

CHOICE AS A FUNCTION OF LOCAL VERSUS MOLAR REINFORCEMENT CONTINGENCIES

BEN A. WILLIAMS

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Rats were trained on a discrete-trial probability learning task. In Experiment 1, the molar reinforcement probabilities for the two response alternatives were equal, and the local contingencies of reinforcement differentially reinforced a win-stay, lose-shift response pattern. The win-stay portion was learned substantially more easily and appeared from the outset of training, suggesting that its occurrence did not depend upon discrimination of the local contingencies but rather only upon simple strengthening effects of individual reinforcements. Control by both types of local contingencies decreased with increases in the intertrial interval, although some control remained with intertrial intervals as long as 30 s. In Experiment 2, the local contingencies always favored win-shift and lose-shift response patterns but were asymmetrical for the two responses, causing the molar reinforcement rates for the two responses to differ. Some learning of the alternation pattern occurred with short intertrial intervals, although win-stay behavior occurred for some subjects. The local reinforcement contingencies were discriminated poorly with longer intertrial intervals. In the absence of control by the local contingencies, choice proportion was determined by the molar contingencies, as indicated by high exponent values for the generalized matching law with long intertrial intervals, and lower values with short intertrial intervals. The results show that when molar contingencies of reinforcement and local contingencies are in opposition, both may have independent roles. Control by molar contingencies cannot generally be explained by local contingencies.

Key words: choice, probability learning, concurrent schedules, matching law, momentary maximizing, local reinforcement contingencies, intertrial interval, lever press, rats

A recurrent theme of theoretical controversy has been the role of local contingencies of reinforcement as determinants of choice. Proponents of momentary maximizing theory (Hinson & Staddon, 1983; Shimp, 1979, 1984; Silberberg, Hamilton, Ziriax, & Casey, 1978) have argued that molar measures of choice are aggregates of responses that individually are determined by the subject choosing the alternative with the highest momentary reinforcement probability. Consequently, the regularities seen at the molar level, such as the matching law, need not reflect fundamental behavioral processes, but instead may apply only to the extent that the individual events governed by the local contingencies combine to produce the molar relations.

There can be little question that behavior can be differentiated by differences in local reinforcement probabilities. Control of the

temporal structure of behavior with schedules (e.g., differential reinforcement of low response rates) provides perhaps the most salient example (e.g., Galbicka & Platt, 1986), but the importance of local contingencies within standard choice procedures themselves is acknowledged by the recognition that a changeover delay (COD) is often necessary to obtain orderly molar relationships (Catania, 1966). Because the highest probability of reinforcement would otherwise occur for changeover responses, a COD is necessary to prevent simple response alternation from dominating. The issue, therefore, is not whether local reinforcement contingencies can determine behavior, but how such contingencies are involved in behavioral regularities at the molar level.

Herrnstein and Loveland (1975) have argued that momentary maximizing theories are not antagonistic to molar theories such as the matching law, but instead are complementary. They apply whenever the available stimuli correlated with differences in reinforcement contingencies (e.g., time since the last response) gain stimulus control over behavior. To the extent that stimulus control fails to occur, these differences in reinforcement con-

This research was supported by NIMH Grant 42197 and NSF Grant BNS 9010323 to the University of California, San Diego. I thank Bertram Ploog and Roger Dunn for technical assistance. Reprint requests should be addressed to the author, Department of Psychology, UCSD, La Jolla, California 92093-0109.

tingencies are not discriminated and the molar matching law then applies. To the extent that stimulus control does occur, new behavioral units are formed, and these new behavioral units themselves obey the matching law. The implication of this perspective is that the role of local reinforcement contingencies may be more or less important in determining choice behavior, depending upon the salience of the stimuli correlated with the local contingencies and perhaps with the differences in the local contingencies themselves.

Evidence that local reinforcement contingencies may or may not control behavior, depending upon the contingencies involved, was provided by Williams (1972). Pigeons were presented with a "win-stay, lose-shift" reinforcement contingency in which the overall probability of reinforcement for two response alternatives was equal, but the momentary contingencies varied with the outcome of the preceding response. When the probability of reinforcement for repeating a just-reinforced response (win-stay) was .65 versus a probability of .35 for switching to the alternative (win-shift), all subjects failed to acquire the win-stay behavior; instead, they developed a position habit. But when the probability of reinforcement for win-stay was increased to .80 (vs. .20 for win-shift), the win-stay behavior was learned. Thus, the establishment of stimulus control by the outcome of the preceding trial depended critically on the size of the difference in the reinforcement probabilities.

Shimp (1976) conducted a procedure similar to that of Williams (1972) and compared the acquisition of a win-stay behavior pattern with a win-shift behavior pattern involving similar differences in local reinforcement contingencies. Both behavior patterns were learned (with a difference in probability of .8 vs. .2, thus replicating the earlier results of Williams, 1972), but the win-shift behavior pattern developed more slowly. Moreover, when the intertrial interval (ITI) was varied from 0.5 to 4.0 s, adherence to both behavior patterns was substantially reduced with the longer ITIs. The apparent implication is that the establishment and maintenance of stimulus control by local reinforcement contingencies depend both on the type of behavior pattern specified by the local contingencies and on the temporal characteristics of the situation. The more general implication is that the role of local reinforce-

ment contingencies in a particular situation cannot be specified beforehand because there are a large number of procedural features that determine when they may, or may not, be discriminated.

When local reinforcement contingencies have been pitted in opposition to molar reinforcement contingencies, separate roles appear to be played by both. Hiraoka (1984) trained rats on a discrete-trial choice procedure in which the overall probabilities of reinforcement for responses to two levers were .61 versus .39, which, according to the matching law, should produce exclusive preference for the lever with the higher probability (because the probability schedules were equivalent to a concurrent random-ratio random-ratio schedule; cf. Herrnstein & Loveland, 1975). This prediction was confirmed for subjects that had no differential local reinforcement contingencies, but was not upheld for two other groups that had different local contingencies embedded within the overall probability structure. For these subjects, 80% of the assigned reinforcers followed a win-stay, lose-shift pattern for one group and a win-shift, lose-stay pattern for a second group. For both of these groups, complete control by the local contingencies would produce an overall choice probability of .50 for each response alternative. The obtained choice proportions (means over subjects) were .82 for the win-stay, lose-shift group and .87 for the win-shift, lose-stay group. Thus, the overall choice proportion was between the predictions of molar matching theory and momentary maximizing theory, suggesting that both types of contingencies were involved, but they were clearly closer to the predictions of the molar theory. Moreover, the local structure of behavior also reflected a similar pattern of joint control. For example, for the win-stay, lose-shift group, the probability of a choice of the majority lever (the lever associated with the .61 probability) after a nonreinforced response to that lever was .59, substantially below the mean probability of choice regardless of the preceding trial outcome, but also much higher than the 0 probability predicted by the local reinforcement contingencies. Similarly for the win-shift, lose-stay group, choice of the majority lever after a reinforced response to that lever was .83, only slightly below the overall choice probability, and much higher than the 0 probability predicted by the local contingencies.

A similar pattern of joint control is evident

in the subsequent results of Zeiler (1987, Experiment 2). He trained pigeons on a concurrent schedule in which different probabilities of reinforcement were assigned to one of two response alternatives. A correction procedure was used, so that the obtained relative reinforcement frequencies equaled the ratio of the reinforcement probabilities. Given a reinforcer on the preceding trial, momentary maximizing theory predicts that the subject should always choose the alternative with the higher scheduled probability. In fact, the choice allocation after a reinforcer on the preceding trial approximately matched the obtained number of reinforcers for each alternative, which implies that the local reinforcement contingencies were not discriminated. Zeiler also recorded the results after nonreinforcement on the preceding trial, when, because of the correction procedure, the probability of reinforcement for repeating the preceding response was 0 and the probability of reinforcement for switching to the other alternative was 1.0. All subjects responded in accordance with the lose-shift contingencies when their preceding response had been to the alternative with the lower overall reinforcement probability, thus demonstrating that some portion of the local contingencies had been discriminated. But they also systematically decreased their adherence to the lose-shift pattern whenever their preceding response had been to the alternative with the higher overall reinforcement probability. That is, their behavior was systematically biased away from the lose-shift response pattern by the overall reinforcement probability of the alternative chosen on the last trial. The apparent implication of Zeiler's result is that the overall reinforcement probability controlled behavior even when the local reinforcement contingencies were discriminated, by providing a response bias that competed with the control by the local contingencies.

The results of Hiraoka (1984) and Zeiler (1987) suggest that control by local reinforcement contingencies and by molar reinforcement contingencies are in competition. The issue posed is how such joint control is to be incorporated into a theory of choice behavior. Proponents of momentary maximizing theory (e.g., Silberberg & Williams, 1974) have acknowledged that control by local reinforcement contingencies may fail to occur but have argued that such failures result from errors of "memory" (i.e., stimulus control by the events on

the preceding trial(s)). According to such an account, behavioral regularities should be less evident in conditions that attenuate control by the local contingencies, but the regularities (e.g., matching), when they do occur, must be due to the local reinforcement contingencies. As noted by Zeiler (1987), however, the systematic biases exerted by the molar reinforcement contingencies seem not easily explainable by such failures of memory. Instead, such data seem to demand some type of hybrid theory in which local contingencies and molar contingencies exert separable effects.

