

*REINFORCER-RATIO VARIATION AND
ITS EFFECTS ON RATE OF ADAPTATION*

JASON LANDON AND MICHAEL DAVISON

UNIVERSITY OF AUCKLAND, NEW ZEALAND

Six pigeons were trained in sessions that consisted of six or seven concurrent-schedule components, each of which could have a different reinforcer ratio arranged in it. The components were unsignaled and occurred in a random order separated by 10-s blackouts. The overall reinforcer rate arranged in each component was 2.22 reinforcers per minute. In Experiment 1, the range of reinforcer ratios in the seven components was varied from a condition in which the ratios were always 1:1, to a condition in which the ratios varied between concurrent variable-interval 27 s extinction (EXT) and concurrent extinction variable-interval 27 s (ratios of 1:EXT, 9:1, 3:1, 1:1, 1:3, 1:9, EXT:1). In Experiment 2, the range of reinforcer ratios was always 27:1 to 1:27, and the presence and absence of the intermediate reinforcer ratios used in Experiment 1 (9:1, 3:1, 1:1, 1:3, 1:9) were investigated. Log response-allocation ratios in components changed rapidly with increasing numbers of reinforcers in components, and Experiment 1 showed that sensitivity to reinforcement was usually higher when the range of reinforcer ratios was greater. When the range of reinforcer ratios was kept constant in Experiment 2, the presence or absence of less extreme reinforcer ratios had no clear effect on sensitivity. At a local level, individual reinforcers had predictable quantitative effects on response ratios: Successive same-alternative reinforcers in a component had rapidly diminishing effects in both experiments. Reinforcers obtained on the opposite alternative to one or more prior reinforcers always had large effects on preference, and these changes were greater when the range of reinforcer ratios was greater. The effects of such reinforcers in changing preference were enhanced, and produced clear preference reversals, when intermediate reinforcer ratios were absent in Experiment 2. Two processes, one local to reinforcers and one with a longer time course, may be necessary to account for these results.

Key words: concurrent schedules, choice, generalized matching, behavior change, reinforcer-ratio variation, key peck, pigeons

Previous research examining choice has been concerned almost exclusively with the analysis of steady-state performance. In these procedures, some independent variable is usually held constant until behavior is stable, that is, until there is no systematic variation in behavior from session to session. Such a criterion typically requires 15 to 30 sessions of training, and data from only the last few sessions are analyzed.

Research on choice using concurrent schedules has shown that the proportion of responses or time allocated to each schedule typically approximates the proportion of re-

inforcers obtained from each schedule. A generalized form of this matching relation, the generalized matching law (Baum, 1974; Staddon, 1968), has been used to describe behavior in a variety of choice situations (see Davison & McCarthy, 1988, for a review). When the logarithmic form of the generalized matching law is used, the log of the ratio of responses emitted (or time spent responding) on the two alternatives is a linear function of the log of the ratio of reinforcers obtained from the two alternatives:

$$\log\left(\frac{B_1}{B_2}\right) = a \log\left(\frac{R_1}{R_2}\right) + \log c, \quad (1)$$

where B_1 and B_2 are the number of responses emitted on and R_1 and R_2 are the number of reinforcers obtained from Alternatives 1 and 2, respectively. The parameter a is sensitivity to reinforcement (Lobb & Davison, 1975), and is a measure of the degree to which response ratios change with variations in the reinforcer ratios. Log c (inherent bias) measures any constant proportional preference for one alternative over the other that is in-

The experiments reported here were completed in partial fulfillment of master's and doctoral degrees at Auckland University by the first author. We thank William Baum, Douglas Elliffe, and Max Jones for continuing input to this research, Mick Sibley for looking after the subjects, and members of the EABRU for helping to conduct the experiments. The data used in the analyses are too extensive to include in an appendix, and may be obtained from the authors on receipt of a 100-Mb zip disk.

Reprints may be obtained from either author, Department of Psychology, University of Auckland, Private Bag 92019, Auckland, New Zealand (E-mail: j.landon@auckland.ac.nz or m.davison@auckland.ac.nz).

dependent of R_1 and R_2 . In concurrent schedules, the ratio of the times spent responding to each of the two alternatives can also be used as the dependent variable in Equation 1 (e.g., Davison & McCarthy, 1988).

Following the quantification of steady-state behavior in this way, researchers began to address aspects of environmental variability that might affect the acquisition and distribution of behavior in concurrent schedules. Davison and Hunter (1979) examined whether preference was affected by previous experimental conditions, and if so, how long this effect lasted. They investigated sessional performance in concurrent variable-interval (VI) VI schedules in which the reinforcer ratio arranged by the two schedules changed every six sessions. Generalized matching analyses showed that the reinforcer ratio from the previous experimental condition had a large effect on performance in the first session after a transition. Some measurable effect of the previous conditions remained after three sessions of training, but this effect disappeared in fewer than six sessions. In this way, Davison and Hunter demonstrated that orderly and quantifiable changes occurred during the transition between stable performances.

Hunter and Davison (1985) also investigated the effects of previous sessions' reinforcers on current-session behavior. Pigeons were exposed to two concurrent schedules, a concurrent VI 60-s VI 240-s schedule and a concurrent VI 240-s VI 60-s schedule. The concurrent schedule that was in effect in a session was unsignaled, and could change each session according to a single 31-step pseudorandom binary sequence. The results showed that, in addition to the current session's reinforcers affecting behavior, previous sessions' reinforcers also affected behavior. As did Davison and Hunter (1979), Hunter and Davison found that the effect of a previous reinforcer ratio remained discernible after three or four sessions (Davison & McCarthy, 1988). Schofield and Davison (1997) used the same procedure. They showed that after many exposures to the random sequences, control became more localized, and the effects of previous sessions' reinforcers ceased to be discernible in the next session.

A different approach to the study of behavior in transition has been adopted by Mazur and his colleagues in a series of experi-

ments (Bailey & Mazur, 1990; Mazur, 1992, 1995, 1996, 1997; Mazur & Ratti, 1991). This research showed large changes in choice within a single session following unpredictable within-session changes in reinforcer ratios. The generic procedure consisted of a few sessions prior to a transition, a transition session in which the reinforcer ratio changed, and approximately four posttransition sessions. The finding that the reinforcers obtained in previous experimental sessions affected current performance (Davison & Hunter, 1979; Hunter & Davison, 1985; Schofield & Davison, 1997) was replicated by Mazur (1995, 1996). Specifically, in sessions following a transition, response proportions reverted towards pretransition levels. Moreover, both Schofield and Davison and Mazur (1992) concluded that rate of approach to asymptotic preference was independent of the magnitude of the change in the reinforcer ratio.

More recently, experimenters have studied choice under conditions in which multiple changes in reinforcer distributions occur regularly within single sessions. The assumption inherent in the generalized matching law is that the variables that control choice are aggregated over relatively large time frames, but as the above research shows, this time frame may be as short as, or even shorter than, a single session. Would performance be sensitive to multiple reinforcer-ratio changes occurring within a session? As Dreyfus (1991) pointed out, behavior allocation might not be sensitive to changes in reinforcer ratios that are either too brief or too rapid. Furthermore, the responsiveness of behavior to local changes in reinforcer ratios could vary as a function of overall reinforcer rate, as has been shown in steady-state research (Alsop & Elliffe, 1988; Elliffe & Alsop, 1996).

Dreyfus (1991) investigated both the effects of local changes in reinforcer ratios on relative time allocation in concurrent schedules and the time frame over which reinforcers controlling behavior were aggregated. He arranged a concurrent-schedule procedure in which relative reinforcer rates varied regularly within sessions. The number and frequency of shifts in relative reinforcer rate were manipulated, as was the overall reinforcer rate. In Experiment 1, Dreyfus arranged component durations of either 10 min or 30

min and sessions of either two or four components in length. In his Experiment 2, there were always four components per session, and the overall reinforcer rate was varied across conditions. His data showed that changes in relative time allocation followed changes in relative reinforcer rates more closely when there were fewer component changes per session and when components were longer. He also found that the correspondence between changes in relative time allocation and changes in relative reinforcer rates was greater when the overall rate of reinforcement was higher. Dreyfus concluded that although his data were silent on whether behavior was controlled at a molar or molecular level, it was clear that "if behavior is controlled by aggregates of reinforcers, that control is manifested over periods of time briefer than the duration of a typical experimental session" (p. 495).

Davison and Baum (2000) used a procedure introduced by Belke and Heyman (1994) and varied the speed of environmental change within sessions to investigate whether pigeons' choice changed more rapidly in more rapidly changing environments. Seven different reinforcer ratios (27:1, 9:1, 3:1, 1:1, 1:3, 1:9, and 1:27) were arranged in a session and were presented to the subject in random order. These components were not differentially signaled but were separated by 10-s blackouts. In Part 1, the speed of environmental change was manipulated by varying the number of reinforcers in each component between 4 and 12 over different conditions. In Part 2, they varied the overall reinforcer rate with components either 4 or 12 reinforcers in length. Davison and Baum used sensitivity to reinforcement as a function of each successive reinforcer delivery as a measure of rate of behavior change within components, and found that it remained constant as the number of reinforcers per component was varied from 4 to 12. However, across all conditions, the speed at which behavior changed within components was rapid, with sensitivity to reinforcement reaching high levels (about 0.60) after just six to eight reinforcer deliveries. They replicated Dreyfus's (1991) finding that the rate of behavioral change was greater when the overall rate of reinforcement was higher (analogous results have been reported in steady-state con-

current-schedule research by Alsop & Elliffe, 1988, and Elliffe & Alsop, 1996). Davison and Baum also reported striking regularities in the effects of individual reinforcers and sequences of reinforcers at a local level. Successive reinforcers obtained from the same alternative in a component had decreasing effects on behavior. Simultaneously, the carryover effects of reinforcers that had been obtained in the previous component decreased with successive reinforcer deliveries, and hence time, in the current component.

Davison and Baum (2000) found that the rate of behavior change did not change as a function of the number of reinforcers delivered in each component. However, a dimension of variation that they did not investigate is the range of within-session reinforcer ratios. This latter dimension was the focus of the present research. Using the same basic procedure as used by Davison and Baum, seven reinforcer ratios were arranged and presented in random order each session. Component length was constant at 10 reinforcers per component. In Experiment 1, the range over which the seven reinforcer ratios varied was systematically manipulated across conditions while the overall rate of reinforcement was held constant. These conditions varied from one in which the reinforcer ratio in components was always 1 to one in which the ratios were varied from concurrent extinction VI 27 s to concurrent VI 27 s extinction.

EXPERIMENT 1

METHOD

Subjects

The subjects were 6 homing pigeons numbered 61 to 66. All subjects apart from Bird 64 were maintained at $85\% \pm 15$ g of their free-feeding body weights. Bird 64 could not be maintained at this level of deprivation with the number of reinforcers arranged, and was therefore maintained at $90\% \pm 15$ g of its free-feeding body weight. Subjects were maintained at their designated body weights by being fed an appropriate amount of mixed grain following the daily training sessions. Water and grit were freely available to all subjects at all times.

Apparatus

Each bird was housed in a cage 380 mm high, 380 mm wide, and 380 mm deep. The back, left, and right walls of each cage were constructed of sheet metal, and the top, floor, and front wall of the cage consisted of metal bars. Each cage contained two wooden perches mounted 50 mm from the chamber floor. One was mounted 95 mm from, and parallel to, the front wall, and the second was mounted 95 mm from, and parallel to, the right wall. Illumination was provided by the ambient room lighting, and there was no sound attenuation.

The right wall of each cage contained three translucent response keys, 20 mm in diameter, centered 100 mm apart, 200 mm above the perches. The center key was always dark and inoperative. The two side keys could be transilluminated yellow and, when lit, operated by pecks exceeding a force of about 0.10 N. A hopper containing wheat was located behind an aperture (50 mm by 50 mm) situated 145 mm below the center key. During reinforcer delivery, the hopper was raised and illuminated for 2.5 s and the keylights were extinguished. From the home cages, other pigeons working on other experiments were visible and audible, but no staff entered the room while the experiments were running. The experimental room was lit from 1:00 a.m. to 4:00 p.m., and sessions began at 1:30 a.m. The subjects were tested successively, and all experimental events were arranged on an IBM-PC compatible computer running MED-PC[®] software, situated remote from the experimental cages. The computer also recorded the time at which every event occurred during the experimental sessions.

Procedure

A two-key concurrent-schedule procedure was used, and sessions began with both side keys lit yellow. Once 10 reinforcers had been obtained, both keylights were extinguished, and a 10-s blackout ensued. At the conclusion of the blackout, the next component began with the side keys again lit yellow. Each session consisted of seven components. Sessions were conducted daily, and ended in blackout following the completion of the seven components or after 45 min had elapsed, whichever occurred first.

