

RAPID ACQUISITION IN CONCURRENT CHAINS: EVIDENCE FOR A DECISION MODEL

RANDOLPH C. GRACE AND ANTHONY P. MCLEAN

UNIVERSITY OF CANTERBURY

Pigeons' choice in concurrent chains can adapt to rapidly changing contingencies. Grace, Bragason, and McLean (2003) found that relative initial-link response rate was sensitive to the immediacy ratio in the current session when one of the terminal-link fixed-interval schedules was changed daily according to a pseudorandom binary sequence (e.g., Schofield & Davison, 1997). The present experiment tested whether the degree of variation in delays across sessions had any effect on acquisition rate in Grace et al.'s (2003) rapid-acquisition procedure. In one condition ("minimal variation"), the left terminal link was always fixed-interval 8 s and the right terminal link was either fixed-interval 4 s or fixed-interval 16 s. In the other condition ("maximal variation"), a unique pair of fixed-interval values was used in each session. Responding was sensitive to the current-session immediacy ratio in both conditions, but across subjects there was no systematic difference in sensitivity. These results challenge the view that initial-link responding in the rapid-acquisition procedure is determined by changes in the learned value of the terminal-link stimuli, and suggests instead that a process resembling categorical discrimination may control performance. A decision model based on the assumption that delays are categorized as short or long relative to the history of delays provided a good account of the data and shows promise in being able to explain other choice phenomena.

Key words: choice, concurrent chains, rapid-acquisition procedure, conditioned reinforcement value, generalized matching, categorical discrimination, key peck, pigeons

What causes matching in concurrent choice procedures? Two possibilities were identified by Williams (1994): responses are emitted in proportion to their strengths, implying that matching itself is fundamental; and a "winner take all" rule in which responses were made to whichever alternative had the greatest value at the present moment, with matching the result in the aggregate (cf. Gallistel & Gibbon, 2000). Four decades of research using steady-state methodologies has stimulated the development of quantitative models that are able to describe individual performance to a high degree of accuracy (e.g., Baum, 1974a; Davison & Jenkins, 1985; Grace, 1994; Mazur, 2001), but have not yielded a clear answer to the question of what causes matching. Thus it is not surprising that acquisition is a topic of increasing research interest, because understanding the dynamics of choice in transition may yield insights into the processes responsible for steady-state behavior.

Most research on choice acquisition has used concurrent schedule procedures, and has

demonstrated that response allocation in non-human subjects can adapt quickly to changing contingencies. For example, Schofield and Davison (1997) trained pigeons on a concurrent variable-interval (VI) VI schedule in which the relative reinforcer rate varied unpredictably across sessions according to a 31-step pseudorandom binary sequence (PRBS). Schofield and Davison performed a regression analysis that quantified the control over response allocation exerted by the current and prior sessions' reinforcer ratios. They found that after three exposures to the PRBS (i.e., 93 sessions), response allocation was determined entirely by the reinforcer ratio for the current session, with no detectable effect of the ratio from prior sessions (cf. Davison & McCarthy, 1988; Hunter & Davison, 1985). Because the sensitivity levels in a generalized-matching analysis were comparable to those obtained in steady-state studies (that typically include five or more conditions, each lasting for approximately 30 sessions), Schofield and Davison suggested that the PRBS procedure might prove to be an attractive alternative to traditional designs for measuring sensitivity to reinforcement.

Grace, Bragason, and McLean (2003) attempted to determine whether similar results could be obtained with a concurrent-chains procedure. In concurrent chains, pigeons

These data were presented at the New Zealand Behaviour Analysis Symposium, Auckland, August 2003.

Address correspondence and reprint requests to Randolph C. Grace, University of Canterbury, Department of Psychology, Private Bag 4800, Christchurch, New Zealand (e-mail: randolph.grace@canterbury.ac.nz).

doi: 10.1901/jeab.2006.72-04

respond during a choice phase in which concurrent schedules operate (initial links). Responses occasionally produce access to one of two mutually exclusive, distinctively signaled outcome schedules (terminal links). After reinforcement is earned during a terminal link, the initial links are reinstated. Response allocation during the initial links can be viewed as a measure of the relative value or effectiveness of the terminal-link stimuli as conditioned reinforcers.

In their Experiment 1, Grace et al. (2003) trained pigeons on a concurrent-chains procedure in which the terminal link associated with the left initial link was always fixed-interval (FI) 8 s. On any session, the terminal-link associated with the right initial link was either FI 4 s or FI 16 s, according to a PRBS similar to that used by Schofield and Davison (1997). Extended training in either of these conditions would produce strong preference for the shorter FI; in an analysis of archival data, Grace (1994) reported that generalized-matching exponents for sensitivity to relative immediacy averaged 2.92 for concurrent chains with FI terminal links. Grace et al. found that by the third PRBS exposure (i.e., Sessions 62 to 93), initial-link response allocation for all pigeons tracked the immediacy ratio in the current session, with estimates of sensitivity ranging from 0.47 to 1.84 across subjects ($M = 1.04$, $SD = 0.57$). Their results demonstrated that pigeons' response allocation could adapt quickly to unpredictable changes in the terminal-link immediacy ratio across sessions, although sensitivity levels were variable across subjects and lower than those obtained in steady-state designs.

In their Experiment 2, Grace et al. (2003) explored the effects of having a wider range and potentially infinite number of delays for the changing terminal link. The same pigeons were used, training began immediately after the conclusion of Experiment 1, and the left terminal-link schedule remained FI 8 s. A new value for the right terminal-link FI schedule was arranged for each session. These values were determined by a pseudorandom number generator, with the constraint that delays ranged between 2 s and 32 s and immediacy ratios were uniformly distributed (in logarithmic terms) between 1 to 4 and 4 to 1. Thus the terminal-link schedule for the right key could take any of an infinite number of values for

a given session. The question was whether pigeons' ability to track daily changes in the immediacy ratio would be disrupted. Grace et al., however, observed no systematic change in sensitivity to relative immediacy over the entire duration of the experiment (61 sessions). Across subjects, sensitivity to the current-session immediacy ratio ranged from 0.52 to 1.55, similar to the values obtained in Experiment 1. This result was unexpected, and suggested that there may be no savings in terms of acquisition by restricting the changing terminal-link to two values (4 s and 16 s). However, this comparison is confounded with a potential order or practice effect, because the two-value condition from Experiment 1 was not replicated.

Because sensitivity did not change over the course of Experiment 2, Grace et al. (2003) summarized the data in terms of scatterplots relating the log initial-link response ratio in each session to the log of the changing terminal-link delay. For 1 subject, there were two distinct clusters of data points, depending on whether the changing delay was less than or greater than 8 s (see their Figure 10). Sensitivity to immediacy was relatively high between clusters, but relatively low within clusters. This clustering suggests a categorical discrimination process, which is provocative because it would be more consistent with a winner take all rather than matching rule for concurrent responding. However, data for the other 3 pigeons did not show the same degree of clustering, so the generality of this result is unclear.

The present research investigated one question posed by results of Grace et al. (2003), namely whether the number of delays used across sessions affects acquisition of preference. Specifically, we compared performance in a condition identical to Grace et al.'s Experiment 1, in which one terminal-link was constant and the other changed unpredictably between two values, with a condition in which different delays were used for both terminal links in each session. We called these conditions minimal variation (one delay changing, two immediacy ratios) and maximal variation (a potentially infinite number of delays and immediacy ratios), respectively. If acquisition depends on changes in the learned value of the terminal-link stimuli, then preference for the shorter terminal link (as measured by

sensitivity) should be greater in the minimal-variation condition. The reason is that the value of the FI 8-s terminal link stimulus should never change, and repeated training with FI 4 s and FI 16 s should result in more rapid acquisition to one of two asymptotic values for the other terminal-link stimulus. By contrast, if response allocation is determined by a categorical decision process (e.g., a “choose shorter” or winner take all rule), then there may be no systematic difference in sensitivity between the minimal- and maximal-variation conditions. We planned to look for differences in terms of both aggregate performance (i.e., overall sensitivity), as well as changes in sensitivity within sessions.

METHOD

Subjects

Subjects were 4 pigeons of mixed breed, numbered 161, 162, 163, and 164, and maintained at 85% of their free-feeding weights (± 15 g) through appropriate postsession feeding. Subjects were housed individually in a vivarium with a 12:12 hr light/dark cycle (lights on at 6 a.m.). Water and grit were freely available in the home cages. Subjects were experienced with a variety of experimental procedures.

Apparatus

Four standard three-key operant chambers were used, measuring 32 cm deep by 34 cm wide by 34 cm high. The keys were 21 cm above the floor and arranged in a row 10 cm apart. In each chamber there was a houselight located above the center key, and a grain magazine with a 5 cm by 5.5 cm aperture was centered 6 cm above the floor. The houselight provided general illumination at all times except during reinforcement delivery. The magazine contained wheat and was illuminated during reinforcement. A force of approximately 0.15 N was necessary to operate each key. Each chamber was enclosed in a sound-attenuating box, and ventilation and white noise were provided by an attached fan. Experimental events were controlled through a microcomputer and MED-PC[®] interface located in an adjacent room.

Procedure

Because subjects were experienced, training started immediately in the first condition.

Sessions were conducted daily at approximately the same time (12 p.m.) with few exceptions.

A concurrent-chains procedure was used throughout. Sessions ended after 72 initial- and terminal-link cycles had been completed or 70 min elapsed, whichever occurred first. At the start of each cycle, the side keys were illuminated white to signal the initial links. An initial-link response produced an entry into a terminal link if (a) it was made to the preselected key, (b) an interval selected from a VI 10-s schedule had elapsed, and (c) a 1-s changeover delay (COD) was satisfied (i.e., at least 1 s had elapsed following a changeover to the side for which terminal-link entry was arranged). Entries were assigned randomly to the left or right terminal link, with the constraint that out of every six cycles, three were assigned to each terminal link.

The VI 10-s initial-link schedule did not begin timing until after the first response in each cycle, to allow postreinforcement pauses after the completion of the previous terminal link to be excluded from initial-link time. The initial-link schedule contained 12 intervals constructed from an exponential progression (Fleshler & Hoffman, 1962). Separate lists of intervals were maintained for cycles in which entries to the left or right terminal link had been arranged. Lists were sampled without replacement so that all 12 intervals would be used three times for both the left and right terminal links each session.

Terminal-link entry was signaled by changing the color of the side key (left key to red, right key to green), coupled with darkening of the other key. Terminal-link responses were reinforced with 3-s access to grain according to FI schedules. During reinforcement, the only illumination in the chamber was the magazine light.