The present study is a further attempt to define how local and molar contingencies interact in choice procedures. The general strategy was first to train rats on an explicit discrimination involving local reinforcement contingencies, in this case a win-stay, lose-shift response pattern. After control by the local contingencies was established, an assessment was made of how such control was affected by the ITI. Such control presumably was based on memory of the events from the preceding trial, which should become less effective with longer times since the preceding trial. After the effects of ITI were established in Experiment 1, Experiment 2 presented a choice procedure in which the two alternatives differed in their overall reinforcement probability but with differences in local probability of reinforcement still in effect, much like the procedure used by Hiraoka (1984). The issue was then how choice behavior was affected by variations in the ITI. To the extent that any molar regularities in choice depend upon discrimination of the local contingencies, these should be less evident as the ITI is lengthened. But if control by the local contingencies and control by the molar contingencies are in competition, and the control by the local contingencies was weakened by the longer ITI, regularities at the molar level should be more evident with longer ITIs.

EXPERIMENT 1

METHOD

Subjects

Four experimentally naive male Sprague-Dawley albino rats, approximately 2 months of age at the start of training, served as subjects. Rats were housed in individual cages with water freely available. Food deprivation was

maintained by allowing 1 hr access to free food after the end of the experimental sessions.

Apparatus

A standard rat conditioning chamber (30 cm wide by 22 cm deep by 20 cm high) was enclosed within a larger sound-insulated wooden box equipped with a ventilating fan, noise from which served to mask extraneous sounds. The ceiling and three walls of the interior chamber were made of Plexiglas, the floor was a standard wire grid, and the front panel was made of sheet metal. Mounted on the front panel, 7.6 cm above the floor, were two retractable levers (BRS/LVE Model RRL-015) that protruded 1.8 cm into the chamber when activated and that required a minimum force of 0.3 N to close the microswitch. The levers were 2.5 cm wide, 1 cm thick, and spaced 7.5 cm apart, measured side to side. A single jewelled panel light was mounted directly between and 5 cm above the levers. Except for this light, the chamber was dark. Also directly between the two levers but just above the floor was mounted a pellet chute into which a single 45-mg Noyes pellet (improved Formula A) could be delivered as the reinforcer from a Gerbrands pellet dispenser.

Procedure

All subjects were first trained to press each of the two levers separately with the other lever withdrawn from the chamber. During this training, each reinforced lever press was followed by a 5-s ITI in which the panel light was extinguished and the lever was withdrawn from the chamber. During the first session, 100 reinforced trials were presented in which only one response was required per trial. During the second session, the number of responses per trial required for reinforcement was increased to five (fixed ratio [FR] 5), and again 100 reinforced trials were presented. This training was presented separately for each lever, with the order of presentation counterbalanced across subjects.

After this pretraining, each trial began with the panel light illuminated and both levers presented. Whichever lever first received five responses on a given trial was designated as the choice for that trial, regardless of the number of responses that occurred to the other lever (which could be zero to four). After five re-

sponses to a given lever had occurred, the food reinforcer was delivered if scheduled, the panel light was extinguished, and both levers were withdrawn from the chamber. Following a 5-s ITI, the levers were returned for the next choice opportunity.

The reinforcement contingencies for a given trial were assigned according to a win-stay, lose-shift response pattern. Given a reinforced choice on the preceding trial, the probability of reinforcement for repeating the same choice was .8, and the probability of reinforcement for a switch to the other alternative was .2. Given that a choice on the preceding trial was nonreinforced, the probability of reinforcement for repeating the same choice was .2, and the probability of reinforcement for a switch to the other alternative was .8. The reinforcement probability for the first choice of the session was .5 for each lever. A total of 24 sessions, each with 100 trials per session, was presented with these contingencies.

During the second phase of training, the ITI was varied. Four sessions were presented with an ITI of 10 s, followed by four more sessions with an ITI of 30 s. Four sessions were then presented with the original 5-s ITI reinstated. Next, the ITI was varied within a session. Four different ITIs, 5, 15, 30, and 60 s, were randomly interspersed within a session. A total of 16 sessions was presented. The data from this second phase of training will not be presented because substantial position habits developed during the presentation of the longer ITIs.

During the third phase, the original 5-s ITI was reinstated for five sessions with the original contingencies. Then, the contingencies for the lose-shift portion of the behavior were changed. The probability of reinforcement for repeating a choice not reinforced on the preceding trial was changed from .2 to 0, and the probability of reinforcement for shifting to the alternative was increased from .8 to 1.0. Ten sessions were conducted under this procedure.

During Phase 4, the ITI was again varied; five consecutive sessions were presented with each of four different ITIs of 5, 15, 30, and 60 s. Throughout this training, the same ITI was used within a given session. During Phase 5, the ITI was again varied within a session; the four different ITIs were randomly interspersed within a session for each of 20 sessions.

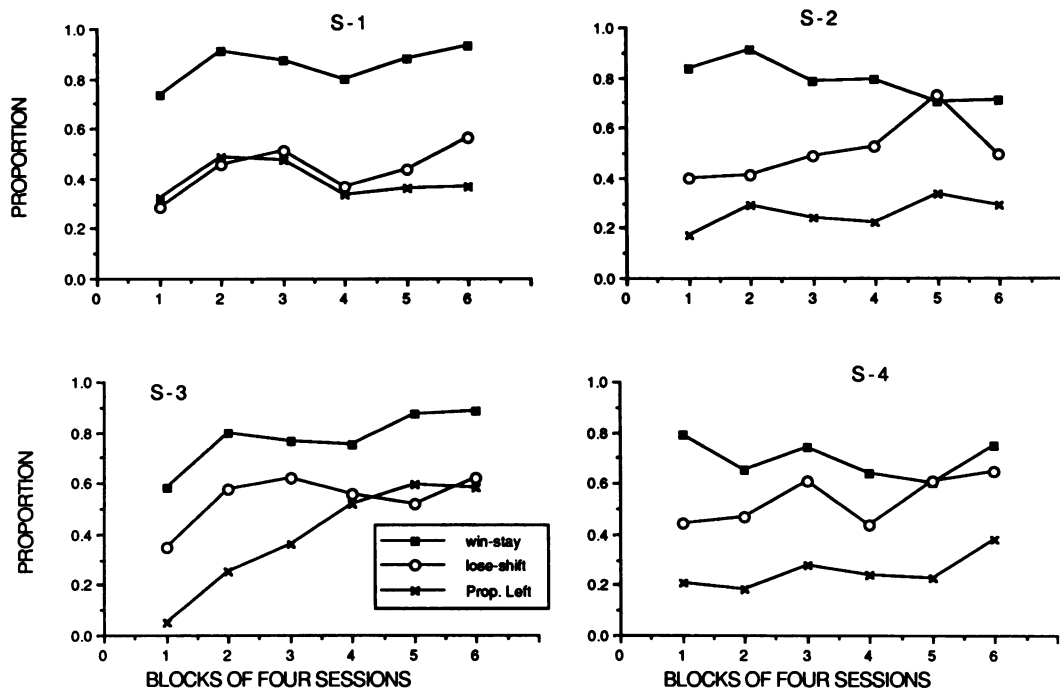


Fig. 1. Acquisition of the win-stay, lose-shift response pattern in Experiment 1, subdivided into the win-stay and lose-shift components. Also shown is the positional bias summed over all responses.

RESULTS

Figure 1 shows the results of the initial exposure to the win-stay, lose-shift contingencies with an ITI of 5 s. Shown separately are the behavior after reinforcement on the preceding trial (win-stay) and behavior after no reinforcement on the preceding trial (lose-shift). Also shown is an index of position bias calculated over all trials (i.e., the percentage of choices to a particular lever regardless of the local contingencies). Adherence to the win-stay and lose-shift response pattern was calculated by averaging the results for the two different response levers, even though the number of responses to each lever could be very different because of position biases. For example, if 80 responses occurred to the left lever and 20 to the right lever, and the adherence to win-stay for responses to the left was 100% and adherence to win-stay for responses to the right was 50%, the overall adherence to the win-stay pattern would be 75%.

Figure 1 shows that adherence to the win-stay and lose-shift portions of the response pattern was very different. For win-stay, all 4

subjects showed substantial adherence from the outset of training, with the proportion of choices to the lever reinforced on the just-preceding trial being .73, .84, .58, and .79 for Subjects S-1 through S-4, respectively, during the first block of training. For S-3, there was some further increase in win-stay adherence with continued training, whereas for S-2 adherence decreased over training. There was no consistent trend over sessions for S-1 and S-4. In contrast, for the lose-shift portion of the response pattern, the initial adherence was very low (.29, .40, .35, and .44 during the first block for the individual subjects) and generally increased over training. Even with that increase, adherence to lose-shift remained below that of adherence to win-stay for 3 of the 4 subjects.