Immediately prior to the beginning of each component, the computer selected the next component randomly without replacement from a list (Table 1). Thus, subjects were exposed to seven different pairs of concurrent VI schedules during each session. Reinforcers were scheduled according to a single exponential VI 27-s schedule ($p = .037/s$). Once a reinforcer was arranged, it was allocated to one of the alternatives according to the probability selected, as shown in Table 1. Reinforcers were dependently scheduled (Stubbs & Pliskoff, 1969): Once a reinforcer was arranged on one alternative, no further reinforcers could be arranged until it had been obtained. A 2-s changeover delay (COD; Herrnstein, 1961) prevented responses from producing an arranged reinforcer on a given alternative until 2 s had elapsed since the first response on that alternative following a response on the other alternative. Once a reinforcer was arranged on an alternative, and if the COD had completed timing, the next effective response to that alternative was reinforced.

The subjects had experience on a foraging procedure similar to that reported by Jones and Davison (1996). Thus, neither magazine nor key-peck training was necessary. However, some preliminary training on the experimental procedure was carried out. This training was used to reduce gradually the probability of reinforcement per second to the experimental value of .037, and to introduce the COD. During preliminary training, the probability of a reinforcer being allocated to either key was always .5.

The sequence of experimental conditions, relative reinforcer probabilities, and the arranged reinforcer ratios in each component are shown in Table 1. Across conditions, the range over which the seven reinforcer ratios varied was manipulated. Experiment 1 consisted of Conditions 1 to 8 and Condition 16. In Conditions 1 and 6, the reinforcer ratios were varied from 1:27 to 27:1, the same range used by Davison and Baum (2000). Condition 2 was discarded due to a programming error. The remaining conditions presented varying ranges of reinforcer ratios, from the least extreme, Condition 8, in which the reinforcer ratio was always 1:1, to Condition 16, in which they varied from EXT:1 to 1:EXT. No stability criterion was in effect. However, given the

Table 1

Sequence of experimental conditions and relative reinforcer probability (shown as probability of reinforcement on the left alternative) for each of the seven components in each condition of both experiments. The overall probability of reinforcement per second was constant at .037 throughout.

Condition	Component	Relative reinforcer probability p (L)	Reinforcer ratio (L:R)
1	1	.0357	1:27
	2	.1000	1:9
	3	.2500	1:3
	4	.5000	1:1
	5	.7500	3:1
	6	.9000	9:1
	7	.9643	27:1
3	1	.1111	1:8
	2	.2000	1:4
	3	.3333	1:2
	4	.5000	1:1
	5	.6667	2:1
	6	.8000	4:1
	7	.8889	8:1
4	1	.0602	1:15.63
	2	.1379	1:6.25
	3	.2857	1:2.5
	4	.5000	1:1
	5	.7143	2.5:1
	6	.8621	6.25:1
	7	.9398	15.63:1
5	1	.2268	1:3.38
	2	.3077	1:2.25
	3	.4000	1:1.5
	4	.5000	1:1
	5	.6000	1.5:1
	6	.6923	2.25:1
	7	.7714	3.38:1
6	1	.0357	1:27
	2	.1000	1:9
	3	.2500	1:3
	4	.5000	1:1
	5	.7500	3:1
	6	.9000	9:1
	7	.9643	27:1
7	1	.4000	1:1.5
	2	.4328	1:1.31
	3	.4663	1:1.14
	4	.5000	1:1
	5	.5337	1.14:1
	6	.5672	1.31:1
	7	.6000	1.5:1
8	1	.5000	1:1
	2	.5000	1:1
	3	.5000	1:1
	4	.5000	1:1
	5	.5000	1:1
	6	.5000	1:1
	7	.5000	1:1

Table 1

(Continued)

Condition	Component	Relative reinforcer probability p (L)	Reinforcer ratio (L:R)
9	1	.0357	1:27
	2	.0357	1:27
	3	.0357	1:27
	4	.9643	27:1
	5	.9643	27:1
	6	.9643	27:1
10	1	.0357	1:27
	2	.0357	1:27
	3	.2500	1:3
	4	.7500	3:1
	5	.9643	27:1
	6	.9643	27:1
11	1	.0357	1:27
	2	.0357	1:27
	3	.2500	1:3
	4	.5000	1:1
	5	.7500	3:1
	6	.9643	27:1
12	1	.0357	1:27
	2	.1000	1:9
	3	.2500	1:3
	4	.5000	1:1
	5	.7500	3:1
	6	.9000	9:1
13	1	.0357	1:27
	2	.0357	1:27
	3	.2500	1:3
	4	.5000	1:1
	5	.7500	3:1
	6	.9643	27:1
14	1	.0357	1:27
	2	.0357	1:27
	3	.2500	1:3
	4	.7500	3:1
	5	.9643	27:1
	6	.9643	27:1
15	1	.0357	1:27
	2	.0357	1:27
	3	.0357	1:27
	4	.9643	27:1
	5	.9643	27:1
	6	.9643	27:1
16	1	0	Ext:1
	2	.1000	1:9
	3	.2500	1:3
	4	.5000	1:1
	5	.7500	3:1
	6	.9000	9:1
	7	1	1:Ext

brief within-session exposures to each reinforcer ratio, 50 sessions were conducted for each condition to ensure that sufficient data were collected to be comparable to five sessions of steady-state data, as is commonly used in steady-state research (Davison & Baum, 2000). The data obtained from the last 35 sessions of each condition were used in the analyses.

RESULTS

The data used in the following analyses consisted of the time of every experimental event. To examine how preference changed within a component, the data were aggregated according to the left-key over right-key response ratios obtained in successive interreinforcer intervals. The logarithms of these ratios were calculated, and the data were plotted as a function of successive reinforcers for each component separately.

Figures 1 to 3 show the changes in individual-subject and group response allocation to the two alternatives as a function of successive reinforcers in a component for Conditions 5 (1:3.38 to 3.38:1), 3 (1:8 to 8:1), and 6 (1:27 to 27:1), respectively. Throughout, the group data are simply the raw data summed across all subjects. These conditions are shown because they are representative of performance across conditions, and show the effects of the range of within-session variation in the reinforcer ratios. In components in which reinforcer ratios other than 1:1 were arranged, both individual-subject and group preference moved towards the alternative providing the higher rate of reinforcement as successive reinforcers were delivered in a component, providing evidence of behavior change within a component. These figures also show that the degree of movement in preference over the course of 10 successive reinforcers in a component was a direct function of the reinforcer ratio arranged in that component: Preference moved further when more extreme reinforcer ratios were arranged. Furthermore, response biases to the right alternative for the behavior of all subjects except Birds 64 and 65 are evident in Figures 1 to 3, and were most notable for Bird 62.

For reasons of brevity, subsequent figures will show only group data, that is, the data summed across individual subjects. Thus, it is important to determine whether the group

data accurately represent the data obtained from individuals. As can be seen, the group data (bold lines) shown in Figures 1 to 3 provide an accurate representation of the general trends evident for all 6 subjects. More precisely, with each successive reinforcer obtained from either alternative, individual preference tended to become more extreme, moving towards the alternative providing the more frequent reinforcers. Thus, it seems reasonable to conclude that the group data provide an accurate account of the major features of the individual data.

Davison and Baum (2000) used multiple linear regression to analyze data obtained from a procedure similar to the present one. Multiple linear regression can be used to assess the linear contribution of the reinforcer ratio in the current component (current-component sensitivity) and the reinforcer ratio in the previous component (previous-component sensitivity or carryover sensitivity) on performance prior to each successive reinforcer delivery in the current component. The equation for this analysis is

$$\log\left(\frac{B_{ln}}{B_{rn}}\right) = a_{pn}\log\left(\frac{R_{lp}}{R_{rp}}\right) + a_{cn}\log\left(\frac{R_{lc}}{R_{rc}}\right) + \log c, \quad (2)$$

where B refers to responses emitted, R to reinforcers obtained, the subscripts l and r refer to the left and right alternatives, p and c refer to the previous and current components, and n refers to the reinforcer number in a component. For these analyses, the number of responses emitted on each of the two alternatives were summed separately for each component, for the periods before the first reinforcer in a component, between the first and second reinforcers in a component, and so on in each condition. Equation 2 was fitted to these data. For Condition 16, the equation was fitted to five data points because the most extreme reinforcer ratios arranged (EXT:1 and 1:EXT) were infinite. Furthermore, data from Condition 8 were not used in these analyses because no variation was arranged in the reinforcer ratios.

Figure 4 shows sensitivity to the current- and previous-component reinforcer ratios plotted as a function of successive reinforcer deliveries in a component. In all cases, sensitivity to the previous-component reinforcer

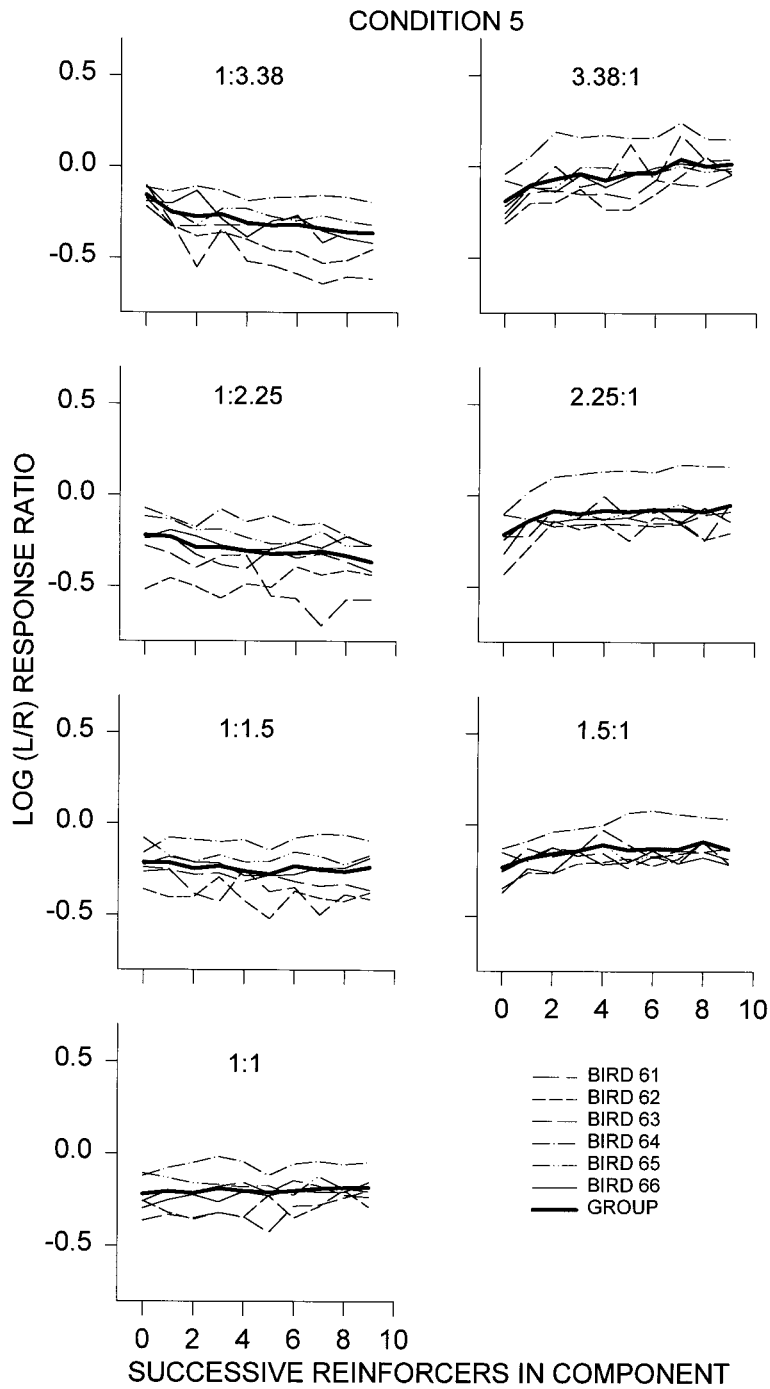


Fig. 1. Log response ratios emitted following each successive reinforcer delivered in each of the seven components of Condition 5. Response numbers for each alternative were summed across all 35 sessions. On each graph the separate plots show data from each individual subject and the group. On the x axis, 0 indicates the log response ratio emitted prior to the first reinforcer delivery in a component.

