2.
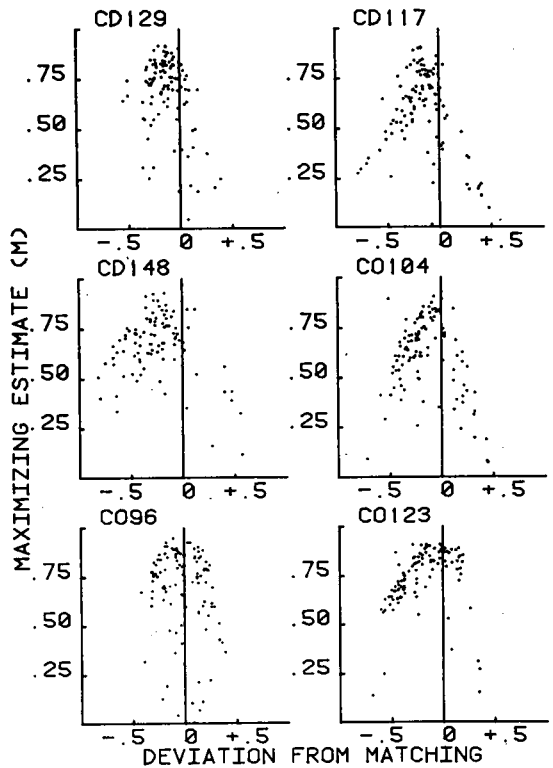

Fig. 2. Plots of *m* versus *q*, the deviation from matching, for all birds on concurrent VI VI across all conditions. The points in each graph correspond to the individual sessions from the last half of each condition for CD129, CD117, and CD148. The last fifteen sessions of each condition are shown for CO96, CO123, and CO104. The abscissa is the log deviation from matching. The ordinate is the corresponding value of *m*.

To provide a numerical estimate of the quality of momentary maximizing, we devised a quantity, *m*, analogous to "total percentage correct," which is defined in Table 2 (see Hinson & Staddon, 1983). For each correct choice, the difference in reinforcement probability for the two alternatives is positive, that is, those entries in cells *a* and *d*, whereas for each incorrect choice the difference in reinforcement probability for the two alternatives is negative, that is, those entries in cells *b* and *c*. The sums of the absolute probability difference for correct choices are then expressed as the proportion of the total amount of absolute probability difference for both correct and incorrect responses. Because *m* is a proportion, *m* is equal to one for perfect momentary maxi-

mizing, zero for momentary minimizing (consistent choice of the lower-probability alternative), and .5 for random choice.

Figure 2 shows the relationship between *m*, our maximizing estimate, and a variable that summarizes deviation of the choice proportion from perfect matching. If, by convention, we take response *x* as the majority (more frequently reinforced) response, then a quantity *q* defined as

$$q = ln\{(x/y)/[R(x)/R(y)]\} \qquad (2)$$

should equal 0 for perfect matching, < 0 for undermatching, and > 0 for overmatching. Figure 2 is a scatter plot of *m* (ordinate) vs. *q* (abscissa) for the latter part of each condition for all animals exposed to concurrent VI VI. For Birds CO96, CO123, and CO104, the plots are for the last 15 sessions of each condition. For Birds CD129, CD117, and CD148,

the plots are for the last half of each condition. Each animal, with the possible exception of CO96, shows an inverted U-shaped scatter: at high m values, q is slightly less than zero (slight undermatching); at lower m values, the outcome is variable, showing both over- and undermatching. Although these plots do not prove that good momentary maximizing forces the animal to approximate molar matching, they do illustrate that increases in the value of m result in decreases in the variability in choice proportion.

A second way to test for a relationship between maximizing and matching is through the hypothesis that molar choice proportion is directly determined by momentary maximizing: If molar choice proportion always follows momentary maximizing, then variations in choice proportion must reflect variations in the animal's estimate of the switching line. A test of this idea requires a way to estimate the animal's switching line that does not artificially force a relationship between the maximizing estimate and choice proportion.