Initial presentation of different ITIs resulted in a decline in both win-stay and lose-shift behavior with longer ITIs and also an increase in positional responding independent of the local contingencies. Because obtained reinforcement rates for the left versus right lever could vary substantially with the positional habit, different ITIs produced different relative rates of reinforcement for the two le-

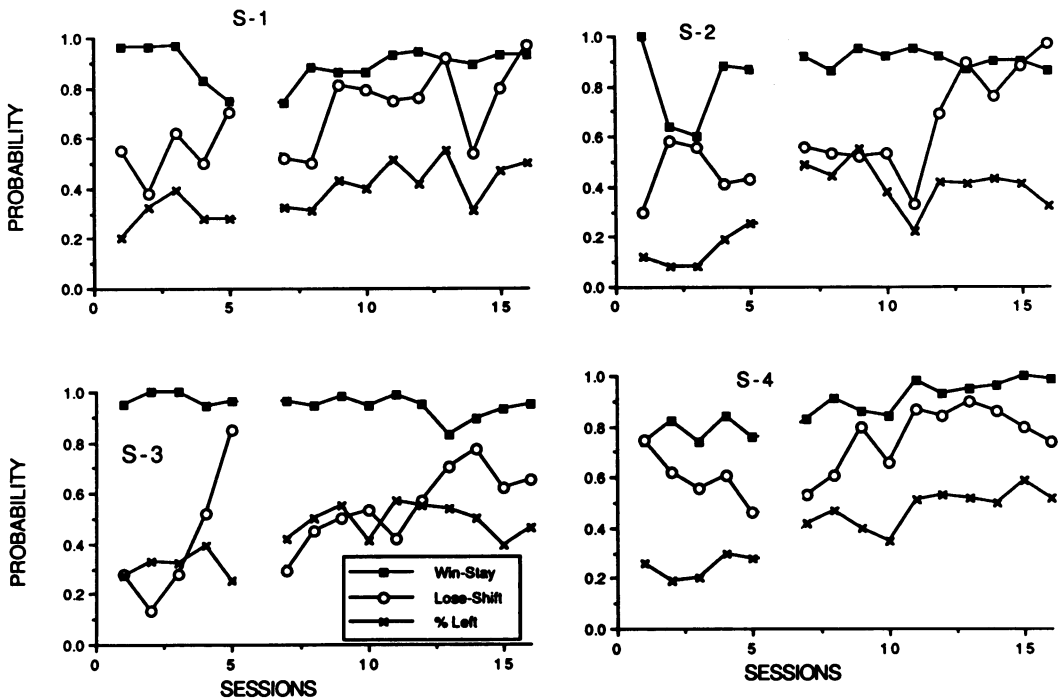


Fig. 2. Effects of changing the lose-shift reinforcement probability from .2 to 0. The first segment of each panel shows the results during retraining with the .2 probability with the 5-s ITI. The second segment shows the results after the change in the lose-shift reinforcement probability.

vers, thus hindering any interpretation of the differences in adherence to the local contingencies. Consequently, the procedure was changed to one that decreased the likelihood of large differences in the obtained number of reinforcers on the two levers. This was accomplished by increasing the strength of the lose-shift contingency. Instead of reinforcement probabilities that were symmetrical with the win-stay behavior (.8 vs. .2), the probability of reinforcement was 0 for repeating a choice nonreinforced on the preceding trial, and the probability for shifting to the alternative was increased to 1.0. Thus, sustained responding to a single position resulted in nonreinforcement until the subject alternated to the other position.

Figure 2 shows the results of this change in procedure. The first portion of the graph shows the results with the original contingency for five sessions of retraining with the 5-s ITI after the exposure to the different ITIs. The general pattern is like that seen in Figure 1: Win-stay behavior was generally superior to lose-shift behavior, although there was substantial variability across sessions. The second portion of

the graph shows the behavior after the change in probabilities; for all subjects, adherence to the lose-shift behavior substantially improved. For S-1 and S-2, the level of lose-shift behavior approximated the level of adherence that occurred to the win-stay behavior. But for S-3 and S-4, adherence to the lose-shift behavior remained at a level below that for win-stay.

Figure 3 shows the results of the variation in ITI after the change in lose-shift contingency. The left panels show the results when the ITI was varied between sessions, and the right panels show the results when the ITI was varied within sessions. For the former, only the data from the last three sessions of each exposure are included, because there was a small improvement in lose-stay performance over the first three sessions of training. However, there were no systematic changes during the last three sessions, so those results are aggregated. All sessions from the within-session form of testing are included because there was no systematic change over testing with that procedure. In general, the results were similar regardless of whether the ITIs were varied between or within sessions, because adherence

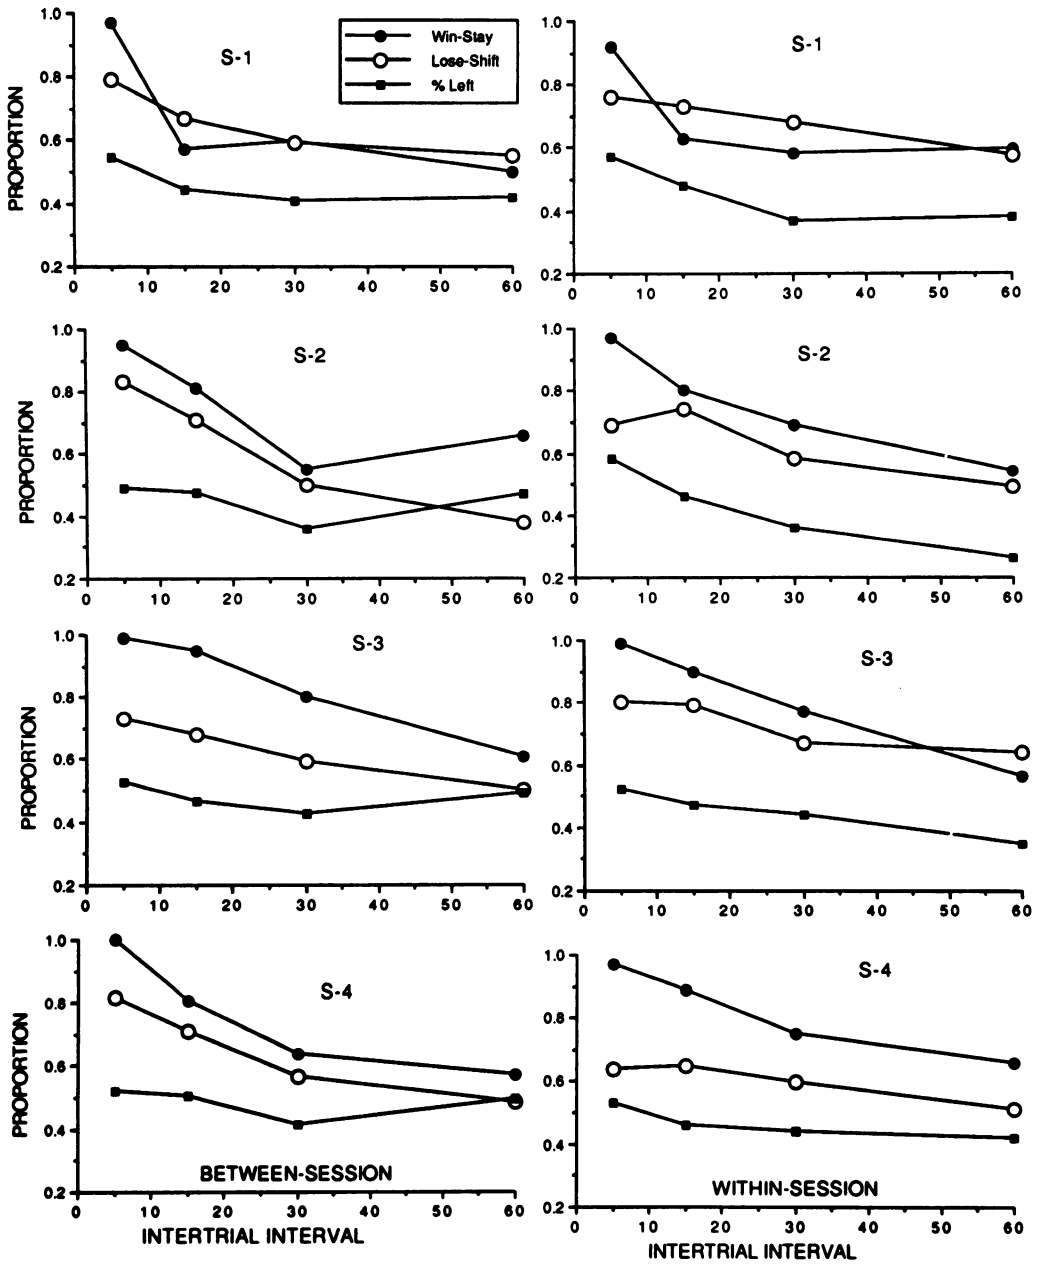


Fig. 3. Effects of ITI variation in Experiment 1 after the change in lose-shift probability to 0. Left panels show the results when ITI was varied between sessions; right panels show the results when ITI was varied within sessions.

to both the win-stay and lose-shift behavior patterns declined regularly with increases in the ITI. For 3 of the 4 subjects, win-stay behavior was still somewhat more accurate than lose-shift behavior at all ITIs, but these differences were much smaller than before the change in lose-shift contingencies. For the re-

maining subject (S-1), there was no systematic difference between the two types of behavior except at the 5-s ITI. The results (Figure 3) were subjected to a three-factor ANOVA (type of test \times response pattern \times ITI): The effect of test type was not significant, $F(1, 3) = 1.59$, whereas the proportion adherence to the win-

stay versus lose-shift response patterns approached significance, $F(1, 3) = 9.37$, $.06 > p > .05$. The effect of ITI was significant, $F(3, 9) = 50.1$, $p < .01$, as was the interaction between response pattern and ITI, $F(3, 9) = 4.01$, $p < .05$. None of the other effects approached significance.