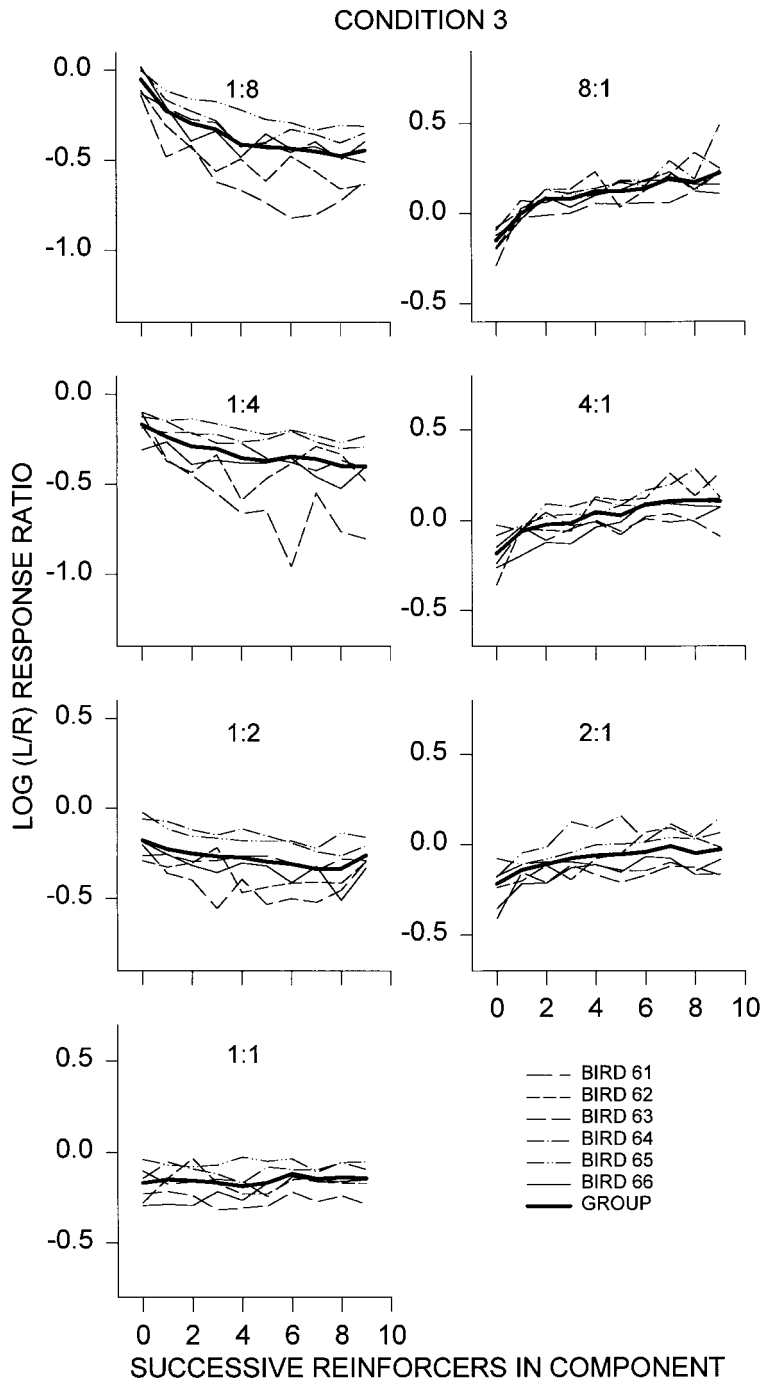


Fig. 2. Log response ratios emitted following each successive reinforcer delivered in each of the seven components of Condition 3. See Figure 1 for details.

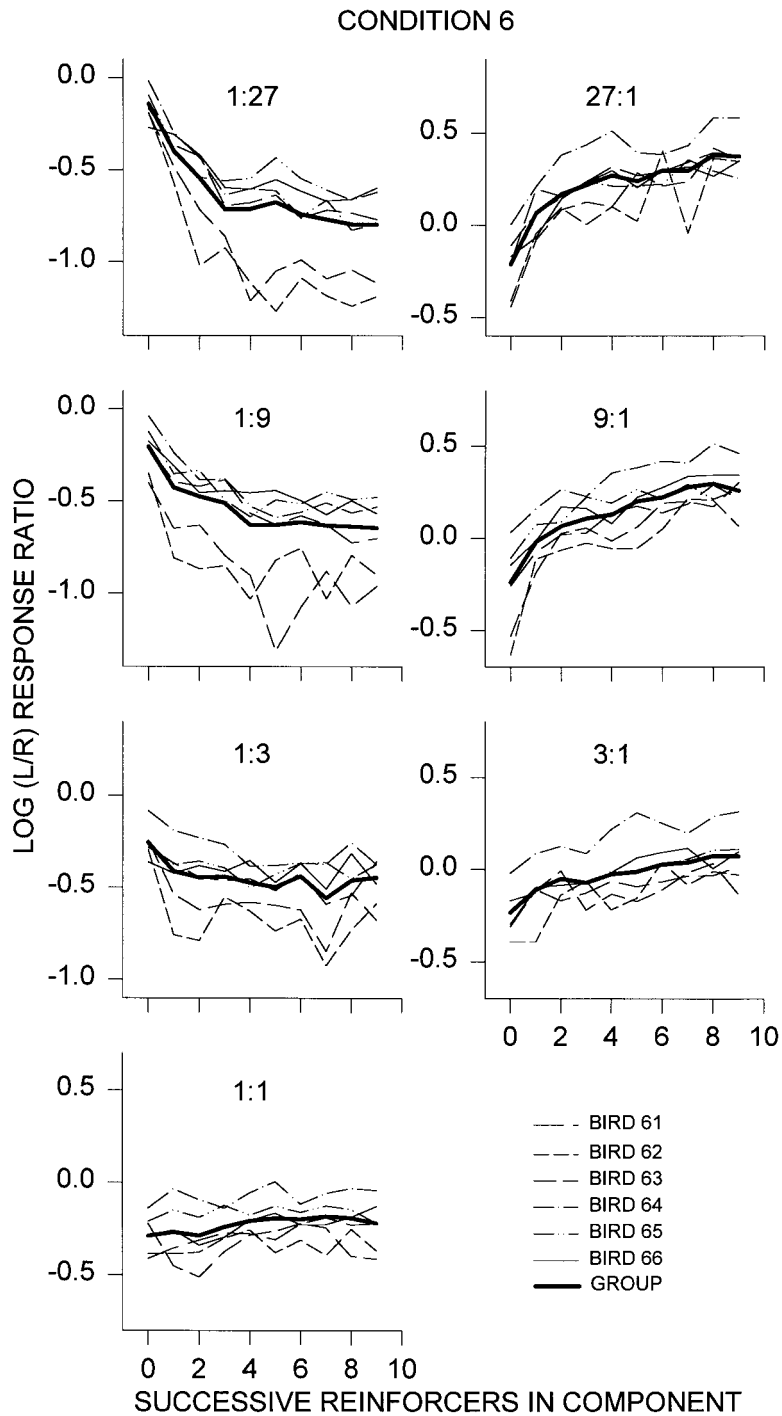


Fig. 3. Log response ratios emitted following each successive reinforcer delivered in each of the seven components of Condition 6. See Figure 1 for details.

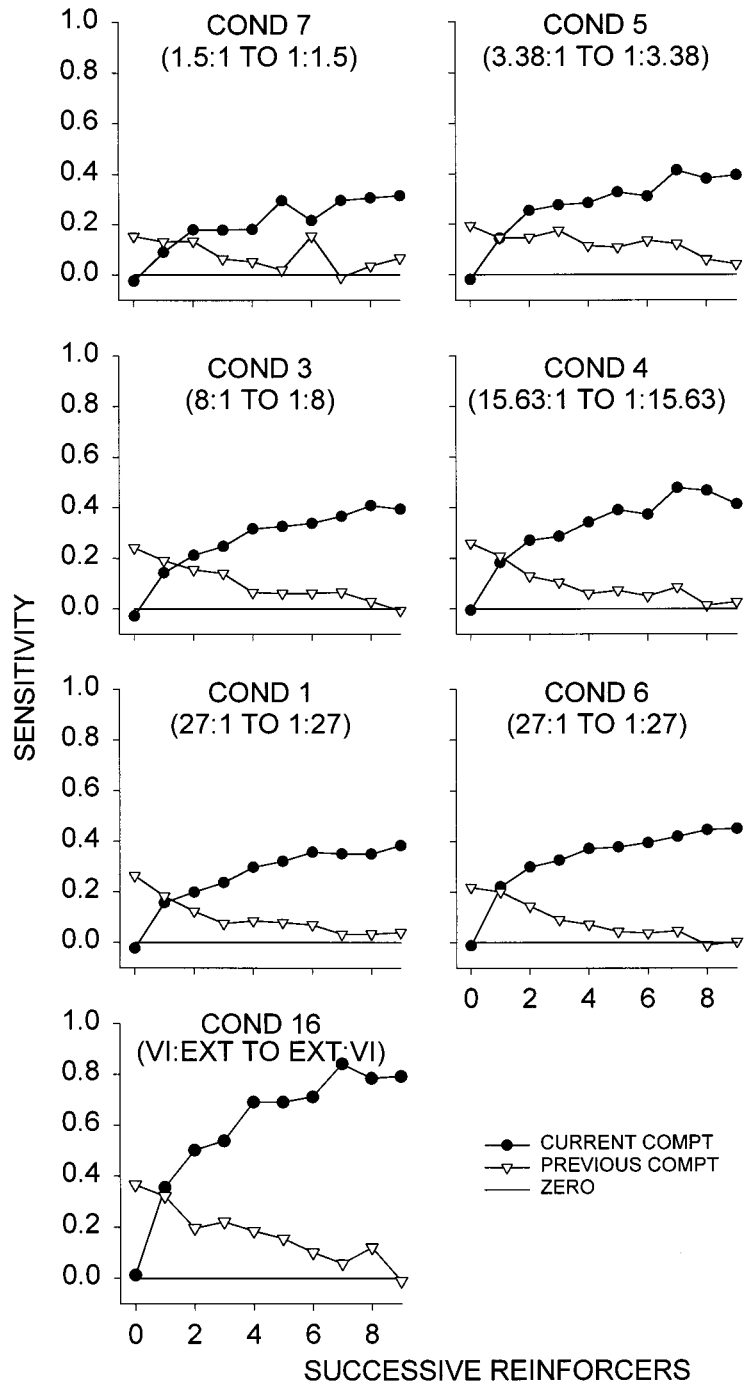


Fig. 4. Sensitivity-to-reinforcement values from multiple linear regressions between log response ratios and arranged log reinforcer ratios (Equation 2) in the previous and current components for each successive reinforcer delivery. Each panel represents data from a different condition. Response numbers were summed across all 6 subjects.

ratio started above 0 (range, 0.15 to 0.37), and fell towards 0 with successive reinforcer deliveries in the current component. This was true for the individual subjects, with one exception: Bird 66 in Condition 7 had an initial sensitivity of -0.08 . After four reinforcer deliveries in the current component, sensitivity to the previous-component reinforcer ratio was, in all cases, still above 0 (range, 0.08 to 0.22). For the group data, this pattern persisted even after six reinforcer deliveries (range, 0.04 to 0.15). Again, if the individual subjects are examined, the data from Condition 7 were more variable, as is evident in Figure 4. Nevertheless, after four reinforcers in a component in Condition 7, only 7 of the 30 sensitivity values calculated were less than 0. For the other conditions, after the fourth reinforcer delivery in a component, 33 of the 36 sensitivity values calculated were greater than 0 (range, -0.08 to 0.31).

Sensitivity to the current-component reinforcer ratio started close to 0 (range, 0 to -0.02), and increased to a mean value of 0.46 (range, 0.32 to 0.79) after nine reinforcer deliveries in a component. Again the individual data also showed similar effects. Initial sensitivity values were close to 0 (range, -0.14 to 0.15), with the more variable data again evident in Condition 7, and sensitivity generally increased with successive reinforcers. As described by Davison and Baum (2000), the difference between final sensitivity to the current-component reinforcer ratio and the starting sensitivity to the previous-component reinforcer ratio shows the effect of the between-components blackout.

The range of reinforcer ratios arranged in a session was increased in order across Conditions 7, 5, 3, 4, 1, and 6, with the most extreme range of reinforcer ratios arranged in Condition 16. To assess any trends in sensitivity across conditions, one-tailed nonparametric trend tests (Ferguson, 1966) were used to test individual-subject sensitivities separately after each reinforcer delivery. Sensitivity values obtained prior to any reinforcers being obtained in a component were omitted from these analyses, and separate trend tests were conducted using Conditions 3, 4, 5, 7, either Condition 1 or 6, and Condition 16.

The first test, using Condition 6, showed significant increasing trends in the value of sensitivity to the current-component reinfor-

er ratio as the range of reinforcer ratios increased following each of the first nine reinforcers in a component (for each test, $N = 6$ subjects, $k = 6$ conditions: $p < .05$, z scores ranged from 1.76 to 4.99). Identical tests, using data from Condition 1 in place of those from Condition 6, showed increasing trends after each reinforcer except the second and third in a component (for each significant test, $N = 6$ subjects, $k = 6$ conditions: $p < .05$, significant z scores ranged from 1.92 to 4.68). Both sets of trend tests indicate that sensitivity to the current-component reinforcer ratio following the first reinforcer increased with increasing range of reinforcer ratios.