The two conditions in the experiment differed in terms of the degree of variation in the terminal-link FI schedule values. In the minimal-variation condition, the FI schedule value for the red (left) terminal link was always 8 s, and the value for the green (right) terminal link was either 4 s or 16 s. The maximal-variation condition was designed to be comparable to the minimal-variation condition in that (a) the left and right alternatives were equally often associated with the shorter-delay terminal link; (b) the expected log immediacy ratios were $\log(1/2)$ and $\log(2)$,

respectively, for sessions in which the right and left alternatives were associated with the shorter delay; and (c) the overall average terminal-link delay was 9 s. The terminal-link schedule values for each session were determined by a pseudorandom number generator subject to the constraints that the log immediacy ratios were uniformly distributed, across sessions, between $\log(4)$ and $\log(1/4)$, and the delays summed to 18 s for each session. Delays for each session were computed as follows. First, a random number x was sampled from a uniform distribution such that $\log(1/4) \leq x \leq \log(4)$. The immediacy ratio then was computed as $IR = 10^x$. The delays were then obtained as $D_1 = 18 \cdot (IR) / (IR + 1)$, and $D_2 = 18 / (IR + 1)$. Whether the shorter or longer delay was assigned to the left key was determined by a 31-step PRBS. In this way, the expected log immediacy ratio for the maximal-variation condition was the same as that in the minimal-variation condition, and log immediacy ratios for individual sessions were "unbiased" in that they were equally likely to be more or less extreme than $\log(2)$ or $\log(1/2)$. Thus the critical difference between the conditions was the degree of variability in immediacy ratios across sessions.

In both conditions, the position of the shorter terminal-link delay was changed across sessions according to a PRBS. The PRBS consisted of 31 steps and was the same as that used by Hunter and Davison (1985). For the 31 sessions, the position of the shorter terminal link was Left (L), Right (R), R, R, L, L, R, L, L, L, R, L, R, L, R, R, R, R, L, R, R, L, R, L, L, R, R, L, L, L, L. Each condition consisted of three PRBS presentations (93 sessions). Pigeons 161 and 162 completed the maximal-variation condition first, followed by the minimal-variation condition. The reverse order was arranged for Pigeons 163 and 164. The maximal-variation condition was replicated (one PRBS presentation) for Pigeons 161 and 162 after the minimal-variation condition had been completed. Owing to a programming error, the minimal-variation condition was not replicated for Pigeons 163 and 164. Statistical tests used the .05 significance level.

RESULTS

Data were analyzed with a generalized-matching model that incorporates the effects

of the current- and prior-session immediacy ratios on initial-link response allocation (cf. Schofield & Davison, 1997):

$$\log \frac{B_{0L}}{B_{0R}} = a_0 \log \frac{1/D_{0L}}{1/D_{0R}} + a_1 \log \frac{1/D_{1L}}{1/D_{1R}} + a_2 \log \frac{1/D_{2L}}{1/D_{2R}} + \dots + \log b. \quad (1)$$

In Equation 1, B and D refer to initial-link response rate and terminal-link delay, respectively, with subscripts for choice alternative (L and R) and lag (0 through 9, where 0 refers to the current session). The parameters $a_0 \dots a_9$ quantify sensitivity to reinforcer immediacy (i.e., reciprocal of delay) at each lag, and b is a bias parameter. Multiple regressions were conducted to obtain estimates of sensitivity coefficients from Lag 0 (i.e., current session) through Lag 9, and bias. Separate analyses were conducted for the three PRBS presentations in each condition, as well as the replication of the maximal-variation condition. The programmed immediacy ratios were used in the multiple regressions, because the obtained delays were quite close to those programmed in all cases.

Figure 1 shows sensitivity to immediacy for Lag 0 through Lag 9 for all subjects and conditions. Overall, sensitivity was always greatest at Lag 0 (the only exception was Pigeon 161's first PRBS presentation in the maximal-variation condition). Further, sensitivity at Lag 0 increased across PRBS presentations within conditions and tended to decrease for the other lags. Lag 0 coefficients were always positive and statistically significant, although occasionally coefficients for other lags also were significant. For example, for Pigeon 161, Lag 0 through Lag 4 coefficients were significant for the first PRBS presentation in the maximal-variation condition, Lag 0 and Lag 1 for the second presentation, and Lag 0 through Lag 3 for the third. When the maximal-variation condition was replicated, only the Lag 0 coefficient was significant for either Pigeon 161 or 162.

Lag 0 coefficients (excluding replications for Pigeons 161 and 162) were entered into a $2 \times 3 \times 10$ repeated-measures analysis of variance (ANOVA) with condition (minimal/maximal), PRBS presentation, and lag as

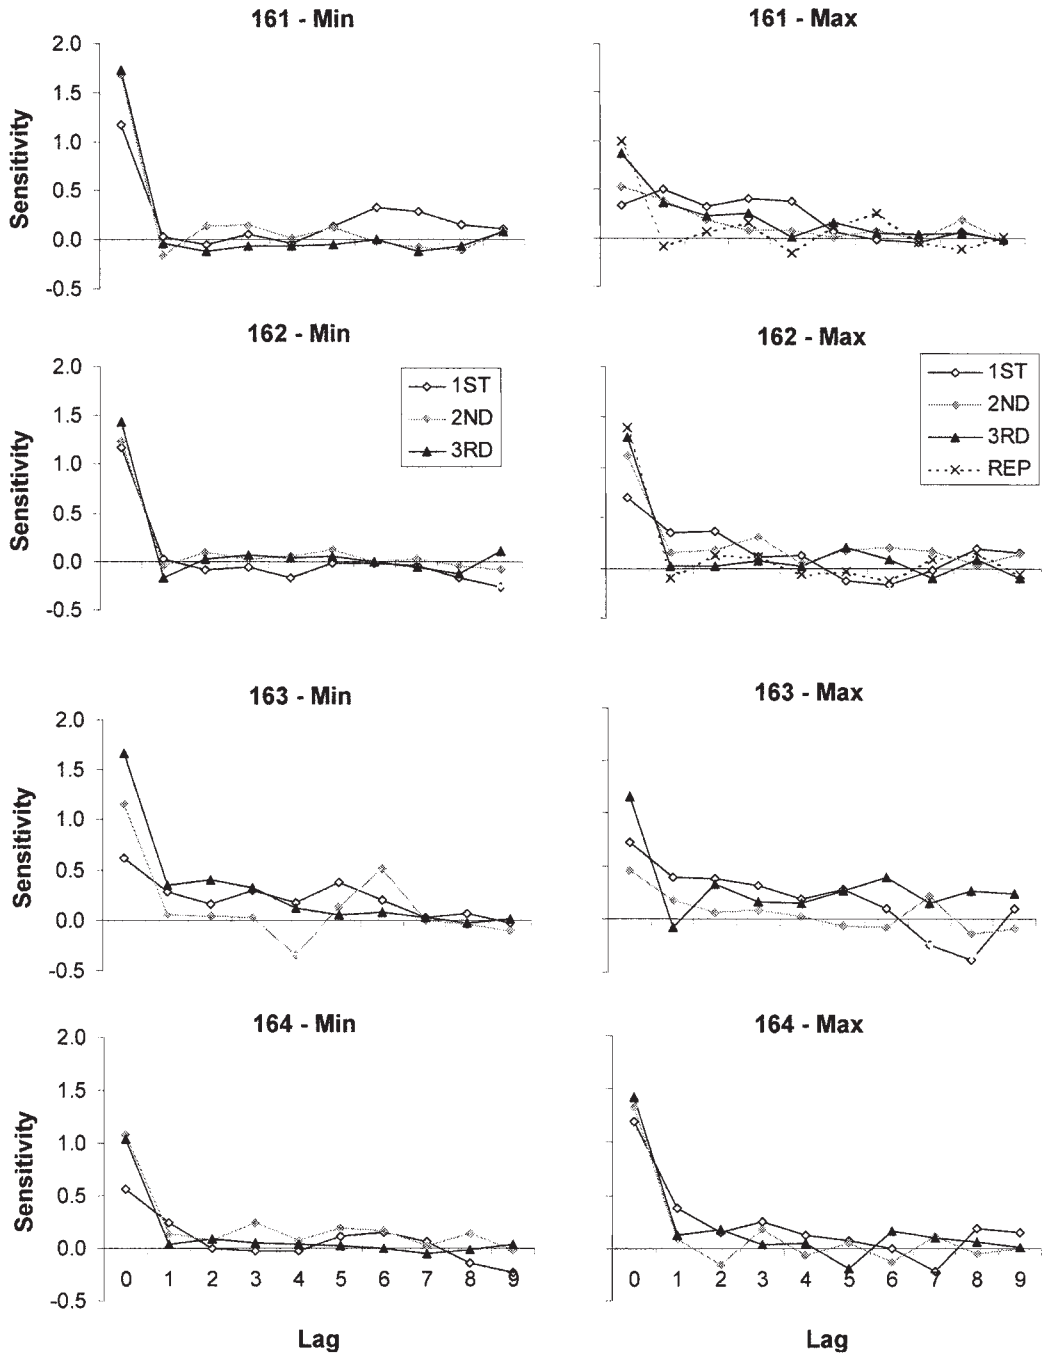


Fig. 1. Sensitivity to immediacy for Lag 0 through Lag 9 when Equation 1 was fitted to data from each PRBS presentation in the minimal-variation (left panels) and maximal-variation (right panels) conditions. The three PRBS presentations in each condition (and replication in the maximal-variation condition for Pigeons 161 and 162) are marked as noted in the legend.




Fig. 2. Lag 0 sensitivity to immediacy coefficients across successive PRBS presentations for all pigeons in the order in which they were conducted. MAX-1, -2, -3, and MIN-1, -2, -3 indicate the three PRBS presentations for the maximal-variation and minimal-variation conditions, respectively. MAX-REP denotes the maximal-variation replication (Pigeons 161 and 162 only).

factors. The main effect of condition was not significant, $F(1, 3) = 1.11$, *ns*, but the main effect of lag, $F(9, 27) = 139.76$, and the lag \times presentation and condition \times lag \times presentation interactions were significant, $F(18, 54) = 5.74$ and $F(18, 54) = 1.85$, respectively. The main effect confirms that sensitivity was greatest for Lag 0, and the lag \times presentation interaction indicates that sensitivity increased across PRBS presentations for Lag 0 but not the other lags. For the three-way interaction, analysis of simple effects found that sensitivity for Lag 2 through Lag 4 was greater in the maximal-variation condition for the first PRBS presentation only. That no significant main effect of condition was found suggests that there was no systematic difference in Lag 0 sensitivity across subjects, although it should not be interpreted to mean that sensitivities were equal in the two conditions.