Of some interest is the degree of adherence to the win-stay and lose-shift contingencies with the longer ITIs used (30 s and 60 s). Averaged over subjects, the adherence to the win-stay pattern was 64% and 57% for the 30-s and 60-s ITIs for the between-session variation of ITI and 70% and 59% for the within-session test. The corresponding values for the lose-shift behavior were 57% and 49% for the between-session test and 63% and 55% for the within-session test. Thus, some degree of control by the local contingencies remained even with the longest ITIs.

DISCUSSION

Experiment 1 demonstrated that rats respond at least partially in accordance with local contingencies of reinforcement even with substantial ITIs separating their successive choices. For the win-stay behavior, for example, some adherence to the local contingencies occurred even with a 60-s ITI. The decline in discrimination accuracy was substantially lower in the present study than that which occurred in the study by Shimp (1976), who reported a major disruption in performance when the ITIs were increased from 1 to 4 s. It should be noted that Shimp used pigeons as subjects, whereas the present study involved rats, and it is plausible that these different subjects vary substantially in their sensitivity to delayed stimulus control. The present use of an FR 5 response requirement is also known to facilitate delayed stimulus control relative to an FR 1 requirement (Williams, 1971a, 1971b, 1972).

Although all subjects eventually came under the control of both the win-stay and lose-shift contingencies, there was a major difference in the speed with which such control occurred. Adherence to the win-stay pattern was evident from the outset of training, and improved relatively little thereafter. In contrast, lose-shift behavior was slow to be established, and generally continued substantially below the level of win-stay performance throughout training. Increasing the strength of the lose-shift con-

tingency did eliminate some of this difference later in training, when the probability of reinforcement for repeating a response that had not been reinforced on the preceding trial was reduced to 0.

The fact that win-stay behavior occurred from the outset of training suggests that it was not necessarily under stimulus control of the response consequence of the preceding trial. That is, the subject need not have discriminated that the probability of reinforcement for repeating a just-reinforced trial was higher than the probability of reinforcement for shifting to the other response. Evidence that repetition of just-reinforced responses does not depend upon differential probabilities of reinforcement is provided by Morgan (1974) and Evenden and Robbins (1984), both of whom presented rats with a probabilistic schedule in which the probabilities of reinforcement for two alternatives were equal after a reinforced response. Despite there being no differential consequence for choosing one or the other response, both studies showed very substantial perseveration (from 80% to 90%) to the response reinforced on the just-preceding trial. The present level of win-stay behavior was similar to that obtained in their studies, which suggests that the present results were also due to the "simple strengthening" effects of reinforcement. Given that interpretation, the present results are noteworthy in showing that such perseveration effects are time dependent, and decay substantially over the range from 5 to 60 s.

The monotonic decline in win-stay performance with increasing ITIs is of interest because it contradicts the results reported by Williams (1983). In that study, pigeons were trained on a similar win-stay, lose-shift response pattern and ITI was varied in blocks of sessions over a range of 3 to 30 s. Lose-shift performance declined regularly with increasing ITI and was at a chance level with 10-s ITI. But win-stay behavior was nonmonotonic, declining from 3 to 10 s and then improving with a further increase to 30 s. No evidence for such a nonmonotonic function was obtained in the present study. Why the conflict in results occurred is unclear, although it should again be noted that the present study used rats as subjects, whereas Williams (1983) used pigeons.

The present results also appear to conflict

Table 1

Experimental conditions for the two groups in Experiment 2. The lever serving as the majority response was counterbalanced across subjects.

	Order	ITI	Prob. (Rf/shift) to		Prob. (Rf/stay) on		Sessions
			Minority	Majority	Minority	Majority	
Group 1	1	5	.4	.6	.1	.2	40
	2	5	.25	.6	.1	.2	32
	3	30	.25	.6	.1	.2	32
Group 2	1	30	.4	.6	.1	.2	40
	2	30	.25	.6	.1	.2	32
	3	5	.25	.6	.1	.2	32

with some previous studies showing that rats are more likely to engage in win-shift and lose-shift behavior than in win-stay behavior. This general conclusion has been supported by analyses of rats foraging in mazes (e.g., Olton, 1979), by foraging behavior of nectar-feeding birds (Cole, Hainsworth, Kamil, Mercier, & Wolf, 1982), and most recently by pigeons in an operant chamber (Plowright & Shettleworth, 1990). However, the predominance of win-stay behavior similar to that observed here was also reported by Shimp (1976) and by Evenden and Roberts (1984). Comparing these studies, the critical variable appears to be the spatial layout of the response alternatives. When widely separated alternatives are used, shift behavior is more likely; with closely spaced alternatives, like those in conventional operant chambers, stay behavior is more likely.

EXPERIMENT 2

The purpose of Experiment 1 was to establish how behavior under the control of local contingencies of reinforcement was affected by changes in the temporal parameters of the situation. Such control deteriorated with increasing ITIs, although some remained even with the longest ITIs. Nevertheless, it is clear that the ITI can serve to modulate the degree of control by the local contingencies, thus allowing an assay of how such control is involved in producing molar regularities of behavior, such as the matching law.

Experiment 2 continued the analysis of how ITI affects choice behavior, but with two major changes in procedure. Whereas the contingencies for the two responses (left vs. right) in Experiment 1 had been symmetrical, now they became asymmetrical, in that one response had

a higher probability of reinforcement than the other. The second major change was in the nature of the local contingencies. Whereas Experiment 1 had a win-stay, lose-shift contingency throughout, now the local contingencies favored shifting regardless of the outcome of the preceding trial. The local contingencies were thus like those of a concurrent variable-interval (VI VI) schedule in which changeover responses have the highest probability of reinforcement. The only remaining change was that the response requirement was reduced from an FR 5 to an FR 1, to be more similar to previous studies (e.g., Williams, 1985).

The issue for Experiment 2 was how variation in the ITI affected the relative degree of control by the local versus molar contingencies of reinforcement. Given the asymmetry in reinforcement probability for the two alternatives, control by the molar contingency of reinforcement predicts that preference should track the overall relative rate of reinforcement for the two responses. For example, given the contingencies used in Condition 1 (Table 1) matching predicts a choice proportion of the majority lever of .70, assuming the probability of a given choice is independent of the events of the preceding trial. To see why this intermediate level of preference is predicted, instead of exclusive preference for the majority lever, note that the distribution of stay versus shift behavior necessarily changes with changes in the overall choice proportion. As preference for the majority lever increases, a higher percentage of responses to that lever will be stay responses; this means that the overall probability of reinforcement for the majority lever will decrease the greater its choice proportion. Conversely, as the percentage of responding to the minority lever decreases, a high percentage

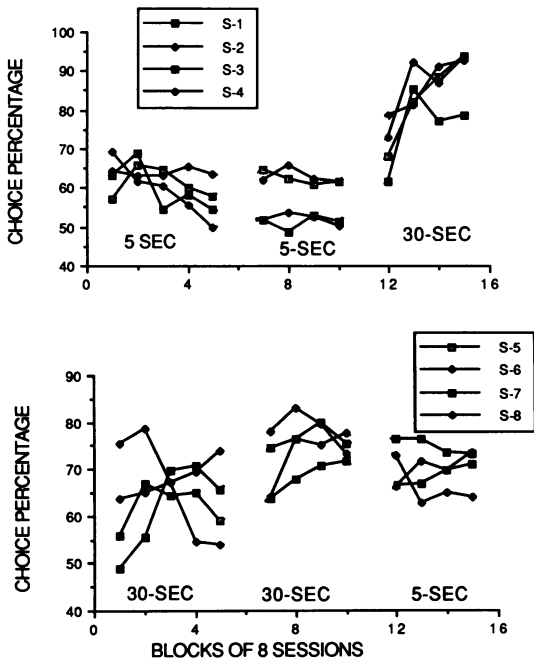


Fig. 4. Mean choice of the majority lever in Experiment 2 across the three conditions. Subjects in Group 1 are shown in the top panel; subjects in Group 2 are shown in the bottom panel.

of those responses become shift responses, which results in a higher probability of reinforcement for minority responses. Equilibrium between choice proportion and reinforcement probability then occurs with intermediate levels of preference in the same way as it does for concurrent VI VI schedules.

In contrast to this prediction for control by the molar contingencies, control by the local contingencies implies a choice proportion of .50, because shift behavior is always the response with the highest probability of reinforcement at any moment. At issue was whether different degrees of the two types of control would occur with different ITIs. According to momentary maximizing theory, control by the molar contingencies is derivative from control by the local contingencies; accordingly, longer ITIs should decrease both types of control. But if in fact the two types of control are in competition, one might expect control by the local contingencies with the shorter ITI, thus producing an overall choice proportion near indifference, but control by the molar contingencies with the longer ITI, thus producing an approximation to matching.