One method is illustrated in Figure 3, which shows a sequence of responses to Alternatives 1 and 2. Consider any minority response, such as the one labeled "A" in the figure. If the animal is consistently following a momentary maximizing rule, then the quantity t_2/t_1 must be greater than the slope of his switching line, that is, greater than λ_1'/λ_2', where λ_1' and λ_2' are the animal's estimates of λ_1 and λ_2, the scheduled reinforcement rates for the two alternatives. Hence, t_2/t_1 for every minority response can provide us with an estimate of the animal's switching line. Because the switching condition requires t_2/t_1 to be greater than the slope of the switching line, our estimate will be biased; but since we are interested in correlations between *changes* in estimate and *changes* in molar choice proportion, this bias in absolute value does not matter.

The simplest way to show changes in the switching-line estimate over time is a cumulative plot of the quantity t_2/t_1. Thus, for the first minority response of the session, a point is plotted at coordinates t_1 and t_2; for the next minority response, the new t_1 and t_2 are plotted with the previous point as origin, and so on. These closely spaced points make up a cumulative curve whose t_2 axis is real time (because every minority response is recorded, so that cumulative t_2 time is total time) and whose slope is directly related to the animal's estimate of the switching line. This cumulative curve can then be compared with cumulative choice proportion, where each point represents the total number of Responses 1 and 2 to that point. We expect that for animals trained on concurrent VI VI, these two curves will parallel one another, whereas animals trained on other procedures (concurrent VI VR, for example) will show no such parallelism.

Figure 4 shows two comparisons (complete 1-hour sessions) between cumulative switching-line estimate (upper panels) and cumulative choice proportion (lower panels) for Birds CO96 and CO123. Hash marks on the lower curves mark 2-min intervals. "−1" and "+1" signify sessions just before and just after a tran-

Fig. 3. Hypothetical event diagrams for two responses showing the technique for generating the cumulative time estimate of the switching line.

Fig. 4. Comparison of curves of cumulative times since each response with curves of cumulative numbers of each response. (Upper) Cumulative times for the last session of concurrent VI 180 VI 60 (−1) and first session of concurrent VI 60 VI 180 (+1) for (left) Bird CO96 and (right) Bird CO123. (Lower) Cumulative responses for the same sessions in the upper graph. Each time axis represents an hour. Downward deflections on the cumulative response curves mark 2-min intervals.

sition from concurrent VI 180 VI 60 to VI 60 VI 180. It is obvious that the four pairs of curves are all highly similar, and this result is typical of all the comparisons we have made. To quantify this result we correlated the slope, minute by minute, of the cumulative switching-line estimate curve with the slope of the comparable cumulative choice-proportion curve for every animal and every session of concurrent VI VI. As a control, we include the same procedure for animals on concurrent VI VR. As shown earlier, time since a VR response is irrelevant to the switching rule for concurrent VR VI. Therefore, any correlation between the t_2/t_1 ratio and the ratio of choices to the alternatives would likely be simply a property of the IRT distributions and not due to momentary maximizing. The grand mean correlations, for the six pigeons exposed to concurrent VI VI and the four exposed to concurrent VI VR, are shown in Table 3: The correlations for concurrent VI VI are high and positive, those for concurrent VI VR much lower; there is almost no overlap between the two sets.

The data in Figure 4 and Table 3 are consistent with the hypothesis that choice proportion on concurrent VI VI is driven by the animal's estimate of the switching line. They do not prove it, for two reasons: the hypothesis is a causal one and we can do no more than exhibit correlations, and many choice data that do not conform to momentary maximizing are bound to show some correlation between our t_2/t_1 measure and choice proportion. For example, if animals make a choice every Δt, with probabilities p and 1-p, then if p varies in indefinite fashion, the ratio t_2/t_1 will tend to vary in a similar fashion. Thus, the alternative hypothesis that switching-line estimate is driven by some stochastic process that affects

choice proportion can never be conclusively refuted with our data.

Nevertheless, although we cannot rule out all alternative hypotheses, we can disconfirm particular hypotheses. For example, assuming that animals choose at random, then a positive correlation (of the sort given in Table 3) between switching-line estimate and cumulative choice proportion is expected. Suppose, however, that we displace the distribution of Response 1 interresponse times by a constant amount Δt (i.e., a "shifted-distribution" analysis of the type discussed at length in Hinson & Staddon, 1983). If the two choice distributions are independent, shifting each element of either distribution by a constant amount cannot have a consistent effect on the correlation between switching-line estimate and cumulative choice proportion. On the other hand, if the correlation depends upon momentary maximizing, a time shift will degrade momentary maximizing and, thereby, the correlation should be progressively reduced by increasing shifts.