We examined Lag 0 coefficients in more detail in a further attempt to find evidence for consistent differences in sensitivity between the minimal- and maximal-variation conditions. Figure 2 displays Lag 0 coefficients for all subjects in the order in which conditions were completed. Data for 2 subjects (Pigeons 161 and 163) showed that Lag 0 sensitivity was greater in the minimal-variation condition. For Pigeon 161, sensitivity increased across the three PRBS presentations in the maximal-

variation condition, continued to increase across the minimal-variation condition, but then decreased when the maximal-variation condition was replicated. Sensitivity increased across the minimal-variation condition for Pigeon 163, but then dropped for the first PRBS presentation in the maximal-variation condition and failed to reach the level achieved at the end of the minimal-variation condition. However, data for the other 2 subjects (Pigeons 162 and 164) did not reveal greater sensitivity in the minimal-variation condition. For Pigeon 162, sensitivity increased across the three PRBS presentations in the maximal-variation condition, but then remained at approximately the same level across the minimal-variation condition and the maximal-variation replication. For Pigeon 164, sensitivity increased from the first to second PRBS presentations in the minimal-variation condition, and then further increased across the maximal-variation condition. Thus across subjects there was no systematic difference in Lag 0 sensitivity between the minimal- and maximal-variation conditions.

Average Lag 0 coefficients for the third PRBS presentation can be compared with results from Grace et al. (2003) to provide an index of overall performance after a consistent amount of training (93 sessions). Across

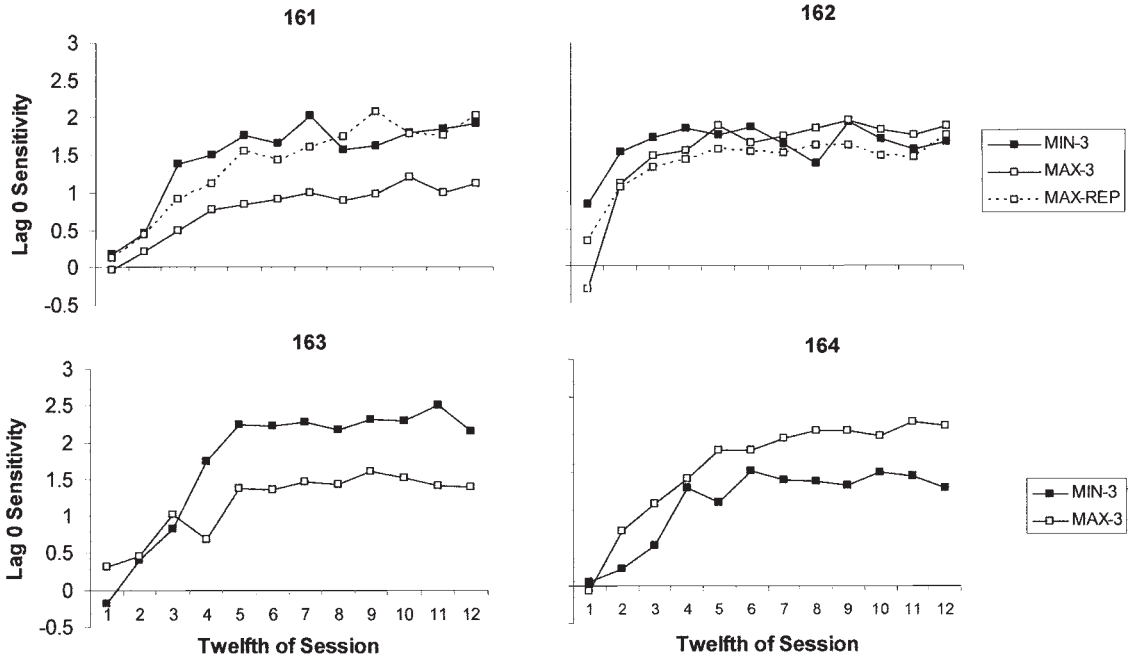


Fig. 3. Lag 0 sensitivity to immediacy coefficients across successive session twelfths for the third PRBS presentations in the maximal- (MAX-3) and minimal- (MIN-3) variation conditions and the maximal-variation replication (MAX-REP).

subjects, average Lag 0 sensitivities in the minimal- and maximal-variation conditions, respectively, were 1.47 [$SD = 0.27$] and 1.19 [$SD = 0.21$]. These values are somewhat greater than the average of 1.04 [$SD = 0.57$] obtained in Grace et al.'s Experiment 1, which was equivalent to the minimal-variation condition.

The preceding analyses were based on data pooled within each session, and it is possible that consistent differences between the minimal- and maximal-variation conditions might emerge if change in preference during sessions is examined. Figure 3 shows within-session changes in Lag 0 sensitivity for all subjects for the third PRBS presentation of each condition and replications. For this analysis, data were aggregated across all 31 sessions for each session twelfth (i.e., six cycles). It is clear from Figure 3 that sensitivity increased within session for all conditions. Moreover, individual differences in Figure 3 are correlated with those in Figure 2. For example, both figures show that sensitivity was greater in the minimal-variation condition for Pigeon 163 but in the maximal-variation condition for Pigeon 164. It also is notable that sensitivities in the maximal-variation

replication for Pigeons 161 and 162 are approximately equal to those in the minimal-variation condition. Thus within-session analyses revealed no systematic differences between the minimal- and maximal-variation conditions.

The failure to find a systematic difference in sensitivity is surprising, and suggested that we should examine performance in the maximal-variation condition in more detail. Figure 4 shows scatterplots of log initial-link response ratio as a function of the log immediacy ratio, for all subjects for the third PRBS presentation in the maximal-variation condition, as well as the replications for Pigeons 161 and 162. Each data point represents performance from a single session. For all subjects, the log initial-link response ratio increased as a function of the log immediacy ratio, but two patterns were evident. For the third PRBS presentation for Pigeons 161 and 163, and for the third PRBS presentation and replication for Pigeon 162, data are scattered approximately randomly around the regression line. However, for the Pigeon 161 replication and for Pigeon 164, the data show systematic deviations. For these subjects, data points tended to fall into two major clusters, depending on whether the left

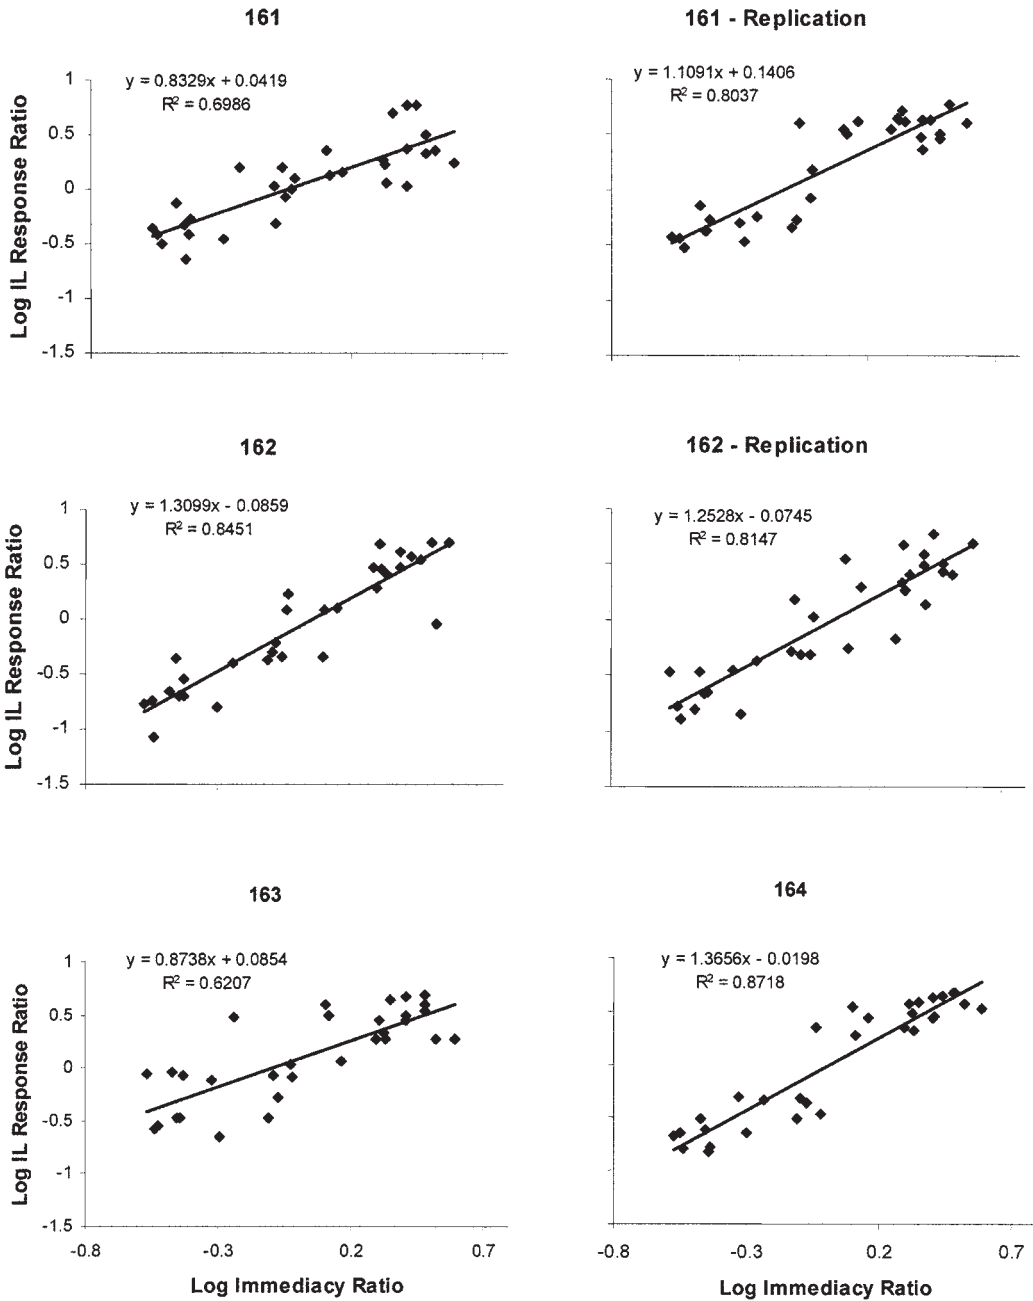


Fig. 4. Log initial-link response ratios as a function of log immediacy ratios for the third PRBS presentation in the maximal-variation condition for all subjects (including replications for Pigeons 161 and 162). Each data point represents performance from a single session. Parameters and variance accounted for by linear regression (heavy lines) also are shown.

or right key was favored. Within each of the clusters, response allocation showed some sensitivity to the immediacy ratio, but the variation within clusters was less than that

between clusters. The result is that the generalized-matching model (i.e., the fitted regression line) is not representative of performance for any range of immediacy ratios.

The systematic deviation was confirmed through polynomial regression of residuals on log immediacy ratios, which found significant cubic components for Pigeon 161 – replication, $\beta = -1.24$, $p < .05$, and Pigeon 164, $\beta = -1.25$, $p < .05$. No significant components were found in regressions performed on residuals for the other subjects.