METHOD

Subjects and Apparatus

The subjects were 8 albino rats, maintained as in Experiment 1. Their experimental history consisted of lever-press acquisition using a delayed reinforcement contingency (see Williams, Preston, & de Kervor, 1990, Experiment 3, for details). The apparatus was also the same as Experiment 1.

Procedure

The rats were begun immediately on the sequence of conditions shown in Table 1. Subjects in Group 1 initially received training with a 5-s ITI; those in Group 2 received training with a 30-s ITI. The first two conditions differed only with respect to the probability of reinforcement for shifting to the minority alternative after a response to the majority lever. All other probabilities remained the same. The ITIs for the two groups were then reversed for the third condition. Within each group, 2 subjects were trained with left as the majority response and the other 2 were trained with right as the majority response. The number of sessions per condition is also shown in Table 1.

RESULTS

Figure 4 shows the development of the overall choice proportion across the three different conditions of the study. The top portion shows the results for Group 1, which first received the 5-s ITI; the bottom portion shows the result for Group 2, which first received the 30-s ITI. The first segment of each panel shows the results when the probability of reinforcement for a shift to the minority response was .4; the remaining two portions show the results when that probability was reduced to .25. The subjects in Group 1 began with a choice percentage in the range of 60% to 70% during the first two blocks of training, but this proportion then decreased toward the 50% to 60% range for 3 of the 4 subjects over the course of training (Figure 4). Little further change then occurred when the probability of reinforcement for a shift to the minority response was decreased (which also produced a decrease in its overall reinforcement probability). When the ITI was then increased to 30 s during the third condition, the choice proportions increased substantially for all subjects.

The results for Group 2 were notably dif-

ferent. The results during Phase 1 were quite variable across subjects, with 1 subject showing a decrease toward the 50% to 60% range over the course of training, another showing a systematic increase to the 70% to 80% range, and the remaining 2 in the middle. Then, during Phase 2, the behavior of the 4 subjects became much more similar, in the range of 70% to 80% for all subjects by the last block of training. When the ITI was reduced to 5 s during Phase 3, a slight decrease in the overall choice proportion occurred, but the terminal level was still substantially above that for the corresponding condition for Group 1 (compare the middle panel of the top portion of the figure). Thus, the data in Figure 4 suggest that the ITI, the local reinforcement probabilities, and the order of presentation were all involved in producing the differences that were evident.

For purposes of statistical analysis, the results from the first two conditions were subjected to a three-way mixed ANOVA. The between-group factor was the ITI; the within-group factors were the reinforcement probabilities and the blocks. In order to have an equal number of blocks for the two phases of training, the first block for Condition 1 was eliminated from the analysis. The effect of the ITI was significant, $F(1, 6) = 15.1, p < .01$, whereas the effect of probability of reinforcement did not quite achieve conventional levels of significance, $F(1, 6) = 5.55, .06 > p > .05$. The interaction between the ITI and probability was also significant, $F(1, 6) = 27.3, p < .01$. No other effects approached significance. Because of the significant interaction, further analysis was done to isolate the nature of the simple effects: The effect of the probability of reinforcement was not significant for Group 1, but was significant for Group 2; the effect of ITI was significant when the probability of reinforcement for shifting to the minority response was .25 but not when that probability was .40. It should be noted that the status of probability of reinforcement can be questioned in the above analysis because it was confounded with order of presentation.

A second two-factor within-subject ANOVA was performed comparing Phases 2 and 3 in which the probability of reinforcement remained constant while only the ITI was changed. Here the effect of ITI was significant, $F(1, 7) = 11.9, p < .01$, as was the effect of blocks, $F(3, 21) = 9.95, p < .05$, and the

interaction between ITI and blocks, $F(3, 21) = 6.05, p < .05$. A test of simple effects showed that the effect of blocks was significant for the 30-s ITI but not for the 5-s ITI, reflecting that choice of the majority lever systematically increased with continued training with the longer ITI but not with the shorter ITI.

Figures 5 and 6 show the same results, but divided according to the four possible trial types on the preceding trial. Figure 5 shows the results for Group 1; Figure 6 shows the results for Group 2. The different functions as a function of preceding trial type should be similar to the extent that the local contingencies were not discriminated. If discrimination did occur, choices of the majority lever should decrease after a choice of the majority lever on the preceding trial (filled data points) and should increase after choices of the minority lever (open data points). Figure 5 shows that considerable discrimination occurred for subjects initially trained with the 5-s ITI: S-2 clearly developed the pattern predicted by the local contingencies, and S-1 and S-3 approximated the pattern during the second condition. During the first condition, however, there was a greater tendency for all subjects to repeat a response to the majority lever when the preceding response had been reinforced than when it was not reinforced, an effect consistent with the reward-perseveration effect seen in Experiment 1. This tendency was very strong for S-4 for both Conditions 1 and 2, and also persisted strongly for S-1. For S-2 and S-3, it disappeared over the course of training. When the ITI was increased to 30 s during Condition 3, the discrimination of the local contingencies seen in Condition 2 substantially decreased and the functions became much more similar for all subjects.

Figure 6 shows the results for the subjects initially trained with the 30-s ITI. Relatively little discrimination of the local contingencies was evident during the first 20 to 25 sessions, but some did develop by the end of training on the first condition for all subjects. Specifically, the highest probability of choosing the majority lever generally occurred after a non-reinforced response to the minority lever on the preceding trial. This discrimination persisted for S-5 during Condition 2 but was less evident for the remaining 3 subjects. However, other evidence for the discrimination of the local contingencies is seen for S-7 and S-8, for