Our pigeons showed the latter result, as illustrated in Figure 5. The figure plots the value of the correlation coefficient, r, against shift in .2-sec increments, from −1 to +1 sec. The figure shows three curves: the leftmost curve is for random data in 3:1 proportion, and the other two curves are for typical sessions from two pigeons on concurrent VI 60 VI 180. Both pigeons show a smoothly unimodal correlation function, with the peak at zero shift for one animal, slightly to the left for the other. The

Table 3

Mean correlation between choice proportion and switching estimate for all birds across all condtions.

| VI VI		VR VI	
CO96	.62	CR101	.20
CO123	.64	CR196	.24
CO104	.43	CR117	.32
		CR129	.49
CD129	.72		
CD117	.69		
CD148	.63		
mean	.63	mean	.31

Fig. 5. Display of IRT distributions shifted positively and negatively in time, with respect to one another. (Far left) Simulated data for random choice time with matching on concurrent VI 180 VI 60. (Middle) Last session of concurrent VI 180 VI 60 for Bird CO96. (Far right) Last session of concurrent VI 180 VI 60 for Bird CO123. The abscissa indicates the amount and direction (positive or negative) of the time shift. The ordinate indicates the correlation coefficient (r) between choice proportion and switching estimate for each time shift.

random data vary over a narrower range and show no clear peak, although, as expected, the correlations are positive—although smaller than for the pigeon data. These two data curves are typical of many sessions we have examined. Thus, the similarity between cumulative switching line and cumulative choice proportion in our concurrent VI VI data does not seem to be an artifact.

Concurrent VR VI

Figure 6 shows molar matching for the concurrent VR VI animals, in the same form as Figure 1: log response ratios vs. log obtained reinforcement ratios. Although we examined only three schedule values, matching is within the normal range; the coefficients of regression lines, which again appear in each graph, show slopes less than one, but intercepts that are generally greater than zero.

The switching line for concurrent VR VI is a horizontal line parallel to the axis for the VR response (Hinson & Staddon, 1983; Staddon, Hinson, & Kram, 1981). Nevertheless, it

is possible to compute m, our momentary maximizing estimate, in the same way as for concurrent VI VI. Figure 7 shows a scatter plot of m plotted vs. q, an estimate of matching, for all sessions for each bird exposed to concurrent VR VI (this figure parallels Figure 2). The results are similar to those shown in Figure 2: When the animal is maximizing well, molar performance is constrained close to matching; otherwise, there is variation on both sides of matching.

Because the switching line for concurrent VR VI is parallel to the VR axis, there is no estimate analogous to the switching-line estimate we used for concurrent VI VI that will permit us to show a correlation between choice proportion and switching line. Part of the problem is that momentary maximizing on concurrent VR VI depends on choice rate, whereas on concurrent VI VI, the expected choice proportion is independent of choice rate. For example, if an animal responds very slowly on concurrent VR VI, at the time of a choice the probability of reinforcement for a VI response will nearly always be higher than on the VR; hence, almost all choices will be to the VI schedule. Conversely, if choice rate is high, the VR will nearly always be the better prospect and nearly all choices should be to the


Fig. 6. Log-log coordinate space for ratios of obtained response and reinforcement rates for each bird in concurrent VR VI. Each square is the average of the last five sessions in each condition. The vertical bars indicate the range of ratios within each average. Best-fitting regression lines are drawn through the data points. The regression equation is supplied for each panel. The ordinate represents the logarithm of VR schedule responses divided by VI schedule responses. The abscissa represents the logarithm of reinforcers obtained from the VR schedule divided by the reinforcers obtained from the VI schedule.


Fig. 7. Plots of m versus q, the deviation from matching, for all birds on concurrent VR VI across all conditions. The points in each graph correspond to the individual sessions from the last half of each condition. The abscissa is the log deviation from matching. The ordinate is the corresponding value of m.

VR schedule. For concurrent VI VI, on the other hand, the choice proportion is independent of choice rate—although, given the smaller probability differences at very low choice rates, when both payoff probabilities are close to unity, one might expect a trend towards indifference (i.e., undermatching) as choice rates fall to low levels.