It is clear from Figure 4 that even for the cases in which deviations were unsystematic, the regression line accounted for less variance than is commonly obtained in traditional steady-state designs—averaged across all regressions in Figure 4, the variance accounted for was 77.6%. It is likely that some (if not most) of the scatter is related to within-session changes in response allocation, given that each data point in Figure 4 represents a single session. Because sensitivity in the maximal-variation conditions reached an asymptotic value approximately midway through the session (see Figure 3), we conducted an analysis to determine whether some of the variance in response allocation in the second half of the session might be related to factors other than the immediacy ratio. Specifically, we were interested to determine whether there would be evidence for perseveration, that is, that a degree of preference, once established within a session, would tend to continue (cf. Killeen, 2003). To test this hypothesis, we performed a series of hierarchical regressions for all conditions displayed in Figure 4, in which the log response ratio from the fourth quarter of the session (i.e., cycles 55 to 72) was first predicted from the log immediacy ratio. The log response ratio from the third quarter of the session (cycles 37 to 54) was then entered into the regression model to test whether it accounted for a significant increase in variance. In all cases, the incremental variance predicted by the log response ratio in the third quarter was statistically significant, indicating perseveration. Variance accounted for by the log immediacy ratio, and by the log immediacy ratio together with the third-quarter log response ratio, respectively, were as follows: Pigeon 161, 75.0% and 94.2% (third quarter log response ratio $\beta = 0.83$), Pigeon 161 – replication, 82.6% and 96.7% ($\beta = 0.85$), Pigeon 162, 89.4% and 91.7% ($\beta = 0.42$), Pigeon 162 – replication, 83.0% and 96.1% ($\beta = 0.83$), Pigeon 163, 52.4% and 70.2% ($\beta = 0.68$), and Pigeon 164, 86.8% and

97.7% ($\beta = 0.88$). This demonstrates that subjects were perseverating, that is, that a pattern of response allocation, once established in a session, tended to persist.

DISCUSSION

We examined acquisition of preference in concurrent-chains using a procedure in which the position of the shorter terminal-link was changed unpredictably across sessions according to a PRBS (Hunter & Davison, 1985). Previous research had demonstrated that, after exposure to the PRBS, pigeons' response allocation was sensitive to the immediacy ratio in the current session with virtually no effect of preceding sessions (Grace et al., 2003). The key question posed by the present study was whether the degree of variation in terminal-link delays would affect sensitivity to the current-session immediacy ratio. In the minimal-variation condition, one terminal-link schedule was always FI 8 s and the other terminal link changed between FI 4 s and FI 16 s across sessions (i.e., the same arrangement used by Grace et al., 2003), whereas in the maximal-variation condition a unique pair of schedule values was used in each session. In both conditions, the shorter delay was associated equally often with the left and right terminal link, and the expected log immediacy ratios were $\log(1/2)$ or $\log(2)$, depending on the location of the shorter delay. The order of conditions was counterbalanced across subjects.

By the third PRBS presentation, response allocation was sensitive to the current session (i.e., Lag 0) immediacy ratio, and insensitive to prior immediacy ratios, in both the minimal- and maximal-variation conditions. Furthermore, changes in sensitivity within session were similar for the two conditions once subjects were experienced with the procedure, indicating similarly rapid adaptation to the immediacy ratio. However, there was no systematic difference in sensitivity between the conditions across subjects. This result is contrary to the proposition that initial-link response allocation in concurrent chains is determined by the learned or conditioned value of the terminal-link stimuli. According to this view, acquisition should have been faster (and hence Lag 0 sensitivity greater) in the minimal-variation condition. Because one ter-

minimal link FI schedule was constant in that condition, and one of only two FI schedules was used for the other, there should have been some savings in acquisition because the values of the terminal-link stimuli presumably did not have to change to the same extent as in the maximal-variation condition (presumably, the value of the FI 8-s stimulus did not have to change at all). The finding of roughly comparable sensitivity (and hence acquisition) in the maximal-variation condition is thus counterintuitive, and raises the question of how the subjects were able to respond similarly in what should have been a far more difficult discrimination. We also obtained evidence of perseveration in response allocation; when a relatively strong preference for a terminal link had been established by the end of the third quarter of the session, preference for that terminal link tended to remain strong in the fourth quarter. The perseveration was independent of any effect of the immediacy ratio.

Grace et al. (2003) reported that results for some subjects in some conditions showed systematic deviations from generalized matching, and similar deviations were found here. Figure 4 revealed two distinct patterns of response allocation in the maximal-variation condition. Either data points from single sessions were scattered unsystematically around the regression line, or fell into two distinct clusters with relatively greater variation between rather than within clusters. The former can be adequately described by a generalized-matching model (e.g., Grace, 1994), but the latter strongly suggests a categorical discrimination. The clustering implies that delays effectively fell into two classes (short and long), with generalization within classes and discrimination between classes. One possibility is that subjects' behavior could be described as following a simple rule or heuristic, for example, "choose the shorter delay." However, this notion is vague and it is unclear whether it also could account for the remaining performances that were consistent with generalized matching.

What is needed is a model able to explain both patterns of response allocation in Figure 4—generalized matching and categorical discrimination—as well as the lack of a systematic difference in sensitivity between the minimal- and maximal-variation conditions. The following section describes a simple decision model that can accomplish these goals.

A Decision Model

The basic assumption of the model is that on any given trial, initial-link response allocation is determined by the relative response strengths of the associated operants (i.e., the relative propensity to respond to the left and right initial-link stimuli):

$$\frac{B_{Ln}}{B_{Rn}} = \frac{RS_{Ln}}{RS_{Rn}}, \quad (2)$$

where B is initial-link response rate, RS is response strength, and the subscripts L , R , and n indicate the left and right alternatives and cycle number, respectively.

Response strength for a given alternative changes after reinforcement has been earned in a terminal link. Only response strength for the alternative associated with the terminal link in which reinforcement has just been earned is updated. The subject is assumed effectively to make a "decision" as to whether the preceding delay was short or long, relative to the history of delays. An important assumption of the model is that the delays that comprise reinforcement history are pooled across both alternatives. This contrasts with theoretical accounts of timing in which subjects are assumed to maintain separate "memories" for the delays associated with each alternative (e.g., Gallistel & Gibbon, 2000; Gibbon, Church, Fairhurst, & Kacelnik, 1988). Changes in response strength are made according to a linear-operator rule. If the delay is judged as short (or long), response strength is incremented (or decremented) by a constant fraction of the difference between the current and maximum (or minimum) response strength. This leads to the following pair of equations that governs increases in response strength for "short" delays, and decreases for "long" delays:

$$RS_{n+1} = RS_n + (\text{Max}_{RS} - RS_n) * \Delta \quad (3a)$$

(for "short" delays);

$$RS_{n+1} = RS_n - (RS_n - \text{Min}_{RS}) * \Delta \quad (3b)$$

(for "long" delays).

In Equations 3a and 3b, RS is response strength, n and $n + 1$ indicate trial numbers, Max_{RS} and Min_{RS} are the maximum and minimum asymptotic response strengths, re-

spectively, and Δ is a learning rate parameter. Note that although a single pair of equations defines the model, separate response strengths for the left and right alternatives are maintained (i.e., either Equation 3a or 3b is used to update response strength for the just-completed alternative). The ratio Max_{RS} / Min_{RS} determines the strongest degree of asymptotic preference that can be obtained. For example, if the terminal links are widely spaced FI schedules (e.g., FI 5 s FI 30 s), such that delays associated with one terminal link always are classified as short whereas delays from the other always are classified as long, then the response strengths will reach asymptote at a rate determined by Δ , and the predicted response ratio will be Max_{RS} / Min_{RS} .

Reinforcement history is represented as a normal distribution of log delays with a mean equal to the log geometric mean of the experienced delays. The variance of the distribution is assumed to be constant for all delays. Although these assumptions are consistent with a logarithmic “representation” of time, note that this distribution is equivalent (after log transformation) to that posited by scalar expectancy theory, which assumes a linear representation of time but with scalar variance (Gibbon, 1981). The distribution of log delays is used both for simplicity and because research in animal psychophysics has shown that the bisection point for temporal discrimination typically occurs at the geometric mean (Church & Deluty, 1977; Stubbs, 1968). Moreover, it should be emphasized that the model does not assume the subject is generating an internal representation of reinforcement history as a distribution of delays in memory; rather, the distribution and associated decision process are proposed as a model for control over response allocation by reinforcement history in the rapid-acquisition procedure (cf. Wixted & Gaitan, 2002).

The mean of the distribution—the temporal bisection point—serves effectively as a “criterion” against which delays are judged. Delays that are less than the criterion are likely (i.e., $p > .5$) to be classified as short, and vice versa for delays that are greater than the criterion. The standard deviation is a parameter in the model, and influences the likelihood of making a correct decision. According to the

model, the probability of making a short decision, p_{short} , is computed as the area under the distribution to the right of the log of the preceding delay (note that $p_{long} = 1 - p_{short}$). The process is illustrated in Figure 5. This figure shows the reinforcement history distribution, with mean (i.e., log criterion; top arrow) equal to the average of the log delays from the maximal-variation condition (0.92, which corresponds to a delay of 8.37 s). Assume that the preceding terminal-link delay was 7.5 s ($\log 7.5 = .875$; bottom arrow in Figure 5). Because this delay is less than the criterion, it is likely that it will be classified as short. According to the model, the probability of a short classification is given by the shaded area in Figure 5, which is 73.7%. The shaded area is equal to the inverse of the cumulative normal distribution.

To generate predictions using the model, p_{short} must be computed for the delay experienced on trial n . Then p_{short} and $1 - p_{short}$ can be used to combine Equations 3a and 3b into a single expression for the expected response strength for trial $n + 1$:

$$RS_{n+1} = RS_n + p_{short} * (Max_{RS} - RS_n) * \Delta - (1 - p_{short}) * (RS_n - Min_{RS}) * \Delta. \tag{4}$$

Equation 4 gives the expected (i.e., average) response strength for an alternative on trial $n + 1$, given that the delay just experienced was classified as short with $p = p_{short}$. In principle, a “real time” version of the model could be used in which the actual sequence of classifications was simulated. Such a model could account for nonmonotonic changes in preference during a session. For present purposes, however, we adopted the simpler approach of computing predictions based on expected response strengths. This yielded a monotonic acquisition curve for each session.