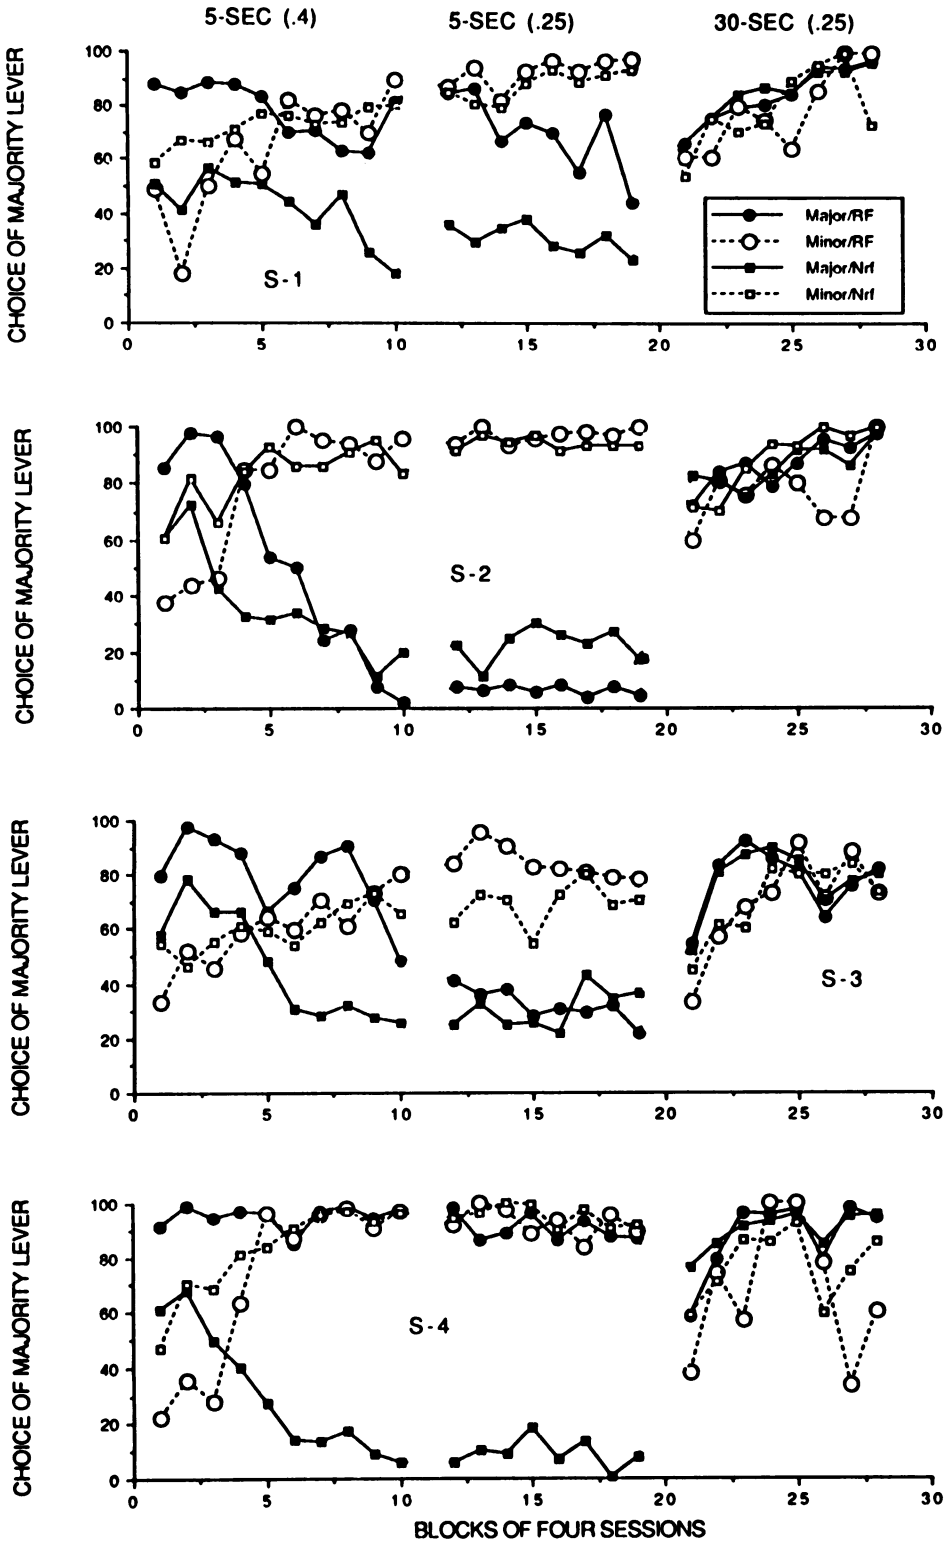


Fig. 5. Percentage of trials on which the majority lever was chosen, subdivided as a function of the events of the preceding trial. Results are from Group 1. Different segments correspond to the different conditions shown in Table 1.

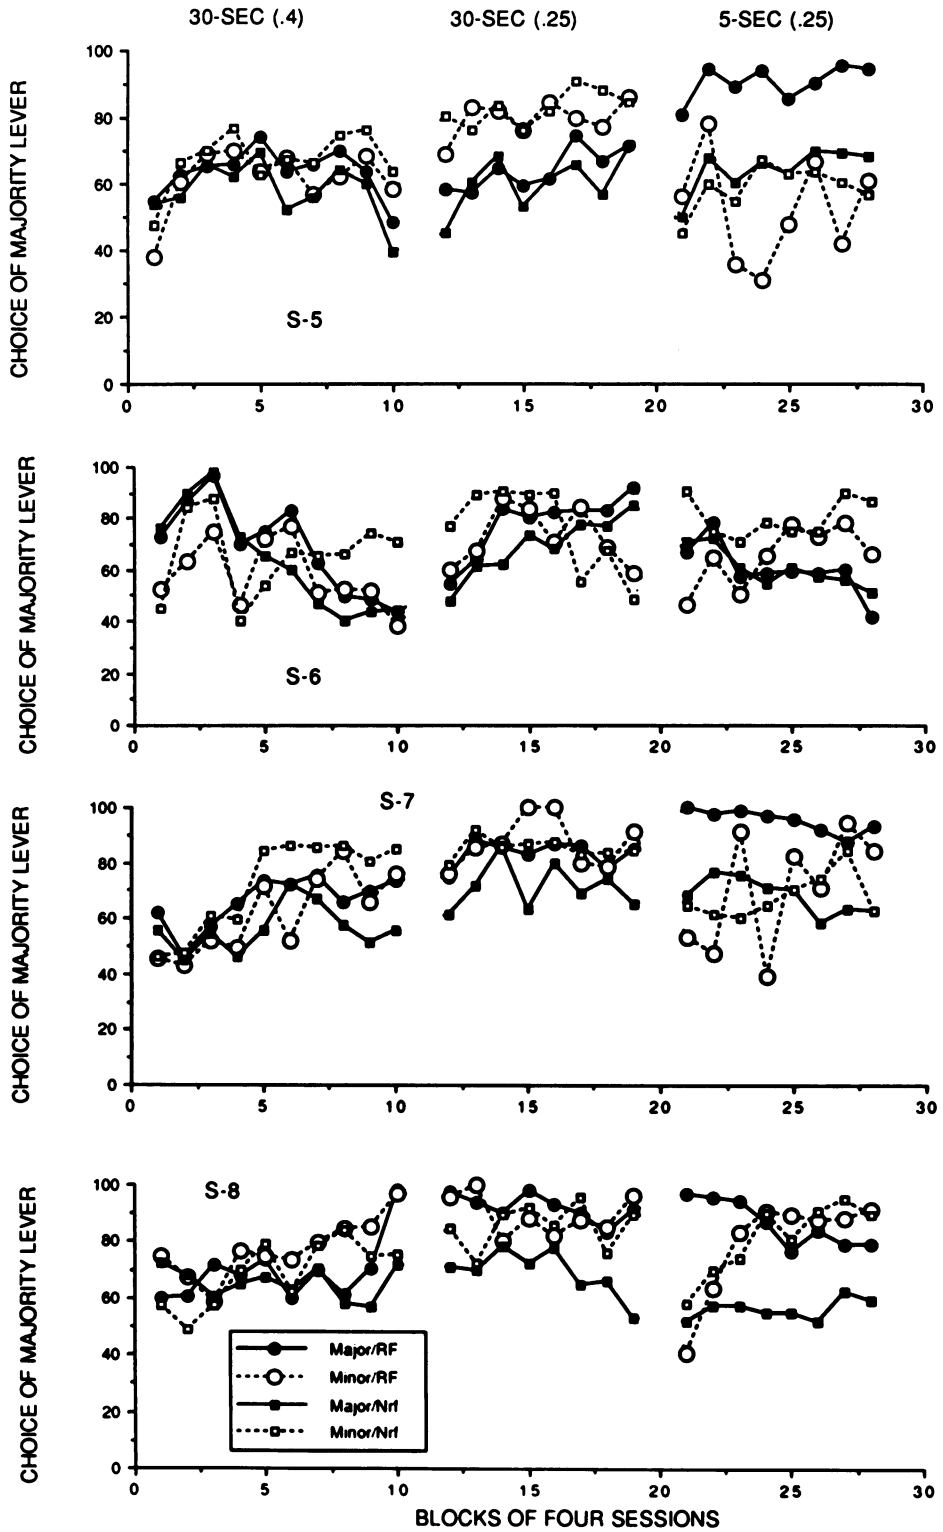


Fig. 6. Percentage of trials on which the majority lever was chosen, subdivided as a function of the events of the preceding trial. Results are from Group 2. Different segments correspond to the different conditions shown in Table 1.

Table 2

Number of reinforcers per session for the last eight sessions of each condition in Experiment 2.

Subject	5 s (.4)	5 s (.25)	30 s (.25)	30 s (.4)
1	39.5	34.9	23.5	
2	48.25	38.6	24.1	
3	37.0	36.5	26.4	
4	40.75	39.0	25.1	
<i>M</i> (Group 1)	41.4	37.25	24.8	
5		25.75	32.5	35.5
6		34.4	26.3	37.1
7		30.4	28.1	34.75
8		33.6	32.4	34.0
<i>M</i> (Group 2)		31.0	29.8	35.3

which the lowest probability of choosing the majority lever occurred when the preceding trial involved a nonreinforced response to the majority lever (filled squares).

The third panel of Figure 6 shows that greater discrimination of the local contingencies occurred after the change in ITI to 5 s. However, this discrimination was relatively poor in comparison to the corresponding condition in Figure 5 (middle panel), and for only 1 subject (S-6) was the pattern of responding in accord with the local contingencies (greater responding to the majority lever after a preceding response to the minority lever than after a response to the majority lever). For the remaining subjects, the probability of responding to the majority lever was high after a preceding reinforced response to the majority lever and was substantially lower after a nonreinforced response to the majority lever, despite the local contingencies after reinforcement and non-reinforcement being identical.

For the purposes of statistical analysis, the choice percentages as a function of the different preceding trial types were aggregated over the last eight sessions of each condition. Two separate ANOVAs were then performed. The first involved the data from Conditions 1 and 2, in which the ITI was a between-group factor, and probability of reinforcement (.4 vs. .25 for reinforcement of a shift to the minority lever) and preceding trial type were within-group factors. The only significant outcomes ($p < .05$) were the preceding trial type, $F(3, 18) = 11.74$, and the interaction between preceding trial type and ITI, $F(3, 18) = 5.37$. A test of simple effects for the interaction showed that

the effect of preceding trial type was significant for the 5-s ITI but was not significant for the 30-s ITI ($F < 1$).

The second ANOVA involved the data from Conditions 2 and 3, with order of presentation as a between-group factor and ITI and preceding trial type as within-group factors. The main effects of ITI, $F(1, 6) = 28.92$, and preceding trial type, $F(3, 18) = 7.65$, were significant, as was their interaction, $F(3, 18) = 5.62$. A test of the simple effects for the interaction term showed that preceding trial type was significant for the 5-s ITI but again was not significant for the 30-s ITI ($F < 1$). The interaction of order of presentation with ITI was also significant, $F(1, 6) = 16.25$, as was the triple interaction among order, ITI, and preceding trial type, $F(3, 18) = 4.93$. These last two interactions provide evidence that the order of presentation was important in determining how the local contingencies were discriminated. The major source of this interaction was that subjects trained with the 5-s ITI after training with the 30-s ITI discriminated the local contingencies less well than those trained initially with the 5-s ITI.