Similar analyses to those reported above were carried out on the values obtained for sensitivity to the previous-component reinforcer ratio. Due to the nature of carryover, these trend tests included the sensitivity values obtained prior to the first reinforcer delivery in a component. Using Condition 6, significant increasing trends with increasing range of reinforcer ratios were evident before the first reinforcer delivery and after the first reinforcer delivery ($N = 6$ subjects, $k = 6$ conditions: $z = 1.99$, $p < .05$, $z = 2.30$, $p < .05$, for performance prior to any reinforcers and performance after the first reinforcer delivery, respectively). When Condition 1 was used, increasing trends were again evident before the first reinforcer delivery and after the first and third reinforcer deliveries in a component ($N = 6$ subjects, $k = 6$ conditions: $z = 4.37$, $p < .05$, $z = 2.38$, $p < .05$, $z = 1.76$, $p < .05$, for performance prior to any reinforcers and performance after the first and third reinforcer deliveries, respectively).

The analyses presented in Figures 1 to 3 showed that preference moved toward the alternative providing the higher rate of reinforcement over the course of the 10 reinforcers in a component. Moreover, the degree of this shift in preference was a direct function of the reinforcer ratio arranged in that component. To examine more closely the effects of individual reinforcers on behavior, a local analysis broke the data into log response ratios emitted in interreinforcement intervals following every possible sequence of reinforcers obtained in a condition. That is, before the first reinforcer in a component, a particular average response ratio was emitted, and a log response ratio could be calculated. After

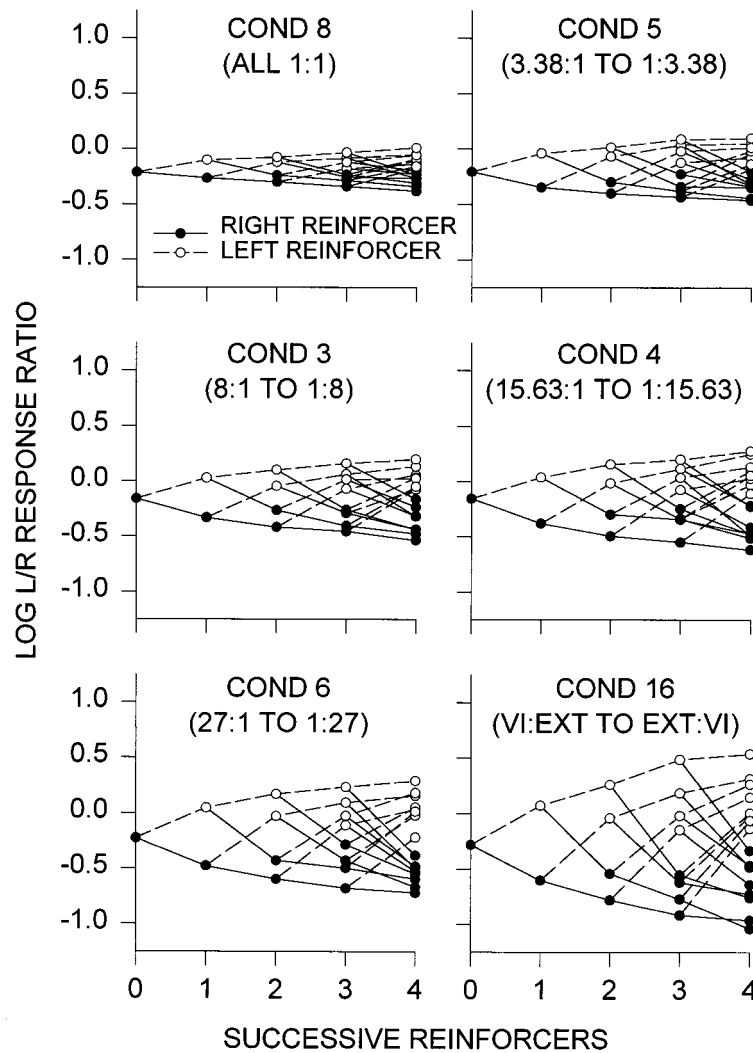


Fig. 5. Log response ratios emitted in interreinforcer intervals following all possible sequences of the first four reinforcers in a component in Conditions 3, 4, 5, 6, 8, and 16.

the first reinforcer and before the second, two log response ratios were available (one following a left reinforcer and one following a right reinforcer). After two reinforcers in a component, there were four log response ratios available, one for each possible two-reinforcer sequence.

Figure 5 shows the results of these analyses, plotting log response ratio as a function of every possible combination of the first four reinforcers in a component from Conditions 8, 5, 3, 4, 6, and 16. Prior to the first reinforcer in a component, the average log response ratios were always slightly negative, reflecting the biases to the right alternative

evident in the behavior of all subjects. Figure 4 showed that there was a similar amount of carryover of control by the previous-component reinforcer ratio over all conditions. It is important to note that the first points plotted in each panel of Figure 5 represent an average of performance affected by the various sequences of reinforcers obtained in the immediately preceding component. Figure 5 shows that, in general, each successive reinforcer produced a shift in subsequent preference, measured in terms of the log response ratio, toward the alternative from which it was obtained.

The effects of a single reinforcer on sub-

sequently emitted behavior were mediated by the context in which it was obtained. At one level, it was clear that preference moved more in conditions in which a greater range of reinforcer ratios was arranged. That is, the effect of a given reinforcer on behavior varied systematically as a function of the variation arranged in reinforcer ratios, as evidenced by the increasing spread of the trees in Figure 5. Thus, across conditions that varied in the range of reinforcer ratios arranged, identical sequences of reinforcers had the same directional effects on behavior but very different quantitative effects on behavior. Within conditions, the effects of a reinforcer on behavior were also dependent on the alternative from which the previous reinforcer was obtained: If a left-alternative reinforcer followed a right-alternative reinforcer, or vice versa ("disconfirmation," Davison & Baum, 2000), then the behavioral changes were greater than if consecutive reinforcers were obtained from the same alternative. Inspection of Figure 5 reveals that there were no occasions on which preference following a right reinforcer (a filled circle) appeared to be above preference following a left reinforcer (an open circle), highlighting the magnitude of the effects of disconfirmations on behavior.

Figure 6 shows the effects of a sequence of right (or left) reinforcers obtained in succession and the effects of a single disconfirmation at each sequential position. Of the 72 disconfirmations shown, all moved preference toward the alternative from which the reinforcer was obtained. To assess any trends in the effects of both disconfirmations and confirmations on behavior, the means of the absolute changes in log response ratio resulting from a disconfirmation and a confirmation were calculated for each condition at each sequential position. Across conditions, a nonparametric test for trend (Ferguson, 1966) showed the change in preference resulting from a disconfirmation increased as a function of the sequential position of the reinforcer in a component ($N = 8$ conditions, $k = 6$ reinforcers: $z = 4.05$, $p < .05$). A similar trend test showed the opposite effect for confirming reinforcers: Their effect on subsequent preference decreased with the sequential position of the reinforcer ($N = 8$ conditions, $k = 7$ reinforcers: $z = 3.12$, $p < .05$).

Figure 6 also shows that, across conditions, the log response ratios following disconfirmations were similar. If the log response ratio emitted prior to any reinforcer deliveries in a component (0 on the x axis) is used as a measure of indifference (naturally, this includes the bias evident in the subjects' behavior), then the log response ratios emitted following a disconfirmation can be compared to that point. Until about the third reinforcer in a component, disconfirmations resulted in shifts in preference beyond this measure of indifference. After further reinforcers, disconfirmations resulted in a log response ratio that was very similar to this measure of indifference. This effect was seen clearly across all conditions, irrespective of the range of reinforcer ratios arranged. Particularly later in a component, the log response ratios emitted following a disconfirmation were virtually superimposable.

DISCUSSION

The purpose of the present experiment was to examine the effects on behavioral adjustment of the range of within-session changes in the reinforcer ratio available from two response alternatives. In addition, the effects of individual reinforcers and sequences of reinforcers on behavior were examined in the context of the different ranges of reinforcer ratios available.

At the present stage of analysis, several findings have emerged. Within components, log response ratios adjusted rapidly as successive reinforcers were delivered in a component. The degree of this shift in preference was directly related to the reinforcer ratio in effect in that component. Using multiple linear regression analyses, the contributions of the previous-component and current-component reinforcer ratios to behavior prior to each successive reinforcer delivery in a component were assessed. Across conditions that arranged an increasing range of within-session changes in reinforcer ratios, increasing trends were evident in the values of sensitivity to the current-component reinforcer ratio calculated. A similar analysis showed evidence of an effect of this variation on sensitivity to the previous-component reinforcer ratio early in a component (Figure 4).

A local reinforcer-by-reinforcer analysis demonstrated regularities, both within and

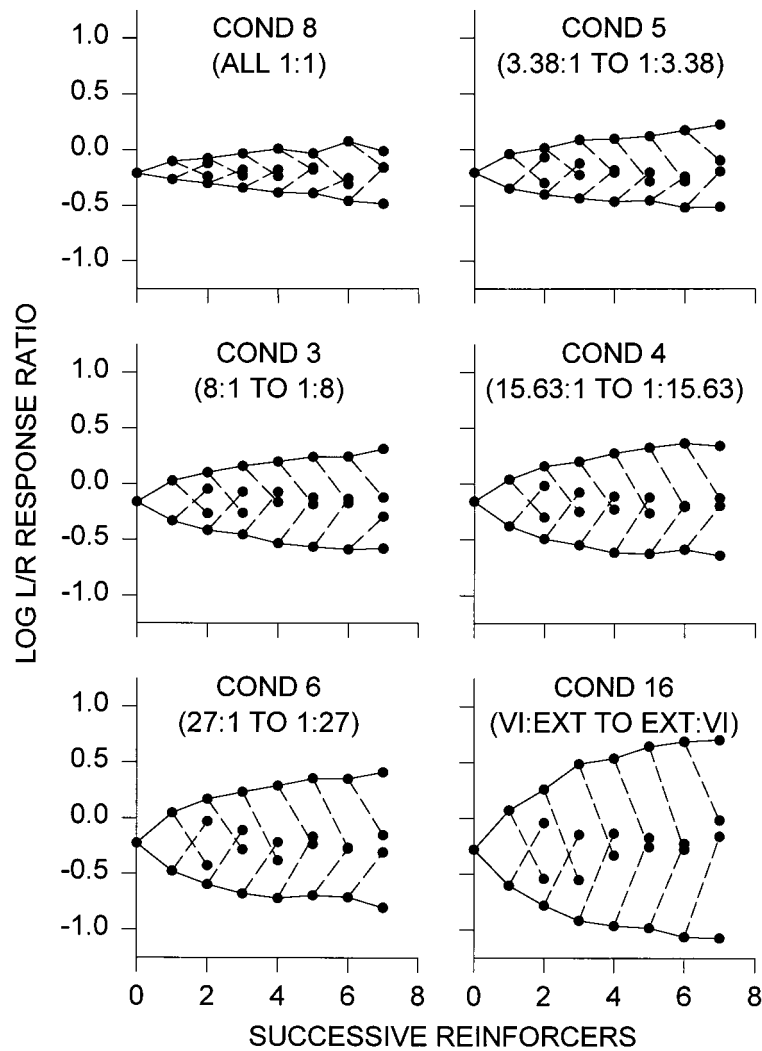


Fig. 6. Log response ratios emitted in interreinforcer intervals following successive same-alternative reinforcers (solid lines) in Conditions 3, 4, 5, 6, 8, and 16. The broken lines join "disconfirmations," where a reinforcer was obtained from the other alternative following sequences of successive same-alternative reinforcers.

across conditions, in the effects of each reinforcer on subsequently emitted behavior. Each successive reinforcer moved subsequent preference towards the alternative from which it was obtained. Across conditions, the effect of individual reinforcers increased as a function of the increasing range in the reinforcer ratios (Figure 6). The effects on preference of disconfirming reinforcers were substantially larger than confirming reinforcers. Moreover, the effects of disconfirmations increased with successive reinforcer deliveries in a component, whereas the effects of confirmations decreased.

Both the large shifts in log response-allocation ratios at the start of components and the values of sensitivity to the current-component reinforcer ratio attained within components demonstrated that behavioral adjustment was rapid under the present contingencies, as also reported by Davison and Baum (2000). In Condition 16, sensitivity values were obtained that were in the range considered normal for steady-state procedures (Taylor & Davison, 1983) in which each reinforcer ratio is in effect for about 15 to 30 sessions (Davison & McCarthy, 1988). Moreover, the largest shifts in preference occurred within the first three

or four reinforcers in a component (Figures 1 to 3), when sensitivity to the current-component reinforcer ratio rose to between 0.20 and 0.30.