Equation 4 was used to generate predictions for a set of terminal-link schedule pairs that covered the range of immediacy ratios used in the experiment, FI 3.6 s FI 14.4 s [4:1] to FI 14.4 s FI 3.6 s [1:4]. For these simulations, the mean of the reinforcement history delay distribution was again equal to 0.92 (i.e., the log geometric mean of the average delays from the maximal-variation condition in training), the learning rate parameters for both alternatives were set equal to 0.2, and the maxi-

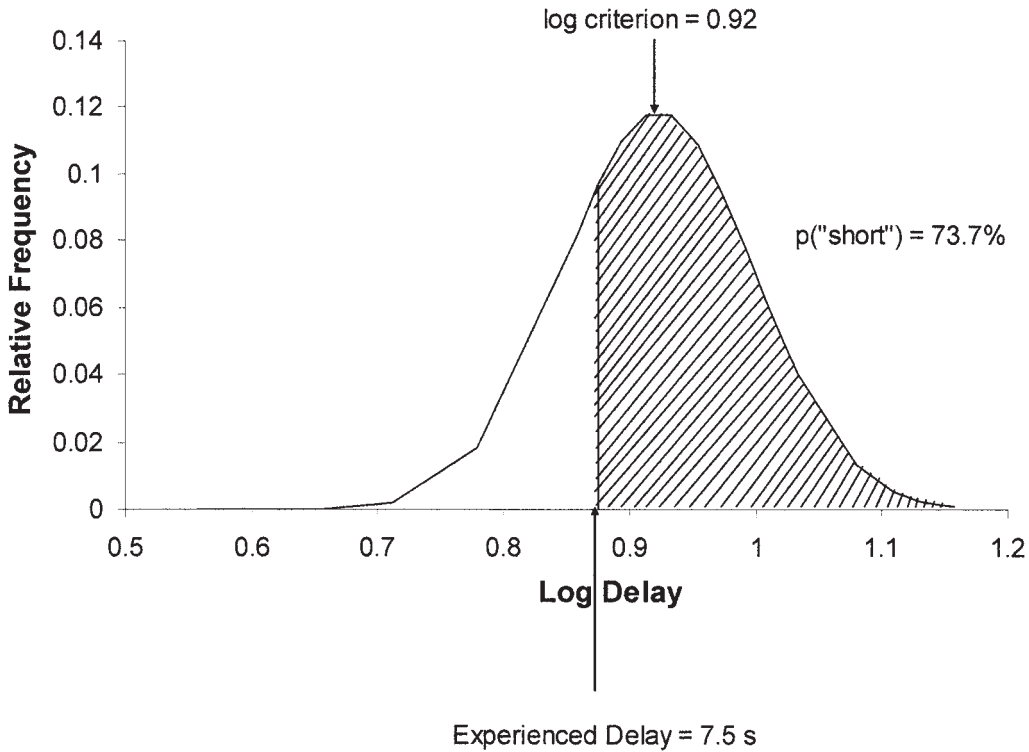


Fig. 5. Illustration of how the decision model computes the probability that a terminal-link delay will be classified as “short.” Shown is the lognormal distribution representing the history of reinforcement delays combined for both alternatives. The probability of a short classification for a 7.5-s delay is the area under the distribution to the right of log 7.5.

imum and minimum response strengths were 0.9 and 0.1, ensuring that the maximum response ratio predictable by the model was 9:1. Response strengths for both alternatives were reset to 0.5 at the start of each session. This corresponds to the assumption that the immediacy ratios from prior sessions had no effect; response allocation was controlled by the terminal-link schedules in the current session only, consistent with the finding that, typically, only Lag 0 sensitivity coefficients were significantly greater than zero by the third PRBS presentation in both conditions (see Figure 1; Grace et al., 2003). Separate versions of Equation 4 then were used to compute changes in response strength across 36 presentations of the left and right terminal links. The predicted overall session response allocation then was based on the average response strength for each alternative across the 36 presentations. Two sets of predictions were calculated, which differed according to whether the standard deviation of the reinforcement

history distribution was relatively large (0.3) or small (0.075).

Results are shown in Figure 6. When the standard deviation was relatively large (left panel), the predicted log initial-link response ratio was an approximately linear function of the log immediacy ratio. The regression line in the left panel accounted for more than 99.5% of the variance. By contrast, when the standard deviation was relatively small (right panel), the predicted data follow a sigmoidal pattern and deviate systematically from the regression line. Thus Figure 6 demonstrates that the decision model can generate predictions that conform to either generalized matching or categorical discrimination. When the standard deviation is low, discrimination of short versus long delays is relatively accurate and response allocation shows evidence of categorical discrimination. However, when the standard deviation is large the discrimination is relatively inaccurate, and response allocation conforms to generalized matching.




Fig. 6. Illustration of how the decision model can predict molar response allocation that is approximately linear (left panel) or sigmoidal (right panel) as a function of the log immediacy ratio, depending on the standard deviation of the lognormal distribution representing the history of reinforcement delays.

To evaluate whether the decision model was able to account for the present results, the model was fitted to the individual data from the third PRBS presentations in both conditions. For each session, log initial-link response ratios were computed for every six cycles (i.e., twelve log response ratios per session), giving a total of 372 (12×31) data points for each subject. The log criterion values were set equal to the log geometric mean of the programmed delays in each condition (0.92 and 0.90 for the maximal- and minimal-variation conditions, respectively). The maximum and minimum response strengths were again assumed to be 0.9 and 0.1, and the starting response strength for both alternatives in each session was set equal to 0.5. Next, the values of the standard deviation and the learning rate parameters for the left and right alternatives (Δ_L and Δ_R) that maximized the variance accounted for by the model were determined by a nonlinear optimization procedure. Then the adequacy of the fits was checked by conducting a regression of the obtained on the predicted values. For 2 subjects in the minimal-variation condition (Pigeons 161 and 163) the model was systematically predicting values that were not extreme enough (as indicated by a regression slope greater than 1.0). The model was reapplied to the data for these pigeons except that the maximum and minimum response strength parameters for the right alternative (i.e., the changing terminal link) were esti-

mated, allowing predicted preference to become more extreme. In this way, adequate fits in terms of acceptable regression parameters—slopes and intercepts not significantly different from 1.0 and 0.0, respectively—were obtained for all conditions. Using these parameter estimates, the predictions of the model for the full session data and variance accounted for were calculated.

The parameter values for the model, variance accounted for in the session-twelfth and full-session data, and regression statistics (obtained on predicted) are shown in Table 1 for all subjects. The decision model accounted for an average of 78% and 69% of the variance in the session-twelfth data in the minimal- and maximal-variation conditions (including replications), respectively. Overall, the fits of the model to the individual data were moderately good ($VAF > 0.75$), with the exception of Pigeon 163 (both conditions), and Pigeon 161 (maximal-variation condition).

Table 1 shows that for all subjects, the estimated standard deviation was greater in the maximal-variation (and replication) compared to minimal-variation conditions. Because the accuracy of correctly classifying a delay as short or long varies inversely with the standard deviation, this might seem to suggest that the model predicts sensitivity to the log immediacy ratio will be greater in the minimal-variation condition—with the other parameters held constant, decreases in the

Table 1

Results of analysis in which the decision model was fitted to data from the third PRBS presentation in the minimal- and maximal-variation conditions. Listed are estimated parameter values (σ , Δ_L , Δ_R), the maximum and minimum response strengths for each alternative (MaxL, MinL, MaxR, MinR), the variance accounted for by the model in the session-twelfth and full-session data (VAF12th, VAFses), the slope and intercept from regressions performed on the obtained and predicted session-twelfth data, and the variance accounted for in the full-session data by the generalized matching law (VAF GML).

Pigeon	Condition	σ	Δ_L	Δ_R	MaxL	MinL	MaxR	MinR	VAF12th	Slope	Intercept	VAFses	VAF GML
161	Minimal	0.0001	0.04	0.62	0.90	0.10	0.80	0.03	0.8813	1.01	-0.01	0.9613	0.9600
	Maximal	0.1982	0.26	0.36	0.90	0.10	0.90	0.10	0.6162	1.00	0.01	0.7065	0.6986
162	Max - rep	0.0320	0.08	0.44	0.90	0.10	0.90	0.10	0.7822	1.03	0.01	0.8691	0.8037
	Minimal	0.0127	0.99	0.88	0.90	0.10	0.90	0.10	0.7723	1.05	-0.03	0.9110	0.9162
163	Maximal	0.1169	0.92	0.28	0.90	0.10	0.90	0.10	0.7328	0.99	-0.04	0.8519	0.8451
	Max - rep	0.1370	0.99	0.21	0.90	0.10	0.90	0.10	0.7226	1.00	-0.03	0.8175	0.8147
164	Minimal	0.0101	0.08	0.41	0.90	0.10	1.36	0.07	0.6564	1.00	0.01	0.7439	0.7611
	Maximal	0.1149	0.06	0.31	0.90	0.10	0.90	0.10	0.4547	0.99	0.00	0.6513	0.6207
164	Minimal	0.0001	0.03	0.37	0.90	0.10	0.90	0.10	0.7940	0.97	-0.02	0.9563	0.9583
	Maximal	0.0016	0.35	0.22	0.90	0.10	0.90	0.10	0.8550	0.98	0.01	0.9005	0.8718

standard deviation produce increases in sensitivity. However, according to the model, sensitivity also depends on the learning rate parameters (Δ_L , Δ_R). Predicted Lag 0 sensitivities were calculated by regressing the model's full-session predictions on the log immediacy ratios, and are listed in Table 2 along with the corresponding obtained Lag 0 sensitivities. The predicted sensitivities were highly correlated ($r = .88$) with the obtained values, and reproduced the pattern of individual differences: sensitivity was greater in the minimal-variation condition for Pigeons 161 and 163, greater in the maximal-variation condition for Pigeon 164, and approximately the same across conditions for Pigeon 162. Thus the model predicted the individual differences in Lag 0 sensitivity, even though for all subjects the standard deviation was smaller (and hence decision accuracy greater) in the minimal-variation condition.

Figures 7 through 9 illustrate the fits of the model in the minimal- and maximal-variation conditions, and maximal-variation replication, respectively. In each figure, the left panels show scatterplots of the obtained versus predicted values for the session-twelfth data. In all cases, the slopes and intercepts of the regression lines are close to 1.0 and 0.0, respectively, indicating that there was little systematic deviation in the model fits. The predictions of the model (unfilled symbols), and the obtained results (filled symbols), are shown for the full-session data in the right panels. The regression equation and variance accounted for associated with a generalized-matching model are provided for sake of comparison. The average variance accounted for by the model in the full-session data was 89% and 80% for the minimal- and maximal-variation conditions (including replications), respectively, and was greater than that accounted for by the generalized-matching model in 7 of 10 cases (the regression lines in the right panels; see also Table 1).

Although the model is able to capture the overall trends, there is a fair amount of scatter, especially in the session-twelfth data (left panels in Figures 7 through 9). Figure 10 provides some insight into why the model was only able to account for about 75% of the variance in relative initial-link responding within- and across-sessions. Shown are log initial-link response ratios by session twelfth

Table 2

Obtained and predicted Lag 0 sensitivity coefficients for all subjects for the third PRBS presentation of each condition (including replications). See text for more explanation.