The degree of control over the local contingencies of reinforcement is also evident from the number of reinforcers acquired during each session. Table 2 shows the average number of reinforcers per session for each subject in each condition. That number should be larger to the extent that subjects discriminated the local contingencies of reinforcement (i.e., to the extent that they followed a momentary maximizing strategy). Complete adherence to the optimal strategy (response alternation) would produce an expected value of 50 reinforcers per session when the probability of reinforcement for a shift to the majority lever was .40 and an expected value of 42.5 when the probability of reinforcement for a shift to the minority lever was .25.

Table 2 shows that the number of obtained reinforcers was generally larger with the 5-s ITI than with the 30-s ITI with the corresponding probability of reinforcement. (Compare Columns 1 and 4 and Columns 2 and 3). Note that the exception was for Group 2 (Columns 2 and 3). That is, its obtained number of reinforcers was only slightly larger with the 5-s ITI, whereas the corresponding difference was much larger for Group 1. The difference between the two groups was presumably due

to the different orders of presentation, which also produced substantial differences in the pattern of behavior (Figures 5 and 6). Thus, after prior training with the .4 probability of reinforcement for shifting to the minority lever, Group 1 substantially continued its adherence to the momentary maximizing strategy; with no prior training with the 5-s ITI, Group 2 showed considerable less adherence to momentary maximizing. A corresponding effect of prior training is also seen with the 30-s ITI, in that Group 2 obtained substantially more reinforcers per session with the .25 probability condition, after prior training with the .4 probability, with the 30-s ITI, in comparison to the number of reinforcers obtained by Group 1, which had no prior history with the 30-s ITI. Thus, adherence to momentary maximizing apparently depended upon the maintenance of a response strategy that was ITI-specific and that required a larger differential in the local probabilities of reinforcement for its initial development.

The above analysis suggests that the control by the molar properties of the reinforcement schedules was greater with the 30-s ITI, in part because of the decrease in control by the local contingencies. A different method for assessing such molar control is to compare the overall choice percentages to that predicted by the matching law. As mentioned in the introduction to Experiment 2, specific response percentages are predicted by the two different schedules of reinforcement, assuming no discrimination of the local contingencies (i.e., the response proportion is independent of the events of the preceding trial) and no response bias. For the schedule used in Condition 1, this prediction was for a response percentage of 70%; for the schedule used in Condition 2, this prediction was 90%. Averaged over all subjects in a given condition with a given ITI, the mean response percentages (from the last eight sessions of each condition) for the 5-s ITI were 57% and 63% for the two reinforcement conditions; for the 30-s ITI, the obtained response percentages were 63% and 82%. Thus, for both ITI values there was undermatching, but considerably less undermatching with the longer ITI value. Note, however, that the assumption on which this analysis is based—that there was no discrimination of the local contingencies—was violated by all subjects in varying degrees. Such violation has the result of producing ob-

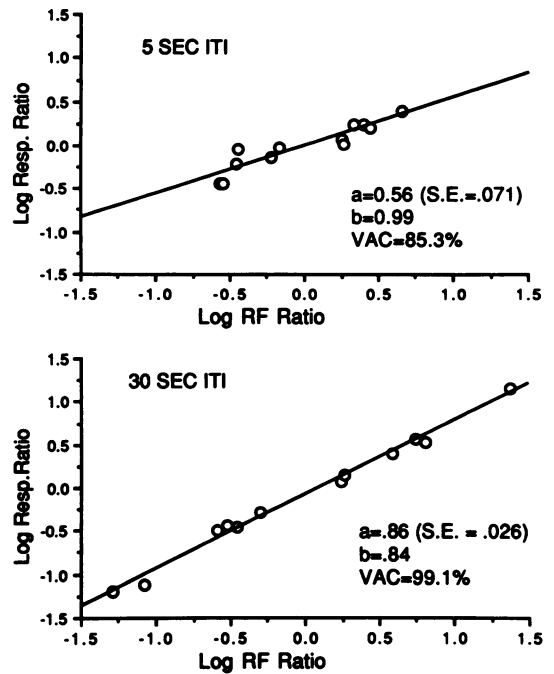


Fig. 7. Log of the response ratios (left vs. right responses) as a function of the log of the reinforcement ratios for all conditions in Experiment 2. Data are aggregated over groups for each ITI condition. Parameter values refer to Equation 1, not its logarithmic version.

tained reinforcement ratios different from those on which the matching analysis was based.

An alternative analysis that incorporates obtained reinforcement ratios for the two response levers is provided by the generalized matching law (Equation 1) in which B is responses to the left and right levers, R is the reinforcers received for those responses, the parameter, b , is bias and the parameter, a , is the sensitivity of the response ratios to the reinforcement ratios:

$$B_l/B_r = b(R_l/R_r)^a \quad (1)$$

Figure 7 shows the fits by this equation to the aggregate data from each ITI condition. The top portion contains the data from Phases 1 and 2 for Group 1 and Phase 3 for Group 2; the bottom portion contains the data from Phases 1 and 2 for Group 2 and Phase 3 for Group 1. Each data point corresponds to the results of an individual subject for each condition. Figure 7 shows that the generalized matching law fit the results with the 30-s ITI extremely well, accounting for more than 99%

of the variance. The fit with the 5-s ITI was less accurate, both in terms of the variance accounted for and the standard error of the slope estimate. Most importantly, the exponent with the 30-s ITI was substantially higher ($a = 0.86$ with the 30-s ITI vs. $a = 0.56$ with the 5-s ITI).

DISCUSSION

The present results indicate that control by molar reinforcement contingencies becomes more evident when the more molecular contingencies fail to be discriminated. This is shown by the application of the generalized matching law to the results obtained with the different ITIs. Better fits of Equation 1 were obtained with the 30-s ITI than with the 5-s ITI; the exponent value of Equation 1 with the 5-s ITI (0.56) was below that usually obtained with matching studies, whereas with the 30-s ITI the exponent value was in the range of 0.80 to 1.0 often reported with pigeon free-operant experiments (cf. Baum, 1979; Wearden & Burgess, 1982). This high exponent value occurred despite there still being some degree of control by the local contingencies even with the longer ITI. That is, several subjects continued to shift away differentially from the majority lever after a nonreinforced response to that lever. It seems likely that control by such local contingencies may be involved in other cases of undermatching as well.

The pattern of choice probabilities seen in Figures 5 and 6 provides direct evidence of the degree of discrimination: Whenever that discrimination was high, control by the molar contingencies was weak; when that discrimination was poor, control by the molar contingencies was strong. Thus, regularities at the molar level (e.g., the matching law) need not be mediated by control by local contingencies of reinforcement; instead, in the present situation control by the molar contingencies was in competition with control by the local contingencies. The results are thus consistent with the previous studies of Hiraoka (1984) and Zeiler (1987).

The differences between the 5-s and 30-s ITI conditions cannot, however, be understood simply as the result of the degree of competition between local versus molar reinforcement contingencies. As in Experiment 1, there was the additional feature of a strong perseverative effect of reinforcement on the preceding

trial with the 5-s ITI. Consequently, there was a substantial tendency for the molar choice proportions to track molar relative reinforcement rates even when there was strong evidence of other types for discrimination of the local contingencies. Because this repetition of a just-reinforced response necessarily built in an approximation to the matching relation, the result was the exponent value of Equation 1 of 0.56 for the 5-s condition instead of the value of 0 that would be expected if complete adherence to the maximizing strategy had occurred. Similar perseverative effects have been reported in some previous studies of choice (e.g., Bailey & Mazur, 1990; Shimp, 1966) but not in others (Nevin, 1969). Why such differences occur is not clear; however, one possibility is the difference in the type of stimulus. Those studies that have reported strong perseverative effects used positional cues, whereas those reporting little or no difference used visual cues. Evidence that this distinction is important comes from the extensive literature on probability learning, in which "reward following" has been commonly observed with spatial problems but not with nonspatial problems (see Sutherland & Mackintosh, 1971, for a review). Such a possibility suggests that what may appear to be minor procedural variations may produce major differences in experimental outcome.

A second complication for interpretation of the present results is the apparent importance of order effects. The two groups of subjects differed substantially in their performance on the two conditions which both received, and the most likely explanation of these differences was the carryover from the preceding condition. Thus, subjects in Group 2 showed poor discrimination of the local contingencies with the 5-s ITI and a reinforcement probability of .25 for shifting away from the majority response, whereas Group 1 showed considerably more discrimination. A similar but smaller difference, in the opposite direction, occurred for the corresponding condition with the 30-s ITI. In both cases, the prior training with the larger probability of reinforcement for shifting (.4 instead of .25) appeared to facilitate the discrimination of the local contingency. This observation is important because it suggests that the discrimination of such local contingencies depends critically on the size of the differential reinforcement probability. Stimulus control is

more likely to be established with larger differences in local reinforcement probability, and, once established, may persist even after the differences in reinforcement probability are reduced to a level that would not enable the establishment of stimulus control if presented from the beginning of training. Thus, different orders of training may be one variable accounting for different experimental outcomes.