The effects of the range of reinforcer ratios on sensitivity to reinforcement are not consistent with previous research that also used interval-based schedules but arranged less frequent changes in reinforcer ratios. For example, Mazur (1992), using a procedure in which a single change in the reinforcer ratio occurred during a transition session, found that the size of reinforcer ratios in concurrent VI VI schedules had no effect on the speed of adjustment. Schofield and Davison (1997) later generalized this finding to random between-session changes in reinforcer ratios. However, the procedures arranged in these previous experiments were fundamentally different from the present one. As outlined earlier, both Mazur and Schofield and Davison arranged single within- or between-session changes in reinforcer ratios. In contrast, the present experiment arranged conditions in which subjects were exposed to a range of reinforcer ratios within a single session. Two important procedural differences are evident: First, the rate of environmental change was more rapid in the present experiment than in these previous studies. Second, the previous experiments arranged only two reinforcer ratios and varied the size of the between- or within-session change, whereas the present experiment arranged a range of reinforcer ratios in a single session. Thus, the present experiment addressed a different dimension of environmental variability, and the present results suggest that, under the contingencies arranged here, the range of within-session changes arranged in reinforcer ratios had an effect on behavioral adjustment early in a component.

Given the widely held assumption that the variables that control choice are aggregated over substantial periods of time, the results of the local analyses (Figure 6) warrant careful consideration. Across all conditions, preference moved towards the alternative from which the previous reinforcer was obtained. There were also quantitative differences in the effects of individual reinforcers across conditions. Specifically, the effects of reinforcers on choice increased systematically with increases in the variation arranged in re-

inforcer ratios. It was also the case that, in all conditions, a reinforcer that followed a reinforcer obtained from the other alternative had a greater effect on behavior than a reinforcer obtained from the same alternative as the previous reinforcer. The effects of disconfirmations on behavior increased with the sequential position of the reinforcer in a component, whereas the effects of confirmations decreased. Moreover, early in a component, disconfirmations resulted in shifts in preference beyond an index of indifference, whereas later in a component disconfirmations returned preference to indifference.

To summarize, the present results showed that the effects of individual reinforcers differed in a number of ways depending on the context in which they were obtained. This context includes more molar factors, such as the range of reinforcer ratios arranged, but also includes local factors, such as the alternative from which the previous reinforcer was obtained and the number of successive reinforcers obtained from that alternative. Thus, the present results showed that each reinforcer had reliable directional and quantitative effects on behavior. The directional effects of each reinforcer could be predicted solely on the basis of which alternative had produced the reinforcer. In contrast, the quantitative effects of each reinforcer jointly depended on the number of reinforcers already obtained in a component, the alternative from which the previous reinforcer was obtained, and the range of variation arranged in the reinforcer ratios.

EXPERIMENT 2

The effects of the range of within-session changes in reinforcer ratios seen in Experiment 1 require further investigation. One distinction that can be made in terms of procedure is between the range of variation in the reinforcer ratios as was manipulated in Experiment 1 and what might be termed the variation in the reinforcer ratios themselves. It is possible for the component reinforcer ratios to be manipulated while the most extreme component reinforcer ratios are held constant. Thus, the range of variation would be held constant while still changing the amount of environmental variation.

Experiment 1 showed what appeared to be

an effect of the range of within-session reinforcer-ratio changes. However, the mechanism underlying this result is unclear. By keeping the most extreme component reinforcer ratios constant across conditions (1:27 or 27:1) and manipulating the other possible component reinforcer ratios, it may be possible to delineate the effects of the range of variation and the variation itself. For example, consider a condition with six components in which the reinforcer ratios are either 1:27 or 27:1, each with equal likelihood. This condition would have the same range of variation as one with seven reinforcer ratios from 1:27 to 27:1 (e.g., 1:27, 1:9, 1:3, 1:1, 3:1, 9:1, and 27:1), but the latter arrangement would contain more variation in the reinforcer ratios themselves. This approach was the focus of Experiment 2.

In this experiment, the subjects were trained on the same procedure as used in Experiment 1. However, in Conditions 9 through 15, the range of within-session changes in reinforcer ratios was constant at 1:27 to 27:1. The variation in the changes in reinforcer ratios was manipulated by changing the reinforcer ratios in effect in the less extreme components. Thus, Conditions 9 and 15 arranged the least variation in reinforcer ratios, because a session consisted of six components in which the reinforcer ratio was either 1:27 or 27:1. In contrast, Condition 12 (a replication of Conditions 1 and 6 from Experiment 1) contained the greatest variation.

METHOD

Subjects and Apparatus

The subjects were the same 6 homing pigeons used in Experiment 1, maintained as described in Experiment 1. The same apparatus used in Experiment 1 was used in Experiment 2.

Procedure

The subjects were trained on the same basic procedure as used in Experiment 1. However, in Conditions 9, 10, 14, and 15, only six components were arranged per session. In these conditions, sessions ended after the completion of the six components or after 45 min had elapsed, whichever occurred first.

The sequence of experimental conditions

for Experiment 2 is shown in Table 1. Across conditions, the range over which the reinforcer ratios varied was held constant. In Conditions 9 and 15, the reinforcer ratios arranged in the six components were either 1:27 or 27:1. In Conditions 10 and 14, six components were arranged, with the reinforcer ratios being 1:27, 1:27, 1:3, 3:1, 27:1, and 27:1. In Conditions 11 and 13, seven components were arranged: the six arranged in Conditions 10 and 14, with the addition of a 1:1 reinforcer-ratio component. Condition 12 was a replication of Conditions 1 and 6 from Experiment 1; thus, the reinforcer ratios were 1:27, 1:9, 1:3, 1:1, 3:1, 9:1, and 27:1. As with Experiment 1, no stability criterion was used, and 50 sessions were conducted for each condition. The data obtained from the last 35 sessions of each condition were used in the analyses.

RESULTS

Figures 7, 8, and 9 show the changes in individual-subject and group response allocation to the two alternatives as a function of successive reinforcers in a component for Conditions 9, 10, and 12, respectively. The data were aggregated in the same manner as those presented in Figures 1 to 3. The conditions shown are representative of performance across conditions. The effects of the reinforcer ratio arranged in a component can again be seen clearly. With successive reinforcer deliveries, both individual-subject and group preference moved rapidly towards the alternative providing the higher rate of reinforcement in a component. The degree of the movement in preference within a component was a function of the reinforcer ratio arranged in that component; that is, the greater the ratio, the larger the degree of movement.

Across conditions, the change in preference in a component with a given reinforcer ratio appeared to be similar, irrespective of the variation arranged in the reinforcer ratios. For instance, the preference changes in the 1:27 components in Conditions 9, 10, and 12 (Figures 7, 8, and 9, respectively) are similar across both the individual-subject and group data. The data for Birds 62 and 63 were sometimes more variable than those for the other subjects. This variability arose in components in which very few responses were

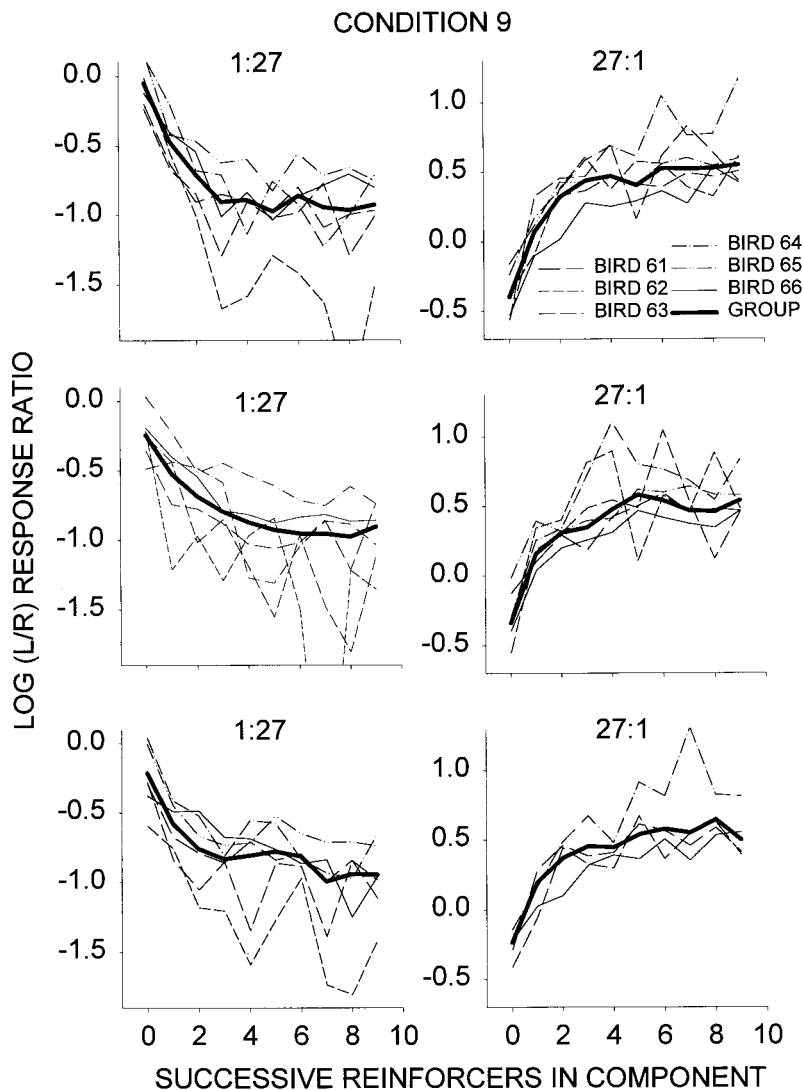


Fig. 7. Log response ratios emitted following each successive reinforcer delivered in each of the six components of Condition 9. See Figure 1 for details.

emitted on the alternative providing the lower rate of reinforcement. Thus, the response ratios were sometimes extreme.

Equation 2 was fitted to the data. Figure 10 shows sensitivity to the current- and previous-component reinforcer ratios plotted as a function of successive reinforcer deliveries in a component. In all cases, sensitivity to the previous-component reinforcer ratio started above 0 both for the group (range, 0.25 to 0.29) and for each individual subject (range, 0.17 to 0.47). Sensitivity then fell towards 0 with successive reinforcer deliveries in the

current component. After four reinforcer deliveries in the current component, sensitivity to the previous-component reinforcer ratio was still above 0 (range, 0.05 to 0.15). For individual subjects after four reinforcer deliveries, 34 of 42 sensitivity values were still above 0 (range, -0.17 to 0.31). Sensitivity to the current-component reinforcer ratio started close to 0 (range, -0.02 to 0.01) and increased with successive reinforcer deliveries to a mean of 0.58 (range, 0.53 to 0.65) after nine reinforcer deliveries in a component. As previously, the individual data showed similar

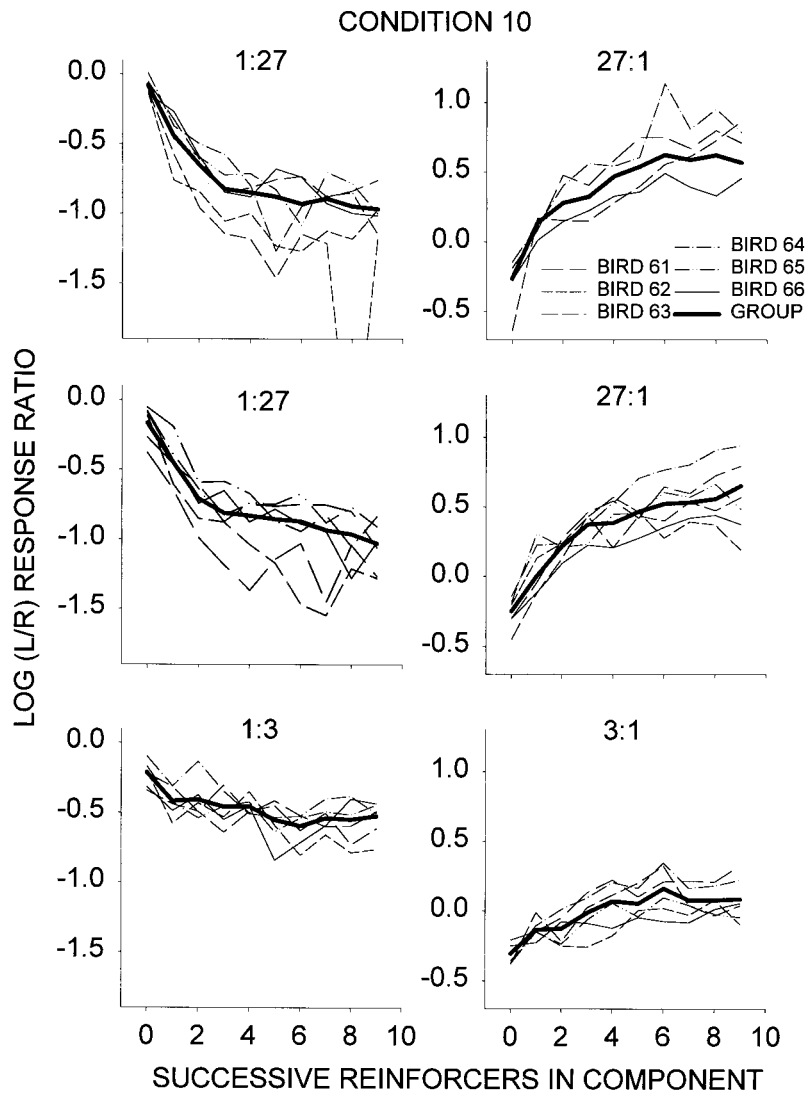


Fig. 8. Log response ratios emitted following each successive reinforcer delivered in each of the six components of Condition 10. See Figure 1 for details.

effects, with initial sensitivity values close to 0 (range, -0.15 to 0.14) and increasing with successive reinforcer deliveries. Overall, the major features of the data were similar to those seen in Experiment 1. Each replication provided a reasonable approximation of the original condition, although in general the sensitivity values obtained in the replications were somewhat larger than those obtained in the original conditions.