Pigeon	Obtained Lag 0 sensitivity			Predicted Lag 0 sensitivity		
	Minimum	Maximum	Maximum replication	Minimum	Maximum	Maximum replication
161	1.73	0.87	1.00	1.57	0.89	1.12
162	1.43	1.30	1.39	1.35	1.37	1.21
163	1.67	1.15		1.50	0.84	
164	1.04	1.42		1.10	1.36	

for Pigeons 161 and 163 in the minimal-variation condition. These subjects were chosen because the model accounted for the highest and the lowest percentage variance in that condition, respectively. In both cases, log response ratios change systematically during the session depending on whether the right terminal-link schedule was FI 4 s or FI 16 s, but it is clear that there are many cases of nonmonotonic changes within sessions. As applied to the data, the decision model is only able to predict monotonic changes (i.e., smooth acquisition curves) because the computations are based on expected response strengths. The model does not predict the actual sequences of “decisions” made by the subject (correct and incorrect), and how response strength might change nonmonotonically throughout the session as a result.

Overall, the decision model satisfies the major goal specified at the outset: to account for response allocation typical of both generalized matching and categorical discrimination. The model accomplishes this through changes in the standard deviation that are intuitively plausible. Given continued training with the PRBS procedure, reinforcement history exerts stronger control over response allocation in that differences between delays in current and prior sessions are judged more accurately. This is reflected by decreases in the standard deviation across successive PRBS presentations within conditions. The model also is able to account for individual differences in sensitivity between the minimal- and maximal-variation conditions. The predicted Lag 0 sensitivity depends not only on the standard deviation, but also on the learning rate parameters (Δ_L , Δ_R). Thus, although standard deviations were always smaller in the minimal-variation condition, suggesting more accurate discrimination, the model was

able to reproduce the obtained pattern of individual differences in sensitivity. Finally, the model also can account for the observed perseveration in response allocation within sessions. To the extent that the sequence of decisions results in relatively large difference in response strength (and hence strong preference) by the end of the third quarter of the session, then preference in the fourth quarter also is likely to be relatively strong.

One major simplifying assumption made here—that response strengths were reset to intermediate values at the start of each session—clearly would not hold for traditional, steady-state designs in which schedule parameters remain unchanged until response allocation is stable. The model is incomplete because the values of the learning rate parameters (as well as the degree of carryover in response-strength changes within session) must depend on the frequency of environmental change (cf. Mazur, 1997). However, it may be that an elaboration of the basic model presented here can account for concurrent-chains choice in both rapid acquisition and steady-state designs (Grace, 2002a). In the remaining sections we will briefly consider some implications of the model for various choice phenomena.

The Provenance of Matching

One of the oldest and most important controversies in research on choice concerns the origins of matching. A large number of different theories have been proposed (see Williams, 1988, 1994, for review), but we will focus on the question of whether matching itself is fundamental or can be explained in terms of a more local process. The model advanced here allows for response allocation to vary continuously and hence show constant sensitivity to the immediacy ratio, even when

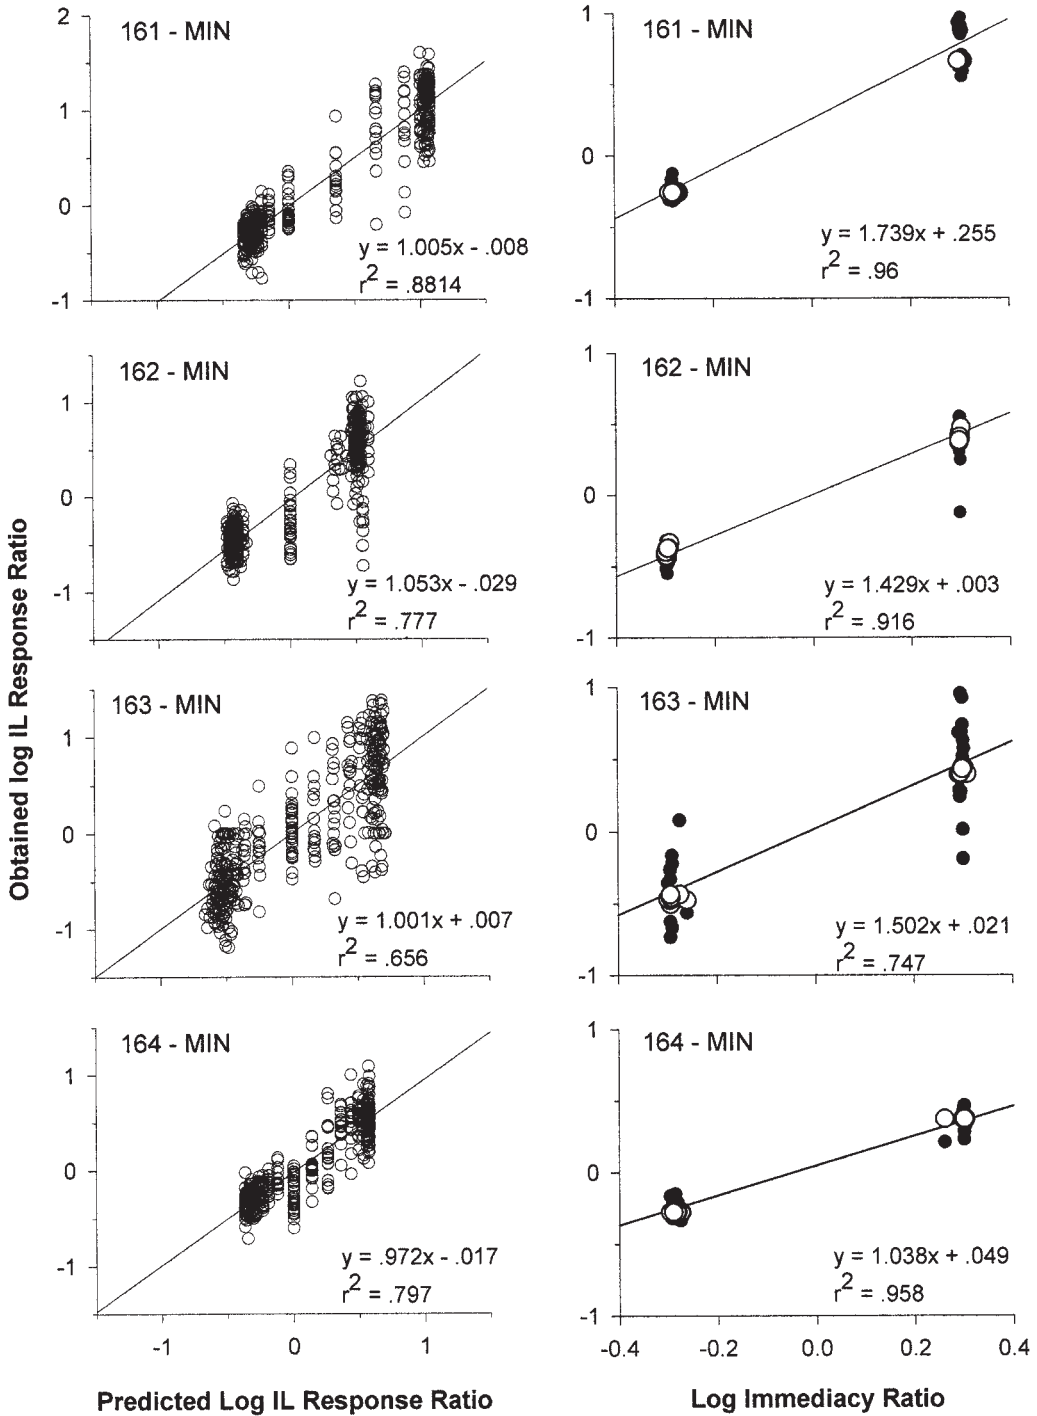


Fig. 7. Fits of the decision model to data from the third PRBS presentation in the minimal-variation condition. The left panels show obtained versus predicted log initial-link response allocations for each session twelfth. The regression lines and associated parameters also are given. The right panels show obtained overall log initial-link response ratios as a function of log terminal-link immediacy ratios (filled symbols) and the log response ratios predicted by the decision model (unfilled symbols). For sake of comparison, the regression lines and associated parameters indicate predictions of a generalized-matching model. See text for more details.