Since choice rate and switching-line estimate are confounded for concurrent VR VI, we cannot show that choice proportion on concurrent VR VI is correlated with the animal's estimate of the switching line in the same fashion as concurrent VI VI. However, momentary maximizing dictates that the proportion of VR responses should be positively correlated (and proportion of VI responses negatively correlated) with the absolute choice rate (i.e., VI + VR response rates). This result is illustrated in a previous paper (Hinson & Staddon, 1983, Table 6): Proportion of VI responses was consistently negatively correlated with overall choice rate on concurrent VR VI; as a control condition, the comparable correlations on concurrent VI VI were close to zero. Thus, in general, momentary maximizing for concurrent VR VI results in predictable changes in choice proportion that are different for different types of schedules.

DISCUSSION

Previous theoretical papers have shown that periodic choice, combined with momentary maximizing, implies a good approximation to molar matching (Shimp, 1969; Staddon, Hinson, & Kram, 1981). The present data show the same pattern in the aperiodic behavior of pigeons on concurrent VI VI and concurrent VR VI schedules: The better the birds follow momentary maximizing, the closer the approximation to molar matching.

In addition, we are able to show a close association on concurrent VI VI between the animal's estimate of the momentary maximizing switching line and molar choice proportion: Changes in choice proportion closely match changes in switching-line estimate. The correlation between the two is much closer on concurrent VI VI than on concurrent VI VR or a random simulation, and the correlation is impaired by manipulations that also impair momentary maximizing (the shifted-distribu-

tion test). Thus the correlation between momentary maximizing estimate and molar choice proportion does not seem to be an artifact of the way we estimate momentary maximizing.

This test cannot be applied on concurrent VR VI because if animals momentary maximize, choice proportion should (and does) vary with overall choice rate. Nevertheless, the feedbacks inherent in ratio schedules typically ensure that such variations in choice proportion do not yield comparable deviations from matching. For example, if choice rate is high, then VI responses occur only when $p(R|t)$ is approximately equal to $p(R)$ for the ratio schedule, that is, 1/the average ratio value. Overall, $p(R|t)$ for the VI must equal $R(x)/x$, where x is the VI response rate. Since 1/mean ratio value is equal to $R(y)/y$, where y is the average ratio response rate, momentary maximizing with a high choice rate implies matching (cf. Staddon, 1980).

On the other hand, if choice rate is low, ratio responses will occur rarely, but when they do, it will be because $p(R)$ for the ratio response is equal to or greater than $p(R|t)$ for an interval response. If choice rate is low, we cannot be sure that when a ratio response occurs, reinforcement probability will in every case be equal to the average for the VI; sometimes it will be considerably greater. Thus, when choice rate is low, fewer responses may be made to the ratio schedule than implied by matching, that is, a bias in favor of the VI. By these arguments, the deviation from matching on concurrent VR VI should tend to bias in favor of the VR when choice rates are high, but in favor of the VI when choice rates are low—at low levels of food deprivation, for example.

There seem to be no relevant data on these predictions. Most of the data on concurrent VR VI schedules involves COD procedures (e.g., Herrnstein & Heyman, 1979). It is important to note that a COD profoundly changes the feedback properties of a schedule. These feedback changes result in concomitant behavioral changes such as greatly increased run lengths to both keys. Unless the theoretical analysis makes qualifying assumptions, it is not possible to predict choice proportion. At present, we have been unable to find a nonarbitrary way to extend our theoretical analysis to COD procedures.

These refinements apart, it is clear that momentary maximizing both implies matching (simple or biased) on concurrent VR VI and implies that matching will be relatively unaffected by variations in absolute choice rate and the associated variations in choice proportion. Moreover, the balance of factors pushing bias towards or away from the ratio side is such that under normal conditions, given typical variation in choice rates, we might expect relatively unbiased matching.

The inevitable conclusion is that on concurrent VI VI and concurrent VR VI, all roads lead to matching. Although correlation is not the same as causation, every test we have applied yields results consistent with the idea that molar choice proportion on concurrent VI VI is driven by momentary maximizing. Feedbacks intrinsic to the procedure ensure that if the animal momentary maximizes, choice proportion will still lead to matching on concurrent VR VI even over the normal range of rate variation on these schedules.