The variation in the reinforcer ratios arranged in a session was increased in order across Conditions 9, 10, 11, and 12, and then

decreased across Conditions 12, 13, 14, and 15. In the same manner as in Experiment 1, identical analyses were conducted on the data obtained from each individual subject to assess any increasing trends in sensitivity values across conditions. One-tailed nonparametric trend tests (Ferguson, 1966) were used to examine the data for any increasing trends in sensitivity to the current-component reinforcer ratio with decreasing arranged variation in the reinforcer ratios after each successive reinforcer delivery. As in Experiment 1, sensitivity values obtained prior to any reinforcers

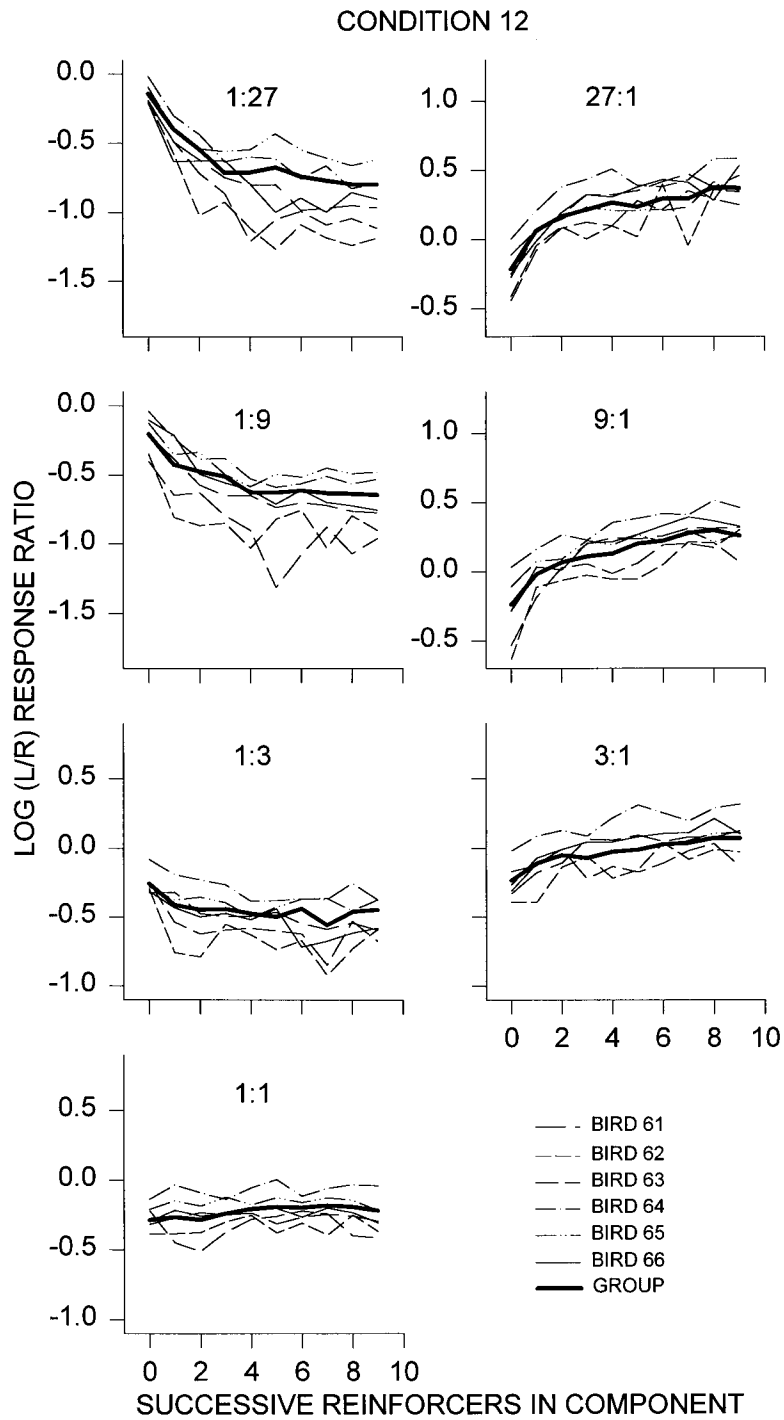


Fig. 9. Log response ratios emitted following each successive reinforcer delivered in each of the seven components of Condition 12. See Figure 1 for details.

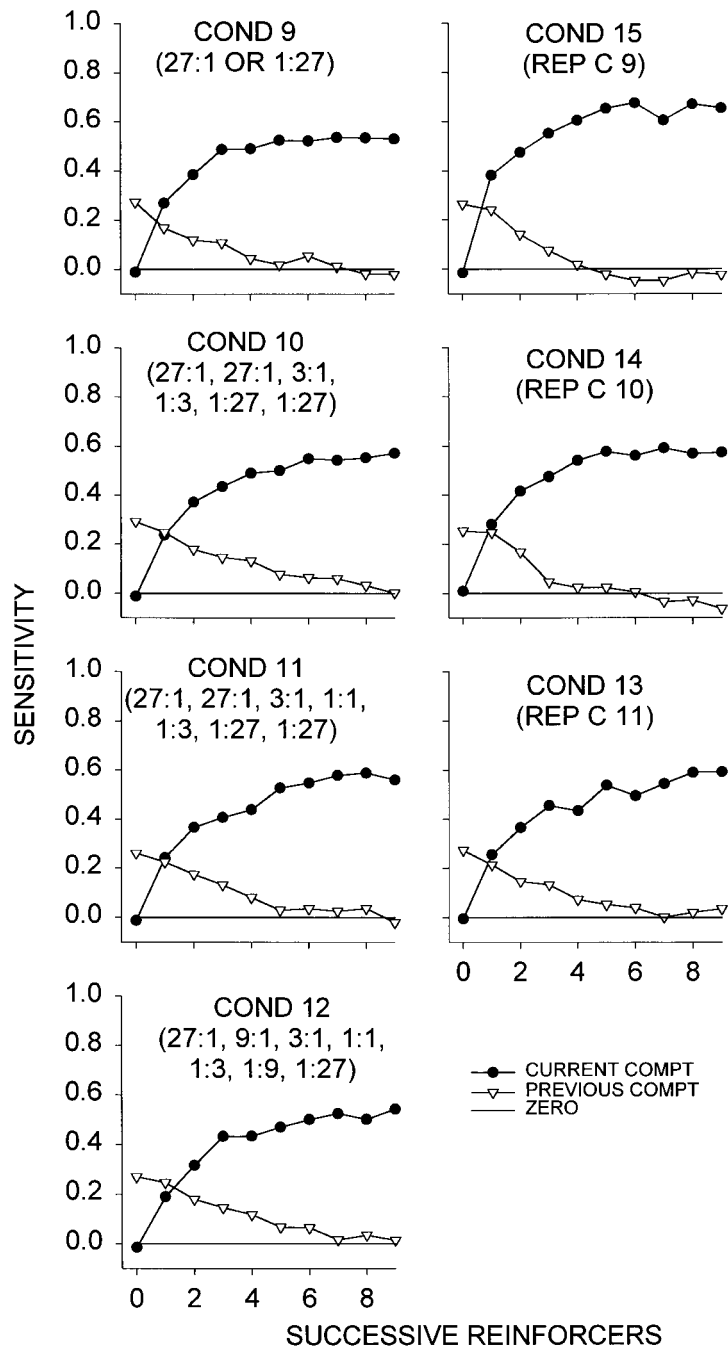


Fig. 10. Sensitivity-to-reinforcement values from multiple linear regressions between log response ratios and arranged log reinforcer ratios (Equation 2) in the previous and current components for each successive reinforcer delivery. Each panel represents data from a different condition. Response numbers were summed across all 6 subjects.

being obtained in a component were omitted from these analyses, and separate trend tests were conducted using Conditions 9, 10, 11, and 12 and Conditions 12, 13, 14, and 15.

The first test, using Conditions 9, 10, 11, and 12, showed a significant increasing trend in the value of sensitivity to the current-component reinforcer ratio as the variation arranged in the reinforcer ratios was increased only following the third reinforcer in a component ($N = 6$ subjects, $k = 4$ conditions: $z = 2.08$, $p < .05$; for performance after the third reinforcer delivery). Identical tests using data from Conditions 12, 13, 14, and 15 showed increasing trends after every reinforcer delivery in a component except the ninth ($N = 6$ subjects, $k = 4$ conditions: $p < .05$; significant z scores ranged from 2.36 to 3.47).

Analyses similar to those reported above were carried out on the values obtained for sensitivity to the previous-component reinforcer ratio. Due to the nature of carryover, these trend tests included the sensitivity values obtained prior to the first reinforcer delivery in a component. Using Conditions 9 to 12, a significant increasing trend with increasing variation of reinforcer ratios was evident only after the fifth reinforcer delivery in a component ($N = 6$ subjects, $k = 4$ conditions: $z = 1.80$, $p < .05$). When Conditions 12 to 15 were used, increasing trends were evident after the fifth, sixth, and eighth reinforcer deliveries in a component ($N = 6$ subjects, $k = 4$ conditions: $z = 2.36$, $p < .05$; $z = 2.36$, $p < .05$; $z = 2.36$, $p < .05$, for performance after the fifth, sixth, and eighth reinforcer deliveries, respectively).

To examine more closely the effects of individual reinforcers on choice, the analyses reported in Experiment 1 were repeated. The data were broken into log response ratios emitted in interreinforcer intervals following every possible sequence of reinforcers obtained in a component. Figure 11 shows the results of these analyses, plotting log response ratios as a function of every possible combination of the first four reinforcers in a component from Conditions 9 to 15. As in Experiment 1, the average log response ratios emitted prior to the first reinforcer in a component were always slightly negative. Again, it was the case that each successive reinforcer produced a shift in subsequent preference to-

ward the alternative from which it was obtained.

The extreme reinforcer ratios arranged in Conditions 9 and 15 resulted in several sequences of reinforcers never occurring in these conditions across the 35 sessions analyzed. Successive same-alternative reinforcers in these two conditions did move preference further than in Conditions 10 to 14. However, across conditions, the trees were generally similar in structure. There was only one occasion on which preference following a right reinforcer was greater than preference following a left reinforcer. This occurred in Condition 9, and was most likely a result of there being very few occurrences of the sequences in the center of the trees.

In the same manner as in Experiment 1, Figure 12 shows the effects of reinforcers later in components for part of the tree structure. It shows a sequence of right (or left) reinforcers obtained in succession and the effects of a single disconfirmation at each sequential position. As in Experiment 1, all disconfirmations moved preference toward the alternative from which the reinforcer was obtained. To assess trends in the effects of both disconfirmations and confirmations on behavior, the means of the absolute changes in log response ratio resulting from a disconfirmation and a confirmation were calculated for each condition at each sequential position. Across conditions, a nonparametric test for trend (Ferguson, 1966) showed the change in preference resulting from a disconfirmation increased as a function of the sequential position of the reinforcer ($N = 7$ conditions, $k = 6$ reinforcers: $z = 5.40$, $p < .05$). A similar trend test showed the opposite effect for confirming reinforcers, their effect on subsequent preference decreasing with the sequential position of the reinforcer ($N = 7$ conditions, $k = 7$ reinforcers: $z = 4.54$, $p < .05$). These findings replicated those from Experiment 1.