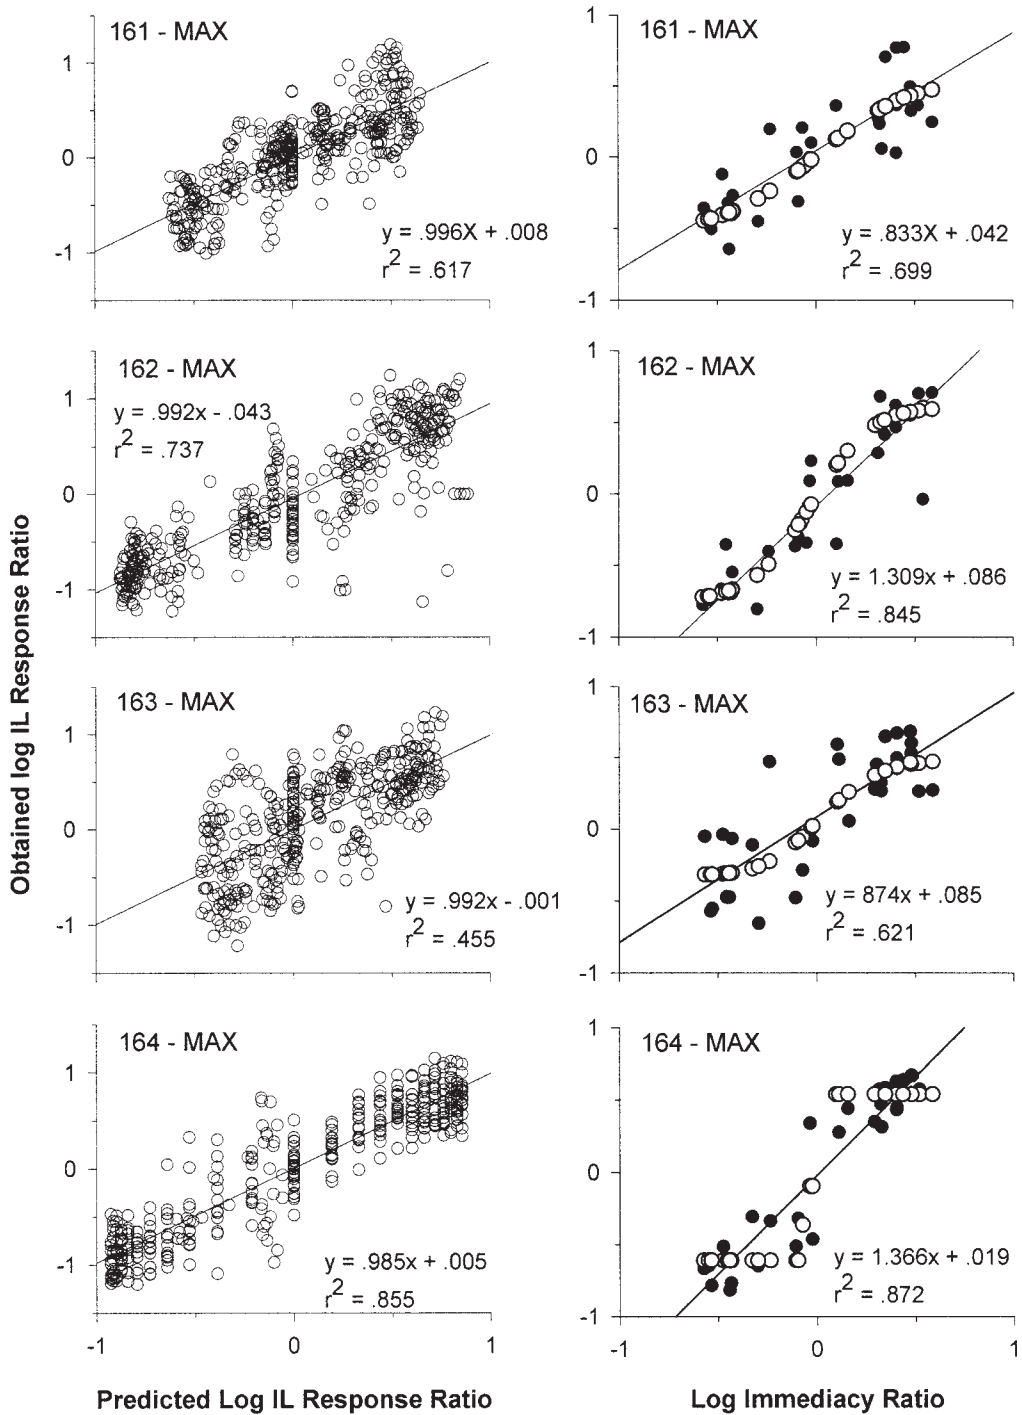


Fig. 8. Fits of the decision model to data from the third PRBS presentation in the maximal-variation condition. The left panels show obtained versus predicted log initial-link response allocations for each session twelfth. The regression lines and associated parameters also are given. The right panels show obtained overall log initial-link response ratios as a function of log terminal-link immediacy ratios (filled symbols) and the log response ratios predicted by the decision model (unfilled symbols). For sake of comparison, the regression lines and associated parameters indicate predictions of a generalized-matching model. See text for more details.

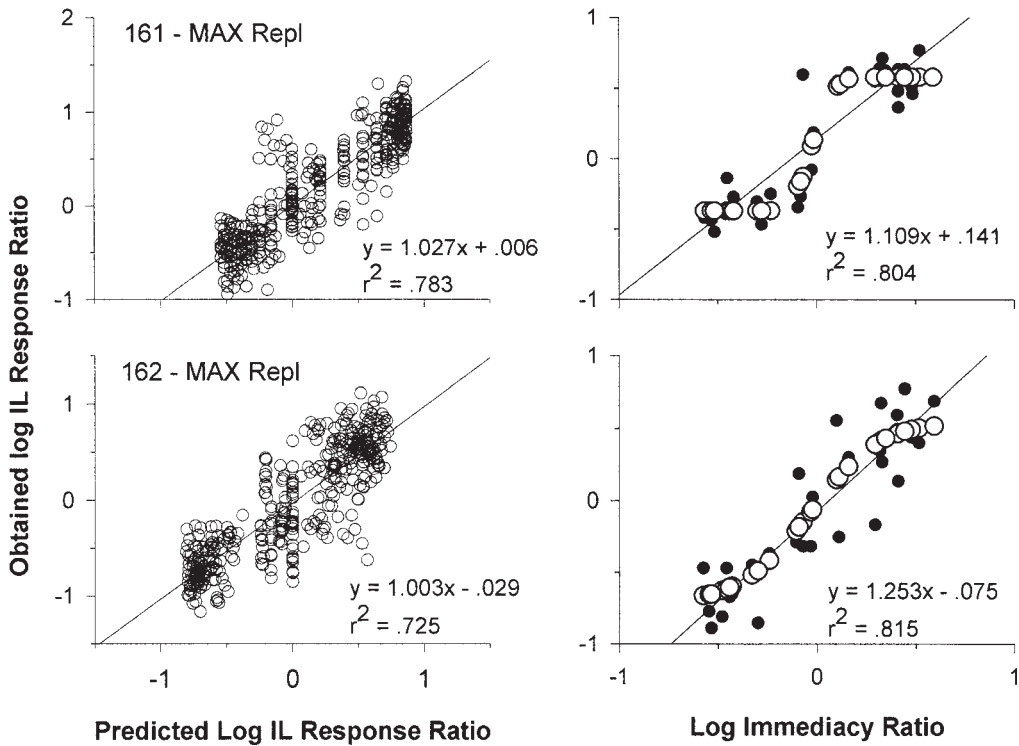


Fig. 9. Fits of the decision model to data from the replication of the maximal-variation condition (for Pigeons 161 and 162 only). The left panels show obtained vs. predicted log initial-link response allocations for each session twelfth. The regression lines and associated parameters also are given. The right panels show obtained overall log initial-link response ratios as a function of log terminal-link immediacy ratios (filled symbols) and the log response ratios predicted by the decision model (unfilled symbols). For sake of comparison, the regression lines and associated parameters indicate predictions of a generalized-matching model. See text for more details.

the decisions are categorical—changes in response strength occur when a delay has been judged short or long, but by *how much* it is short or long is irrelevant. Thus the model suggests that generalized matching in concurrent chains—that is, response allocation that is a power function of the immediacy ratio—can be understood in terms of a categorical decision process.

However, it is important to note that the model only produces an approximation to generalized matching. As the left panel of Figure 6 shows, even when the model predicts that response allocation is approximately a linear function of the log immediacy ratio, there are systematic deviations from generalized matching in the direction predicted by a sigmoidal function. Because these deviations are relatively small, they may be hard to detect empirically.

Can a similar model account for matching to relative reinforcement rate in concurrent schedules? Recently, Davison, Baum and colleagues (e.g., Davison & Baum, 2000; Landon, Davison, & Elliffe, 2003) have reported a series of experiments in which reinforcement parameters (e.g., rate or magnitude) were varied within sessions in a multiple-component concurrent schedules procedure. When data were aggregated across many sessions (typically 50 per condition) and analyzed in terms of reinforcer sequences, response allocation showed systematic changes; specifically, after each reinforcer preference shifted towards the most recently reinforced alternative, which led Davison and Baum (2000) to conclude that “every reinforcer counts.” If every reinforcer counts, implying that each one has an equivalent impact on choice behavior, then matching would be the expected result at the molar

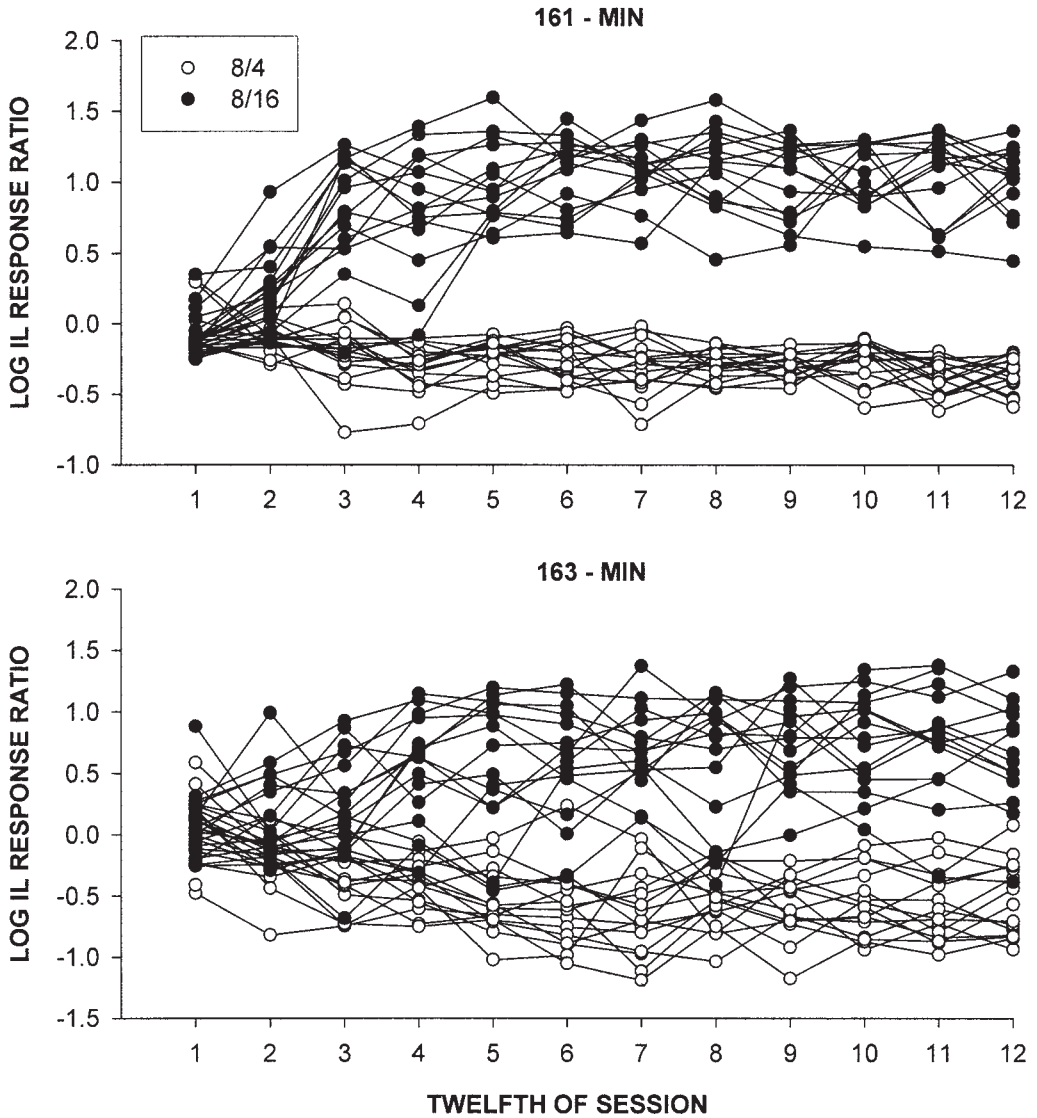


Fig. 10. Log initial-link response ratios as a function of session twelfth for the third PRBS presentation in the minimal-variation condition for Pigeons 161 (upper panel) and 163 (bottom panel). Data points from the same session are connected by lines. Filled symbols indicate sessions in which the terminal links were FI 8 s FI 16 s; unfilled symbols indicate sessions in which the terminal links were FI 8 s FI 4 s. See text for explanation.

level. A process in which responding is strengthened incrementally by reinforcement is consistent with the decision model offered here.