It should be noted that the current theoretical analysis applies to discrete responses; it is therefore directed toward response rather than time allocation matching. However, as the model of momentary maximizing is based on the time since each response, its prediction should be about the same for time-based results. Such an extension of the model would require arbitrary assumptions about behavior in order to generate predictions. Since we have not worked out the particulars of the time-based case, we prefer not to speculate further about possible connections between response and time-based matching.

Other data consistent with momentary maximizing as a widespread hill-climbing strategy have been reviewed elsewhere (e.g., Staddon, 1980; Staddon, Hinson, & Kram, 1981), but we should add as particularly compelling evidence an experiment by Lea (1976) using a concurrent ratio schedule with a titration procedure. The ratio value for Alternative 1 was held constant, whereas the ratio on Alternative 2 decreased when Alternative 1 was chosen. Thus, by varying the frequency of Alternative 1 and Alternative 2 choices, the animal could select almost any Alternative 2 ratio. Overall maximizing under these conditions implies that Alternative 1 should be chosen frequently, so as to keep Alternative 2 at a low value. Mo-

mentary maximizing implies a pattern of choice such that Alternative 2 should average the same probability of reinforcement as Alternative 1, which is what Lea found.

Momentary maximizing is also consistent with the shifts in preference reported by Herrnstein and Vaughan (1980) in concurrent-ratio experiments in which pairs of different ratio values were alternated between the two choices: Although exclusive responding to either choice is consistent with matching, their animals consistently shifted towards the lower ratio. Momentary maximizing is also consistent with the results of a complex procedure (Herrnstein & Vaughan, 1980) in which pigeons could achieve higher overall payoff rates by allocating responses disproportionately to one alternative; as in Lea's experiment, unconstrained overall maximizing predicts disproportionate allocation, whereas momentary maximizing predicts what was found—approximate matching.

Other Hill-Climbing Rules

Momentary maximizing is not the only possible hill-climbing rule appropriate for concurrent schedules. We do not wish to explore an infinite field, but it may be instructive to conclude with a brief discussion of one other class of hill-climbing rules and show that it, too, leads to matching of one sort or another on concurrent schedules.

Consider the conflicting demands placed on an animal freshly placed in a choice situation: He needs to explore each alternative, and he needs to be sensitive to the consequences of his choices. A simple switching rule that accomplishes both objectives is

$$t_1 + A/T_1 = t_2 + A/T_2, \qquad (3)$$

where t_i is time since the last response to alternative i, T_i is time since the last reinforcement for alternative i, and A is a constant that represents the relative importance to choice of times t and T. Suppose that constant A is zero at the beginning of training; the resulting switching line, $t_1 = t_2$, prescribes simple alternation ("least-recent" choice) between the two alternatives—an appropriate exploratory strategy. If the occurrence of food causes some increase in the value of A, then immediately after getting fed for response i, the animal is much more likely to repeat response i, because the

quantity A/T_i will be large; as time elapses, this quantity declines and the least-recent rule gains control over behavior—until the next food delivery, when the tendency to repeat response i will again increase. This model does not result in good momentary maximizing as defined earlier (Hinson & Staddon, 1983). The values of m are consistently low for this choice strategy due to perseveration of responding after a reinforcer.

We simulated on a computer the effect on molar choice of responding according to Equation 3, with various values of parameter A. Choice time was random throughout. We simulated choice performance, defined by A parameter and absolute choice rate (defined as the ratio of minority VI reinforcement rate to choice rate: .02 is about the midpoint of the normal range), over five different concurrent VI VI schedules, ranging from 1:1 to 1:6 in relative scheduled reinforcement rate. At each value of A and choice rate, regression lines were fitted to the logarithms of choice and reinforcement ratios from 500,000 program cycles. Some of the results are shown in Table 4. In every case, the fit of the regression line was excellent: Equation 1 fitted impeccably all the data generated by this class of hill-climber. Matching was always unbiased, as one might expect: Constant a in Equation 1 was always close to one. And in general the tendency was towards undermatching, $b < 1$, with b lower at lower A values, when choice approached simple alternation.

The point of this simulation is not that Equation 3 fits actual performance on, for example, concurrent VI VI—in our previous report (Hinson & Staddon, 1983) we in fact found no evidence whatsoever for a reinforcement-recency effect of the sort implied by Equations 3 with $A \neq 0$—although such effects

are seen frequently under other conditions. Our point is that essentially any hill-climbing rule sufficient to prevent exclusive choice on concurrent VI VI and concurrent VR VI must yield data consistent with Equation 1. Hence, the fact of molar power-ratio matching tells us almost nothing about what sort of choice rules are determining performance. The evidence discussed in the rest of this paper, and in our earlier report, suggests that, for the most part, pigeons' choices are determined by momentary maximizing.