The effects on behavior of the disconfirmations shown in these graphs indicated some differences from those seen in Experiment 1. Again, using the average log response ratio emitted prior to any reinforcer deliveries in a component as an index of indifference, behavior following a disconfirmation can be compared to a measure of indifference. In both Conditions 9 and 15 (27:1 or

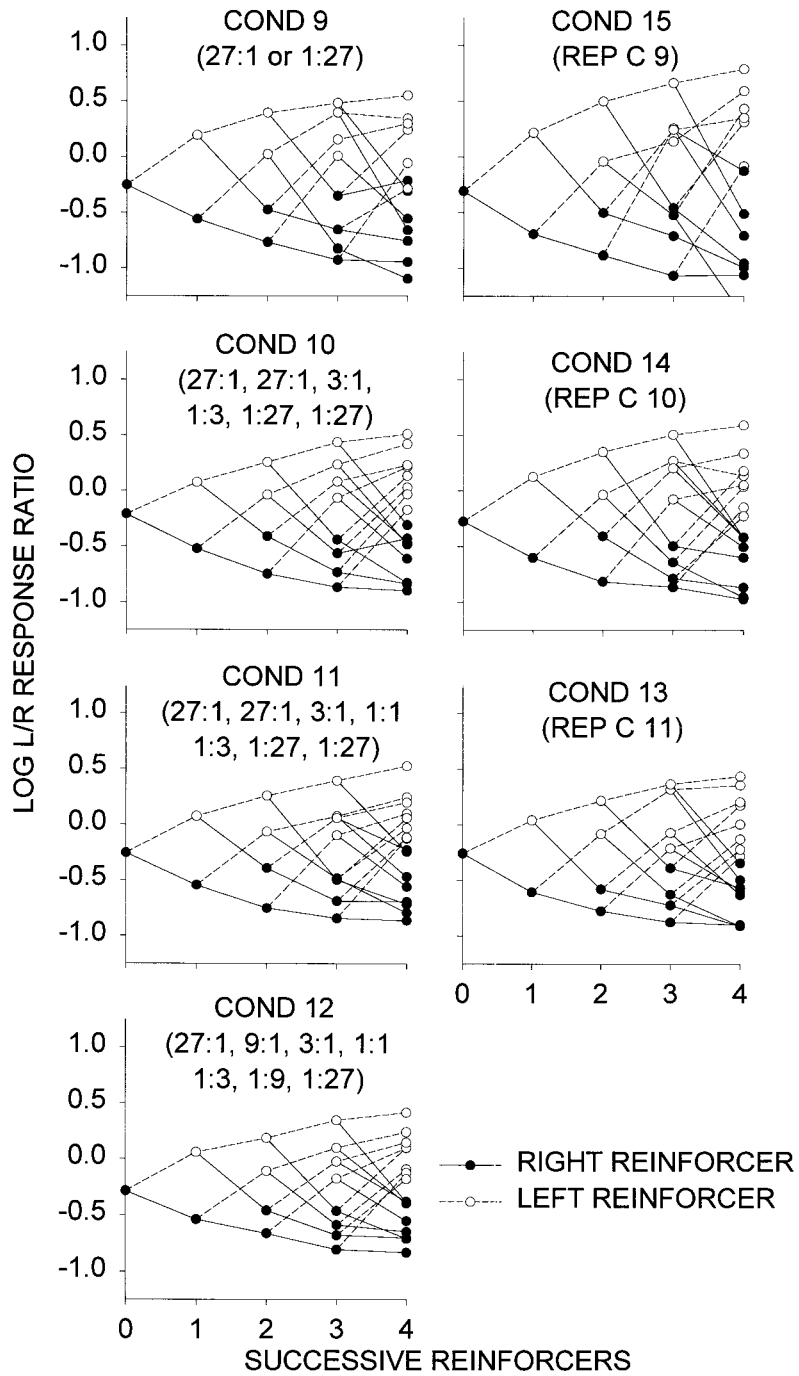


Fig. 11. Log response ratios emitted in interreinforcer intervals following all possible sequences of the first four reinforcers in a component in Conditions 9, 10, 11, 12, 13, 14, and 15.

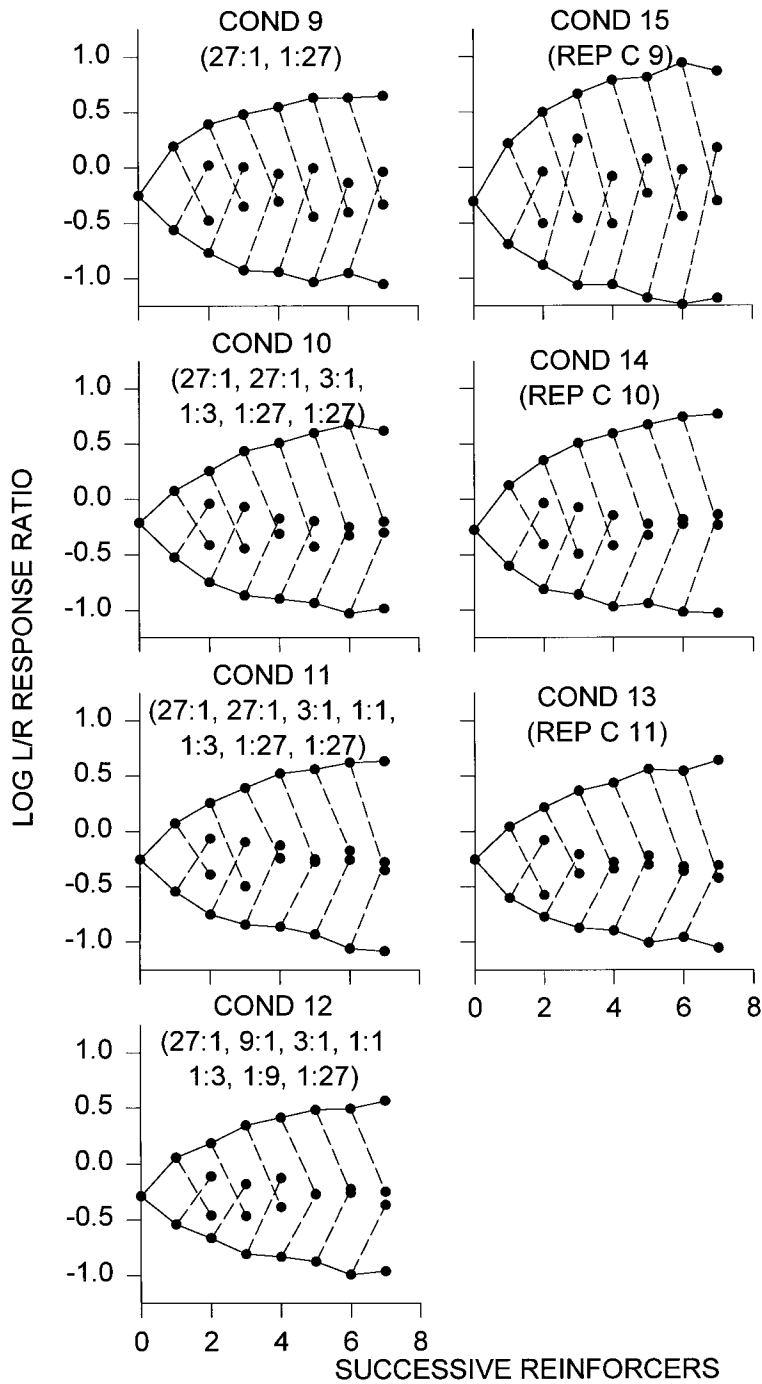


Fig. 12. Log response ratios emitted in interreinforcer intervals following successive same-alternative reinforcers (solid lines) in Conditions 9, 10, 11, 12, 13, 14, and 15. The broken lines join "disconfirmations," where a reinforcer was obtained from the other alternative following sequences of successive same-alternative reinforcers.

1:27), each disconfirmation moved preference beyond this measure of indifference towards the alternative from which it was obtained. This occurred regardless of the sequential position of the disconfirmation. In Conditions 10 and 14 (27:1, 27:1, 3:1, 1:3, 1:27, or 1:27), these crossovers ceased occurring when the disconfirmation was the fifth or sixth reinforcer in a component, respectively, and the crossovers were not as substantial beyond the third reinforcer in a component. In the remaining conditions, the crossovers ceased after either the third or fourth reinforcer in a component.

GENERAL DISCUSSION

Experiment 2 was an investigation of the effects of the variation in the within-session changes in reinforcer ratios. Over seven conditions, the range of reinforcer ratios arranged was always 1:27 to 27:1, while the reinforcer ratios in effect in the less extreme components were varied. Thus, the variation of the reinforcer ratios was varied from conditions in which six components were arranged and the reinforcer ratios were either 1:27 or 27:1 (Conditions 9 and 15) to one in which seven components were arranged and the reinforcer ratios were varied from 1:27 to 27:1 (Condition 12, a replication of Conditions 1 and 6 of Experiment 1).

Multiple linear regression analyses showed the degree of control exerted over behavior by the reinforcer ratio in effect in both the current and previous components. For Conditions 9 to 12, an increasing trend with changes in the variation of the reinforcer ratios was evident only following the third reinforcer delivery in a component. However, for Conditions 12 to 15, there were significant increasing trends in sensitivity to the current-component reinforcer ratio throughout a component as a function of decreasing variation in the reinforcer ratios. Furthermore, increasing trends in sensitivity to the previous-component reinforcer ratio with increasing variation in reinforcer ratios were evident later in components. Given these discrepancies, any conclusions on the basis of the results at this level of analysis would be tenuous at best.

The analyses at a local level showed a context effect on the effects of individual rein-

forcers on behavior. Sequences of same-alternative reinforcers moved preference to more extreme levels when their occurrence was more frequent (i.e., Conditions 9 and 15). These analyses showed regularities in the effects of individual reinforcers similar to those seen in Experiment 1: There were diminishing effects of successive same-alternative reinforcers and increasing effects of disconfirming reinforcers.

In Experiment 1, a disconfirming reinforcer that followed a sequence of same-alternative reinforcers reversed preference when it occurred early in a component and returned preference to close to indifference when it occurred later in a component (Figure 6). This was not the case in Experiment 2. When there were relatively few components in which less extreme reinforcer ratios were arranged (i.e., Conditions 9, 10, 14, and 15) and therefore successive same-alternative reinforcers were more common, preference reversals caused by disconfirmations persevered much later into components (Figure 12). This was most clearly seen in Conditions 9 and 15, in which the reinforcer ratios arranged in the six components were always either 1:27 or 27:1 (see Figure 12). Naturally, this arrangement led to a high frequency of runs of reinforcers on one alternative or the other. In these conditions, however, preference reversals always occurred following a disconfirming reinforcer in a sequence of same-alternative reinforcers: Even following six successive reinforcers on, say, the left alternative, a single right reinforcer moved preference to the right alternative.

The present experiments addressed the question of how one aspect of environmental variation, namely range of variation in reinforcer ratios, affects the acquisition of choice. Previous research had shown that the more frequently an environment changed, the faster behavior changed in response to those changes (e.g., Davison & Hunter, 1979; Hunter & Davison, 1985; Schofield & Davison, 1997; Shettleworth & Plowright, 1992). Using the same procedure as the present experiment, Davison and Baum (2000) found no effect of speed of environmental change on behavior in transition, implying a limit to environmental control over speed of behavioral adjustment. Research that had manipulated the magnitude of within- or between-session

changes in reinforcer ratios (e.g., Mazur, 1992, 1995; Schofield & Davison, 1997) found no effect of the size of these changes on behavioral adjustment. However, the present procedure manipulated the range of reinforcer ratios available, rather than just the size of a single change in reinforcer ratios.

To summarize, the present experiments have provided information on how range and variation of reinforcer ratios affect preference in this procedure. At a molar level of analysis, sensitivity to the current-component reinforcer ratio increased as the range of reinforcer ratios increased. Likewise, early in a component, sensitivity to the previous-component reinforcer ratio was higher when the range of reinforcer ratios was greater.

More compelling, however, were the analyses conducted at a local reinforcer-by-reinforcer level (Figures 5, 6, 11, and 12). These analyses showed that a reinforcer obtained from one alternative or the other resulted in a shift in preference towards that alternative in all experimental conditions. However, the size of this shift in preference was affected by numerous factors. The range of reinforcer ratios arranged had a clear effect on the log response ratios seen after various sequences of reinforcers: As the range of reinforcer ratios was increased, so too did the size of the effect of a given reinforcer on preference (see Figures 5 and 6). If several successive reinforcers were obtained from the same alternative, each reinforcer had diminishing effects on behavior in all conditions in both experiments. In contrast, disconfirming reinforcers in these sequences of same-alternative reinforcers had comparatively large effects on preference, and these effects increased in size throughout a component (see Figures 6 and 12). However, the effects of these disconfirmations depended on the distribution of reinforcer ratios arranged: When seven different reinforcer ratios were arranged, evenly distributed around 1:1 (e.g., those in Experiment 1), disconfirmations early in a component reversed preference, whereas later in a component disconfirmations returned preference to approximate indifference. In Conditions 9 and 15, in which the reinforcer ratios were either 1:27 or 27:1 (and, to a lesser extent, in Conditions 10 and 14), however, disconfirmations reversed preference regardless of where they occurred in

a component. Across all conditions, any given reinforcer had a larger effect on preference if it was obtained from the alternative opposite to the previous reinforcer than if it had been obtained from the same alternative as the previous reinforcer.