A major issue raised by the present results is how they are to be related to conventional choice experiments. Most such studies involve temporal parameters more similar to the 5-s ITI rather than to the 30-s ITI conditions of the present study, which suggest that local contingencies of reinforcement should be an important variable. For example, Nevin (1969) used a 6-s ITI with a discrete-trial concurrent VI VI schedule, whereas Williams (1985) used a 6-s ITI with a discrete-trial concurrent variable-interval variable-ratio (VI VR) schedule. In neither study was there any evidence of control by local contingencies of reinforcement; the probability of choosing a given response did not increase with time since the last choice of that response, despite the large increase in the scheduled probabilities of reinforcement associated with the longer interchangeover times.

The pattern of results obtained here is perhaps most related to those of Silberberg and Williams (1974), who trained pigeons on a discrete-trial probability learning task in which the ITI was either 1, 22, or 120 s. The contingencies were like a conventional concurrent VI VI in that the probability of reinforcement for a given response increased with the number of responses on the alternative response, but were unlike a conventional schedule in that only alternation responses were reinforced. Subjects with the 1-s ITI learned to alternate very quickly from the onset of training, but subjects in the 120-s condition never learned to alternate at a high level. Subjects with the 22-s ITI were perhaps most interesting in that early in training their relative response rates approximately matched their relative reinforcement rates, but then they gradually departed from matching in favor of consistent alternation over the course of 100 sessions of training. Silberberg and Williams (1974) interpreted this pattern of results as showing the importance of temporal factors determining the power of local reinforcement contingencies, but

it is also important to recognize the trade-off between control by the local contingencies and control by the molar contingencies. That is, in the absence of control by the local contingencies, control by the molar contingencies did occur, with an approximation to matching as the result. Thus, like the present results, their findings suggest that the two types of control are in competition, and that control by the molar contingencies cannot be reduced to a more molecular basis.

Given the substantial control by the local contingencies of reinforcement with the 5-s ITI in the present case and in that of Silberberg and Williams (1974), the question becomes why such local control failed to occur in the studies of Nevin (1969) and Williams (1985). One possible reason for the disparity is the nature of the cues involved in producing the local control. In the studies of Nevin and Williams, differences in the local probabilities of reinforcement could be discriminated only as a function of the time since the last response to a given alternative. But here time was not a relevant factor, and the subject needed only to discriminate the location of the preceding response. The same was true in the study of Silberberg and Williams, because their reinforcement contingency allowed reinforcement only for response alternation. Interestingly, Silberberg and Williams found very poor temporal control, while at the same time showing that alternation was readily learned. That is, the functions relating run length to changeover probability were flat (and in some cases decreasing) despite there being a strong tendency to alternate. This pattern suggests that a positional discrimination was much easier than a time-based discrimination, so that control by the local contingencies was more easily established. In any event, it is apparent that control by local contingencies of reinforcement may or may not occur, depending upon the particular experimental conditions.

Despite the complexities of control revealed by the present study, it is important to recognize the underlying fundamental issue. Given different response alternatives with different reinforcement histories that produce different "response strengths," what is the rule that determines the allocation of responding to the two alternatives? The rule sponsored by momentary maximizing theories is that the response with the greatest strength will always

be chosen, so that variation in the choice allocation is due to changing stimulus sets that control different levels of response strengths at different moments in time. Intermediate levels of choice thus result from the averaging of the effects of the different stimulus sets and do not directly reflect the choice rule per se. Alternatively, choice allocation may follow some proportional rule, with matching being the simplest but not the only example, by which different probabilities of occurrence are specified by the different strengths. Intermediate choice levels thus reflect the direct operation of the proportionate rule, which cannot be reduced to more molecular mechanisms, in the same sense that rate constants govern the rate of emission of photons in physics. The present results are relevant to this distinction because they demonstrate that when the discrimination of the different local probabilities is reduced by increasing the ITI, that behavior becomes more orderly, not less, in terms of control by the molar reinforcement probabilities. Thus, the present data support the proportionate rule for response allocation when longer ITIs separate the choices. On the other hand, the short-ITI conditions appear to favor momentary maximizing (with the complication of the perseverative effects of reinforcement). But, as noted by Herrnstein and Loveland (1975), this may only reflect an inadequate specification of the response units. If "response alternation" becomes a functional unit, in addition to "left and right" responding, then a different set of calculations would be necessary to show that the proportionate rule of response allocation is inappropriate. Given the clear evidence for the proportionate rule with the long-ITI conditions, parsimony seems to dictate that the analysis provided by Herrnstein and Loveland be taken seriously. Accordingly, the present results appear to favor the proportionate rule for response allocation, although it remains an open question whether a different type of molecular theory (e.g., melioration) can provide a more fundamental account of the molar results.

REFERENCES

- Bailey, J. T., & Mazur, J. E. (1990). Choice behavior in transition: Development of preference for the higher probability of reinforcement. *Journal of Experimental Analysis of Behavior*, *53*, 409-422.
- Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, *32*, 269-281.
- Catania, A. C. (1966). Concurrent operants. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 213-270). New York: Appleton-Century-Crofts.
- Cole, S., Hainsworth, F. R., Kamil, A. C., Mercier, T., & Wolf, L. L. (1982). Spatial learning as an adaptation in hummingbirds. *Science*, *217*, 655-657.
- Evenden, J. L., & Robbins, T. W. (1984). Win-stay behaviour in the rat. *Quarterly Journal of Experimental Psychology*, *36B*, 1-26.
- Galbicka, G., & Platt, J. R. (1986). Parametric manipulation of interresponse-time contingency independent of reinforcement rate. *Journal of Experimental Psychology: Animal Behavior Processes*, *12*, 371-380.
- Herrnstein, R. J., & Loveland, D. H. (1975). Maximizing and matching on concurrent ratio schedules. *Journal of the Experimental Analysis of Behavior*, *24*, 107-116.
- Hinson, J. M., & Staddon, J. E. R. (1983). Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior*, *39*, 25-47.
- Hiraoka, K. (1984). Discrete-trial probability learning in rats: Effects of local contingencies of reinforcement. *Animal Learning & Behavior*, *12*, 343-349.
- Morgan, M. J. (1974). Effects of random reinforcement sequences. *Journal of the Experimental Analysis of Behavior*, *22*, 301-310.
- Nevin, J. A. (1969). Interval reinforcement of choice behavior in discrete trials. *Journal of the Experimental Analysis of Behavior*, *12*, 875-885.
- Olton, D. S. (1979). Mazes, maps, and memory. *American Psychologist*, *34*, 583-596.
- Plowright, C. M. S., & Shettleworth, S. J. (1990). The role of shifting in choice behavior of pigeons on a two-armed bandit. *Behavioural Processes*, *21*, 157-178.
- Shimp, C. P. (1966). Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, *9*, 443-455.
- Shimp, C. P. (1976). Short-term memory in the pigeon: The previously reinforced response. *Journal of the Experimental Analysis of Behavior*, *26*, 487-493.
- Shimp, C. P. (1979). The local organization of behaviour: Method and theory. In M. D. Zeiler & P. Harzem (Eds.), *Advances in the analysis of behaviour: Vol. I. Reinforcement and the organization of behaviour* (pp. 261-298). Chichester, England: Wiley.
- Shimp, C. P. (1984). Timing, learning, and forgetting. In J. Gibbon & L. Allen (Eds.), *Timing and time perception* (pp. 346-360). New York: New York Academy of Sciences.
- Silberberg, A., Hamilton, B., Zirriax, J. M., & Casey, J. (1978). The structure of choice. *Journal of Experimental Psychology: Animal Behavior Processes*, *4*, 368-398.
- Silberberg, A., & Williams, D. R. (1974). Choice behavior on discrete trials: A demonstration of the occurrence of a response strategy. *Journal of the Experimental Analysis of Behavior*, *21*, 315-322.
- Sutherland, N. S., & Mackintosh, N. J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.
- Wearden, J. H., & Burgess, I. S. (1982). Matching since Baum (1979). *Journal of the Experimental Analysis of Behavior*, *38*, 339-348.

- Williams, B. A. (1971a). Color alternation learning in the pigeon under fixed-ratio schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, **15**, 129-140.
- Williams, B. A. (1971b). Non-spatial delayed alternation by the pigeon. *Journal of the Experimental Analysis of Behavior*, **16**, 15-21.
- Williams, B. A. (1972). Probability learning as a function of momentary reinforcement probability. *Journal of the Experimental Analysis of Behavior*, **17**, 363-368.
- Williams, B. A. (1983). Effects of intertrial interval on momentary maximizing. *Behaviour Analysis Letters*, **3**, 35-42.
- Williams, B. A. (1985). Choice behavior in a discrete-trial concurrent VI-VR: A test of maximizing theories of matching. *Learning and Motivation*, **16**, 423-443.
- Williams, B. A., Preston, R. A., & de Kervor, D. E. (1990). Blocking of the response-reinforcer association: Additional evidence. *Learning and Motivation*, **21**, 379-398.
- Zeiler, M. D. (1987). On optimal choice strategies. *Journal of Experimental Psychology: Animal Behavior Processes*, **13**, 31-39.

Received December 26, 1990
Final acceptance June 6, 1991