REFERENCES

Bacotti, A. V. Matching under concurrent fixed-ratio variable-interval schedules of food presentation. *Journal of the Exeprimental Analysis of Behavior,* 1977, **27**, 171-182.

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

Baum, W. M. Choice, changeover, and travel. *Journal of the Experimental Analysis of Behavior,* 1982, **38**, 35-49.

Davison, M. Choice between concurrent variable-interval and fixed-ratio schedules: A failure of the generalized matching law. In C. M. Bradshaw, E. Szabadi, & C. F. Lowe (Eds.), *Quantification of steady-state operant behaviour.* Amsterdam: Elsevier/North-Holland, 1981.

Fleshler, M., & Hoffman, H. S. A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior,* 1962, **5**, 529-530.

Herrnstein, R. J. Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior,* 1961, **4**, 267-272.

Herrnstein, R. J. On the law of effect. *Journal of Experimental Analysis of Behavior,* 1970, **13**, 243-266.

Herrnstein, R. J., and Heyman, G. M. Is matching compatible with reinforcement maximization on concurrent variable interval, variable ratio? *Journal of the Experimental Analysis of Behavior,* 1979, **31**, 209-223.

Herrnstein, R. J., & Vaughan, W., Jr. Melioration and behavioral allocation. In J. E. R. Staddon (Ed.), *Limits to action: The allocation of individual behavior.* New York: Academic Press, 1980.

Hinson, J. M., & Staddon, J. E. R. Maximizing on interval schedules. In C. M. Bradshaw (Ed.), *Recent developments in the quantification of steady-state operant behavior.* Amsterdam: Elsevier/North-Holland, 1981.

Hinson, J. M., & Staddon, J. E. R. Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior,* 1983, **39**, 25-47.

Hunter, I. W., & Davison, M. C. Response rate and changeover performance on concurrent variable-in-

Table 4

Parameters of matching lines (Equation 1) for switching rule of Equation 3.

		Choice Rate					
		.1			.01		
		Slope (b)	*Intercept (log a)*	*r²*	*Slope (b)*	*Intercept (log a)*	*r²*
A	60	.78	−.05	.99	.44	−.01	.99
(thousands)	40	.73	−.04	.99	.38	−.01	.99
	20	.62	−.04	.99	.26	.01	.98

terval schedules. *Journal of the Experimental Analysis of Behavior,* 1978, **29,** 535-556.

Lea, S. E. G. Titration of schedule parameters by pigeons. *Journal of the Experimental Analysis of Behavior,* 1976, **25,** 43-54.

Lea, S. E. G. Correlation and contiguity in foraging behaviour. In P. Harzem & M. D. Zeiler (Eds.), *Advances in analysis of behaviour* (Vol. 2). *Predictability, correlation, and contiguity.* Chichester and New York: Wiley, 1981.

Minsky, M. Steps toward artificial intelligence. *Proceedings of the Institutes of Radio Engineers,* 1961, **49,** 10-30.

Rachlin, H., Green, L., Kagel, J. H., & Battalio, R. C. Economic demand theory and psychological studies of choice. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 10). New York: Academic Press, 1976.

Shimp, C. P. Probabilistically reinforced choice behav-

ior in pigeons. *Journal of the Experimental Analysis of Behavior,* 1966, **9,** 443-455.

Shimp, C. P. Optimal behavior in free-operant experiments. *Psychological Review,* 1969, **76,** 97-112.

Staddon, J. E. R. Optimality analyses of operant behavior and their relation to optimal foraging. In J. E. R. Staddon (Ed.), *Limits to action: The allocation of individual behavior.* New York: Academic Press, 1980.

Staddon, J. E. R., Hinson, J. M., & Kram, R. Optimal choice. *Journal of the Experimental Analysis of Behavior,* 1981, **35,** 397-412.

Staddon, J. E. R., & Motheral, S. On matching and maximizing in operant choice experiments. *Psychological Review,* 1978, **85,** 436-444.