The context effect shown by reinforcers having larger effects on behavior when the range of reinforcer ratios arranged was more extreme presents a problem for the quantitative model proposed by Davison and Baum (2000). Because this model is an accumulation model, the effect of each reinforcer is determined by the reinforcer and the current levels of the reinforcer accumulations and not by the context in which the accumulations occurred.

The comparatively large effects of disconfirming reinforcers on log response ratios, and the fact that, across conditions, with the notable exceptions of Conditions 9 and 15, the log response ratios emitted following a disconfirming reinforcer were virtually identical, highlight some of the strong regularities seen across conditions. Of immediate interest is what happens to preference with successive reinforcers following a disconfirmation. To examine this, various sequences of reinforcers were isolated and compared. The first was a sequence of three successive left- or right-alternative reinforcers from the beginning of a component. The second was three successive left- or right-alternative reinforcers following a left-right (LR) or right-left (RL) reinforcer sequence, respectively. Thereafter, the same sequences were examined following LLR, RRL, LLLR, RRRL, and finally LLLL and RRRRL sequences. Sequences were examined both where the three successive reinforcers were obtained from the opposite alternative to the disconfirmation and from the same alternative as the disconfirmation.

Figure 13 shows the results of these analyses for Conditions 3, 6, and 8. These conditions are representative of performance in other conditions. It should be noted that in Experiment 2 there were many fewer occurrences of these sequences, and therefore no data from Experiment 2 are shown. The left panels show the effects on preference of three successive same-alternative reinforcers obtained from the opposite alternative to the disconfirmation which started the sequence, and the right panels show the effects of three

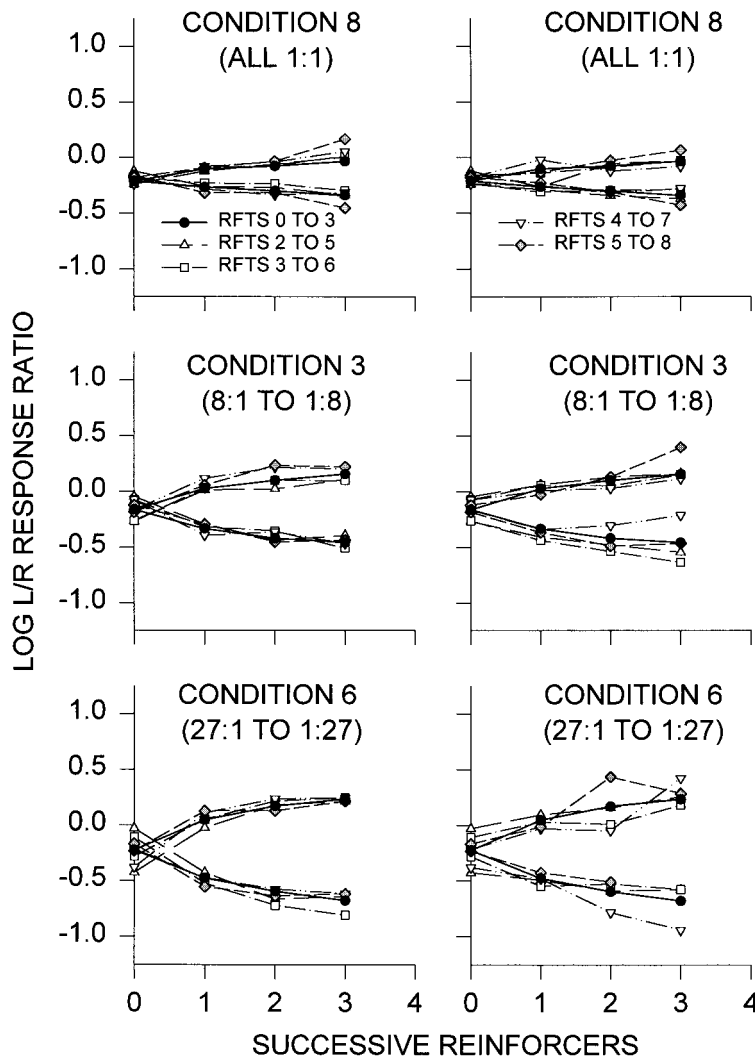


Fig. 13. Log response ratios emitted in interreinforcer intervals following sequences of three successive same-alternative reinforcers in Conditions 3, 6, and 8. The three successive reinforcers occurred at the beginning of a component or following “disconfirmations,” which were the second, third, fourth, or fifth reinforcer in a component following a sequence of successive reinforcers obtained on the other alternative (left panels) or the same alternative (right panels).

successive reinforcers obtained from the same alternative as the disconfirmation. For comparison purposes, the data are all plotted from 0 to 3 on the *x* axis, and the symbols indicate where in a component the sequences were actually obtained.

Figure 13 shows again the crossing over of preference following the first two disconfirmations in a component seen earlier (Figure 6). In general, the crossing over was most prevalent following the first disconfirmation (open triangles), followed by preference after

the second disconfirmation (open squares). Thereafter, preference following a disconfirmation was more similar to preference seen at the beginning of a component (filled circles). Figure 13 shows that log response ratios with each successive reinforcer following a disconfirmation were largely superimposable within each condition. Thus, the crossover effect following a disconfirmation was immediately negated by the next reinforcer, irrespective of whether the next reinforcer was obtained from the same or the opposite al-

ternative. Each reinforcer thereafter moved preference to a similar level to that which would have been attained if the reinforcer had been obtained at the beginning of a component, or elsewhere.

These analyses underline the local nature of control under these conditions and show how preference changed in a regular way in the center of the tree structure. The tree diagrams as a whole show that control was most local in Conditions 9 and 15 (1:27 or 27:1). In these conditions, for example, following six successive reinforcers from the same alternative, a single reinforcer obtained from the other alternative resulted in a reversal of preference. Obviously, given the contingencies arranged, this reinforcer was unlikely to be followed by more reinforcers on the same alternative. This pattern of responding suggests a highly local locus of control. In other conditions, only the first two or three disconfirmations following a sequence of same-alternative reinforcers resulted in a reversal of preference. Although this too reflected a local locus of control, behavior did not follow individual reinforcers to the same extent as seen in Conditions 9 and 15. However, it is interesting to note that these differences, evident at the local level of analysis (e.g., Figures 12 and 13), were not seen in the more molar analyses (Figure 10): No differences were evident in terms of the control exerted by the previous-component reinforcer ratios.

Although it was the case that behavior followed individual reinforcers more closely in Conditions 9 and 15, this does not eliminate the possibility of some molar contingencies controlling behavior. If the reinforcer ratios arranged in these conditions (either 1:27 or 27:1) are considered at a molar level and the sequential position of reinforcers in a component is ignored, a single reinforcer on one alternative was likely to be followed by more reinforcers on that alternative. However, at a more local level, a single right-alternative reinforcer following more than one successive left-alternative reinforcer was not a good predictor of further right-alternative reinforcers because the component in effect was most likely to have arranged 27:1 in favor of the left alternative. Thus, the apparent local control of responding evident in the size of the disconfirmations could be due, to some extent, to molar contingencies.

Another possibility is exemplified by the approach taken by the linear operator model (Bush & Mosteller, 1955). Its basic assumption is that reinforcement increases response probability by a constant proportion of the difference between the current probability and the maximum probability. Nonreinforcement similarly decreases response probability by a constant proportion of the difference between the current probability and the minimum. Thus, a linear operator approach predicts decreasing marginal changes in response ratios with successive confirming reinforcers. Equally, it predicts that the effects of disconfirming reinforcers on the response ratio will increase following longer sequences of confirming reinforcers. Although a quantitative fit of the model to the present data was not assessed, both the predictions are qualitatively consistent with the results shown in Figures 6 and 12. That is, the effect of a confirming reinforcer on subsequent preference decreased with the sequential position of the reinforcer. In contrast, the effect of a disconfirming reinforcer on subsequent preference following a sequence of successive confirming reinforcers increased with the sequential position of the reinforcer.

A reasonable way to view performance under these conditions may be in terms of dual control by both local and more molar contingencies (e.g., Davis & Staddon, 1990). There is evidence across conditions in the present experiment for the presence of both local and molar control. The mechanism by which these two sources of control interact is unclear at this stage, and further empirical investigation is required. However, the data suggest increasing molar control as the range of reinforcer ratios was decreased (Experiment 1) and as the less extreme components were added in Experiment 2. The relative frequency of various sequences of reinforcers was clearly important, and it may be that performance can be reasonably approximated on the basis of the relative probabilities of sequences of just two or three successive reinforcers being obtained on the same alternative. Support for this can be seen in Conditions 9 and 15, which showed that the sequential position of disconfirming reinforcers is somewhat less important than might have been expected. It is, however, difficult to envisage these results being driven by a

mechanism that is solely local or solely molar in nature.

As Davison and Baum (2000) pointed out, the present procedure allows the investigation of behavior at a resolution not previously possible. The regularity in the local effects of individual reinforcers on behavior, and the regularity of changes in these effects with experimental manipulations, are compelling. The data allow analyses at numerous levels, and it remains possible that regularities might also be evident at even more local levels of analysis.

REFERENCES

- Alsop, B., & Elliffe, D. (1988). Concurrent-schedule performance: Effects of relative and overall reinforcer rates. *Journal of the Experimental Analysis of Behavior*, *49*, 21–36.
- Bailey, J. T., & Mazur, J. E. (1990). Choice behavior in transition: Development of preference for the higher probability of reinforcement. *Journal of the Experimental Analysis of Behavior*, *53*, 409–422.
- Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, *22*, 231–242.
- Belke, T. W., & Heyman, G. M. (1994). Increasing and signaling background reinforcement: Effect on the foreground response–reinforcer relation. *Journal of the Experimental Analysis of Behavior*, *61*, 65–81.
- Bush, R. R., & Mosteller, F. (1955). A mathematical model of simple learning. *Psychological Review*, *58*, 313–323.
- Davis, D. G. S., & Staddon, J. E. R. (1990). Memory for reward in probabilistic choice: Markovian and non-Markovian properties. *Behaviour*, *114*, 37–64.
- 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. C., & Hunter, I. W. (1979). Concurrent schedules: Undermatching and control by previous experimental conditions. *Journal of the Experimental Analysis of Behavior*, *32*, 233–244.
- Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum.
- Dreyfus, L. R. (1991). Local shifts in relative reinforcement rate and time allocation on concurrent schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, *17*, 486–502.
- Elliffe, D., & Alsop, B. (1996). Concurrent choice: Effects of overall reinforcer rate and the temporal distribution of reinforcers. *Journal of the Experimental Analysis of Behavior*, *65*, 445–463.
- Ferguson, G. A. (1966). *Statistical analysis in psychology and education*. New York: McGraw-Hill.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, *4*, 267–272.
- 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.
- Jones, B. M., & Davison, M. (1996). Residence time and choice in concurrent foraging schedules. *Journal of the Experimental Analysis of Behavior*, *65*, 423–444.
- Lobb, B., & Davison, M. C. (1975). Performance in concurrent interval schedules: A systematic replication. *Journal of the Experimental Analysis of Behavior*, *24*, 191–197.
- Mazur, J. E. (1992). Choice behavior in transition: Development of preference with ratio and interval schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, *18*, 364–378.
- Mazur, J. E. (1995). Development of preference and spontaneous recovery in choice behavior with concurrent variable-interval schedules. *Animal Learning & Behavior*, *23*, 93–103.
- Mazur, J. E. (1996). Past experience, recency, and spontaneous recovery in choice behavior. *Animal Learning & Behavior*, *24*, 1–10.
- Mazur, J. E. (1997). Effects of rate of reinforcement and rate of change on choice behavior in transition. *Quarterly Journal of Experimental Psychology*, *50B*, 111–128.
- Mazur, J. E., & Ratti, T. A. (1991). Choice behavior in transition: Development of preference in a free-operant procedure. *Animal Learning & Behavior*, *19*, 241–248.
- Schofield, G., & Davison, M. (1997). Nonstable concurrent choice in pigeons. *Journal of the Experimental Analysis of Behavior*, *68*, 219–232.
- Shettleworth, S. J., & Plowright, C. M. S. (1992). How pigeons estimate rates of prey encounter. *Journal of Experimental Psychology: Animal Behavior Processes*, *18*, 219–235.
- Staddon, J. E. R. (1968). Spaced responding and choice: A preliminary analysis. *Journal of the Experimental Analysis of Behavior*, *11*, 669–682.
- Stubbs, D. A., & Pliskoff, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, *12*, 887–895.
- Taylor, R., & Davison, M. (1983). Sensitivity to reinforcement in concurrent arithmetic and exponential schedules. *Journal of the Experimental Analysis of Behavior*, *39*, 191–198.

Received July 24, 2000

Final acceptance February 13, 2001