Molar Choice in Concurrent Chains

Ultimately, any successful comprehensive model for concurrent chains must be able to explain performance in traditional, steady-state designs. The model described here

incorporates only the effects of reinforcement delay, and would need to be elaborated further to account for factors such as reinforcement magnitude (Grace, 1995), probability (Spetch & Dunn, 1987), signaled versus unsignaled terminal links (Williams & Fantino, 1978), and terminal-link response contingencies (Moore & Fantino, 1975; Nevin, Grace, Holland, & McLean, 2001). In its present form, the model also is unable to account for

the effect of absolute initial- and terminal-link duration on preference (e.g., Fantino, 1969; Grace & Bragason, 2004; MacEwen, 1972). However, there are several results that are easily explained by the model as it stands.

The model predicts preference for variability—that is, for VI over FI schedules that have the same average delay to reinforcement (Herrnstein, 1964). Because the distribution of intervals for the VI schedule is skewed, the geometric mean of the pooled delay distribution will be less than the FI value. Thus the probability of the FI delay being classified as short is less than 0.5. Conversely, a delay sampled from the VI schedule is likely ($p > .5$) to be shorter than the average value, and thus is more likely to be classified as short than the FI delay.

The model also predicts that preference should be more sensitive to the ratio of average delays with FI than with VI terminal links (Killeen, 1970). The reason is that the delay distributions for the VI schedules overlap, whereas the distributions for the FI schedules do not. Thus the probability of a short delay from the longer VI schedule is greater than the probability of a short delay from the longer FI schedule, and conversely, the probability of a long delay from the shorter VI schedule is greater than the probability of a long delay from the shorter FI schedule. This leads to weaker preference for pairs of VI schedules compared to FI schedules with the same immediacy ratio.

Temporal Discounting as a Derived Phenomenon

Temporal discounting is the process whereby the value of a reward decreases according to its delay and is a topic of increasing research interest with both humans and nonhumans (e.g., Grace, 1999; Kirby, Petry, & Bickel, 1997; see Green & Myerson, 2004, for review). One of the most interesting and provocative results has been that data from experiments with both humans and nonhumans are better described by hyperbolic rather than exponential functions (e.g., Myerson & Green, 1995).

One of the implications of our model is that the temporal discounting may be a derived phenomenon, at least in nonhumans. Using the concurrent-chains procedure, a discounting function may be derived by assuming that initial-link response allocation matches the relative value of the terminal-link stimuli, with

value determined according to a discounting function. According to this view, the slope of the generalized-matching relation between log relative initial-link response and log terminal-link immediacy ratios provides an index of discounting (Grace, 1999). Moreover, the determiners of “value” in other procedures involving choice between delayed reinforcers, such as Mazur’s (1984) adjusting-delay procedure, appear to be the same as in concurrent chains (Grace, 1996). If the generalized-matching relation, which is a continuous function of the immediacy ratio, is obtained through an accumulation of ordinal “decisions,” then subjects’ behavior is not determined by the delay ratio directly. The implication is that temporal discounting rate, as indexed by the slope of the generalized-matching relation, may be a derived phenomenon. Ultimately, it may be possible to predict the shape of the discounting function from the basic principles of a model similar to the one proposed here.

Conditioned Reinforcement Value

Models for concurrent chains typically have assumed that response allocation during the initial links is determined by the relative value of the terminal-link stimuli as conditioned reinforcers, although how value should be defined has been a controversial issue (e.g., Fantino, Preston, & Dunn, 1993; Grace, 1994, 2002b; Mazur, 2001; but cf. Grace & Nevin, 1999). As noted in the Introduction, according to this view acquisition should have been faster in the minimal-variation condition, and so the lack of a systematic difference in Lag 0 sensitivity is evidence against learned or conditioned value as an explanation for the present results. However, it is important to realize that “response strength” in the decision model (i.e., propensity to respond to an alternative) could be relabeled “conditioned value” without altering the model’s predictions or its ability to describe the data. How the model’s construct should best be interpreted is a theoretical question, not an empirical one. We believe that the assumption that the construct be reset (to 0.5) at the start of each session, which was necessary to model the data, is difficult to reconcile with traditional views of conditioned reinforcement. Thus we favor an interpretation of the construct as response strength or “propensity to respond,” although

it is possible that alternative perspectives on conditioned reinforcement (e.g., as “situation transition”; see Baum, 1974b) might be consistent with the present data and model.

CONCLUSION

Overall, the decision model shows promise as an account of concurrent-chains performance. It is able to describe changes in response allocation within-sessions under conditions in which terminal links change unpredictably across sessions. A plausible extension of the model is consistent with many results from the steady-state literature (e.g., preference for variability, overmatching with FI schedules). Ultimately, a model based on the decision process outlined here may be able to provide a complete account of concurrent-chains performance.

REFERENCES

- Baum, W. M. (1974a). 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. (1974b). Chained concurrent schedules: Reinforcement as situation transition. *Journal of the Experimental Analysis of Behavior*, 22, 91–101.
- Church, R. M., & Deluty, M. Z. (1977). Bisection of temporal intervals. *Journal of Experimental Psychology: Animal Behavior Processes*, 3, 216–228.
- Davison, M., & Baum, W. M. (2000). Choice in a variable environment: Every reinforcer counts. *Journal of the Experimental Analysis of Behavior*, 74, 1–24.
- Davison, M., & Jenkins, P. E. (1985). Stimulus discriminability, contingency discriminability, and schedule performance. *Animal Learning & Behavior*, 13, 77–84.
- Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum.
- Fantino, E. (1969). Choice and rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, 12, 723–730.
- Fantino, E., Preston, R. A., & Dunn, R. (1993). Delay reduction: Current status. *Journal of the Experimental Analysis of Behavior*, 60, 159–169.
- Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 5, 529–530.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107, 289–344.
- Gibbon, J. (1981). On the form and location of the psychometric bisection function for time. *Journal of Mathematical Psychology*, 24, 58–87.
- Gibbon, J., Church, R. M., Fairhurst, S., & Kacelnik, A. (1988). Scalar expectancy theory and choice between delayed rewards. *Psychological Review*, 95, 102–114.
- Grace, R. C. (1994). A contextual model of concurrent-chains choice. *Journal of the Experimental Analysis of Behavior*, 61, 113–129.
- Grace, R. C. (1995). Independence of reinforcement delay and magnitude in concurrent chains. *Journal of the Experimental Analysis of Behavior*, 63, 255–276.
- Grace, R. C. (1996). Choice between fixed and variable delays to reinforcement in the adjusting-delay procedure and concurrent chains. *Journal of Experimental Psychology: Animal Behavior Processes*, 22, 362–383.
- Grace, R. C. (1999). The matching law and amount-dependent exponential discounting as accounts of self-control choice. *Journal of the Experimental Analysis of Behavior*, 71, 27–44.
- Grace, R. C. (2002a). Acquisition of preference in concurrent chains: Comparing linear-operator and memory-representational models. *Journal of Experimental Psychology: Animal Behavior Processes*, 28, 257–276.
- Grace, R. C. (2002b). The value hypothesis and acquisition of preference in concurrent chains. *Animal Learning & Behavior*, 30, 21–33.
- Grace, R. C., & Bragason, O. (2004). Does the terminal-link effect depend on duration or reinforcement rate? *Behavioural Processes*, 67, 67–79.
- Grace, R. C., Bragason, O., & McLean, A. P. (2003). Rapid acquisition of preference in concurrent chains. *Journal of the Experimental Analysis of Behavior*, 80, 235–252.
- Grace, R. C., & Nevin, J. A. (1999). Timing and choice in concurrent chains. *Behavioural Processes*, 45, 115–127.
- Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, 130, 769–792.
- Herrnstein, R. J. (1964). Aperiodicity as a factor in choice. *Journal of the Experimental Analysis of Behavior*, 7, 179–182.
- Hunter, I., & Davison, M. (1985). Determination of a behavioral transfer function: White-noise analysis of session-to-session response-ratio dynamics on concurrent VI VI schedules. *Journal of the Experimental Analysis of Behavior*, 43, 43–59.
- Killeen, P. (1970). Preference for fixed-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 14, 127–131.
- Killeen, P. R. (2003). Complex dynamic processes in sign-tracking with an omission contingency (negative automaintenance). *Journal of Experimental Psychology: Animal Behavior Processes*, 29, 49–60.
- Kirby, K. N., Petry, N. M., & Bickel, W. K. (1997). Heroin addicts have higher discount rates for delayed rewards than non-drug using controls. *Journal of Experimental Psychology: General*, 128, 78–87.
- Landon, J., Davison, M., & Elliffe, D. (2003). Choice in a variable environment: Effects of unequal reinforcer distributions. *Journal of the Experimental Analysis of Behavior*, 80, 187–204.
- MacEwen, D. (1972). The effects of terminal-link fixed-interval and variable-interval schedules on responding under concurrent chained schedules. *Journal of the Experimental Analysis of Behavior*, 18, 253–261.
- Mazur, J. E. (1984). Tests for an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 426–436.
- Mazur, J. E. (1997). Effects of rate of reinforcement and rate of change on choice behaviour in transition. *Quarterly Journal of Experimental Psychology*, 50B, 111–128.
- Mazur, J. E. (2001). Hyperbolic value addition and general models of animal choice. *Psychological Review*, 108, 96–112.

- Moore, J., & Fantino, E. (1975). Choice and response contingencies. *Journal of the Experimental Analysis of Behavior*, *23*, 339–347.
- Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior*, *64*, 263–276.
- Nevin, J. A., Grace, R. C., Holland, S., & McLean, A. P. (2001). Variable-ratio versus variable-interval schedules: Response rate, resistance to change and preference. *Journal of the Experimental Analysis of Behavior*, *76*, 43–74.
- Schofield, G., & Davison, M. (1997). Nonstable concurrent choice in pigeons. *Journal of the Experimental Analysis of Behavior*, *68*, 219–232.
- Spetch, M. L., & Dunn, R. (1987). Choice between reliable and unreliable outcomes: Mixed percentage-reinforcement in concurrent chains. *Journal of the Experimental Analysis of Behavior*, *47*, 57–72.
- Stubbs, D. A. (1968). The discrimination of stimulus duration by pigeons. *Journal of the Experimental Analysis of Behavior*, *11*, 223–238.
- Williams, B. A. (1988). Reinforcement, choice, and response strength. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology: Vol. 2. Learning and cognition* (2nd ed., pp. 167–244). New York: Wiley.
- Williams, B. A. (1994). Reinforcement and choice. In N. J. Mackintosh (Ed.), *Animal learning and cognition* (pp. 81–108). San Diego: Academic Press.
- Williams, B. A., & Fantino, E. (1978). Effects on choice of reinforcement delay and conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*, *29*, 77–86.
- Wixted, J. T., & Gaitan, S. C. (2002). Cognitive theories as reinforcement history surrogates: The case of likelihood ratio models of human recognition memory. *Learning & Behavior*, *30*, 289–305.

Received August 5, 2004

Final acceptance September 1, 2005