CONTEXT SPECIFICITY OF CONDITIONED-REINFORCEMENT EFFECTS ON DISCRIMINATION ACQUISITION

BEN A. WILLIAMS AND ROGER DUNN

UNIVERSITY OF CALIFORNIA, SAN DIEGO STATE UNIVERSITY

Pigeons were trained on a series of reversals of a simultaneous form discrimination in which the trial outcomes were separated from the choice responses by an 8-s delay interval. Different conditions were defined by the stimuli occurring during the two halves of the delay interval. Discrimination learning was greatly facilitated by having differential stimuli during the delay following correct versus incorrect choices. When the differential stimuli appeared only at the midpoint of the delay, some facilitation occurred relative to when no different stimuli occurred, but there was substantially less facilitation than when the differential stimuli occurred immediately contingent on choice. A reversed-stimulus condition, in which the stimulus at the onset of the delay following a correct choice was the same as that during the last segment of the delay following an incorrect choice, and the stimulus at the onset of the delay following an incorrect choice was the same as that preceding food during the last segment of the delay following a correct choice, also facilitated discrimination learning relative to the nondifferential stimulus conditions.

Key words: conditioned reinforcement, delay of reinforcement, discrimination, reversal learning, conditional discrimination, stimulus segmentation, key peck, pigeons

---

When animals are presented a simultaneous discrimination in which the choice responses are temporally separated from the trial outcomes, the rate of learning is strongly affected by the stimuli present during the delay interval. In the absence of any differential stimuli during the delay, discrimination acquisition is greatly retarded, and may completely fail to occur if the delay interval is extended beyond a few seconds (Grice, 1948). But when differential stimuli are present, the rate of learning is relatively unaffected by the delay value.

A critical issue in the history of learning theory has been how to characterize such stimulus effects. The traditional interpretation has been in terms of the concept of conditioned reinforcement, in that the stimulus during the delay is assumed to acquire conditioned value in its own right because of its pairing with the reward that occurs at the end of the delay interval. The stimulus then serves as an effective reinforcer for the choice of the stimulus (S+) with which it is temporally contiguous.

Although the conditioned-reinforcement interpretation of such stimulus effects is sup-

ported by a variety of data, alternative interpretations have also been proposed. These include *marking,* the concept that any stimulus contingent on a response will enhance its memorability and thus make it more easily associated with the delayed reward (Lieberman, McIntosh, & Thomas, 1979), and *bridging,* the concept that the stimulus provides an associative link between the response and reinforcer by causing them to "go together" (Rescorla, 1982). Distinguishing between these alternative interpretations by unambiguous empirical tests has proven to be difficult.

A critical feature of both the marking and bridging hypotheses is that the effect of the stimulus during the delay is to enhance the associative linkage between the choice and outcome. It thus should not depend on the value of the stimulus itself. A potential method to contrast the different hypotheses is thus to pit the contingency between the choice and stimulus (conditioned reinforcer) against the contingency between the choice and delayed outcome. An important example of such a procedure was reported by Cronin (1980). Pigeons were presented a simultaneous discrimination with a 60-s delay of reinforcement. When the same stimulus appeared during this delay regardless of the choice response, acquisition of the discrimination failed to occur over a course of 30 sessions of training. When differential stimuli extended throughout the delay interval following the choice of S+ ver-

sus the stimulus not associated with reward (S−), rapid acquisition occurred. The critical condition was that of a "reversed cue" group, which received one color of houselight (e.g., yellow) for the first 10 s after an S+ choice and then a different houselight color (e.g., blue) during the last 10 s of the delay prior to food. After an S− choice, the houselight presentations were reversed: The blue houselight occurred immediately after an S− choice, and the yellow houselight occurred during the last 10 s of the 60-s delay before the onset of the next intertrial interval (ITI). This reversed-cue condition not only failed to produce discrimination acquisition but also produced a consistent choice (approximately 90%) of the S− stimulus. Presumably this was because choice of the S− had the immediate effect of producing the stimulus (blue houselight) that preceded food on the trials on which the S+ was chosen.

Although Cronin's (1980) results provide strong support for the conditioned-reinforcement interpretation of stimulus effects on delayed reinforcement contingencies, subsequent research from the same laboratory has shown these stimulus effects to have complex determinants, in that choice of the S− failed to occur when the stimuli at the beginning and end of the delay interval were very brief and also failed to occur with short delay-of-reinforcement intervals (Winter & Perkins, 1982). The previous results also leave unclear how such effects change with continued training, because all of the prior work has been done with naive subjects learning a single simultaneous discrimination, and for some subjects the effect of the reversed-stimulus condition was to produce choice of the S− early in training, but then a switch to choice of the S+ later in training. Because of the general importance of these results for an understanding of conditioned-reinforcement effects, it is worthwhile to explore the determinants of such effects in further detail and to examine how the reversed-stimulus condition used by Cronin compares with other stimulus arrangements during delay-of-reinforcement intervals. The present study provided such an investigation, using a repeated-reversal procedure that has proven to be a useful within-subject procedure for studying other variables (e.g., Williams, 1981; Williams & Dunn, 1991b).

All subjects received all of the different stim-

ulus conditions shown in Table 1 at different points in training. Of primary interest was the condition indicated as RG. During the first half of the 8-s delay interval following an S+ choice, the keylights were changed to red, followed by green keylights during the second half of the interval. Following an S− choice, the sequence of colors was reversed. Thus, the stimulus temporally contiguous with food (green) was also immediately contingent on the S− choice, which means that the conditioned-reinforcement contingency was in direct competition with the delayed primary-reinforcement contingency. The procedure is essentially similar to that used by Winter and Perkins (1982, Experiment 2) in which the subjects initially began choosing the S−, only to reverse their preference with continued training. The present study goes beyond the previous study, however, by including a control condition in which no differential stimuli appeared during the delay interval, to provide a determination of whether the reversed-stimulus condition improved or retarded discrimination performance even after continued training. The present procedure also differs in using a repeated-acquisition procedure (in contrast to the acquisition of a single discrimination in the earlier studies).

A second comparison of interest was the performance during Conditions GG versus BG. In the latter condition, the stimulus presentation during the delay interval was segmented into two chunks, whereas in the former there was a continuous stimulus throughout. Given previous data showing that preference is greater for nonsegmented stimuli in concurrent-chains procedures (Leung & Winton, 1988), it was of interest to determine whether a comparable effect would occur with a repeated-acquisition procedure.

The remaining conditions were added in later phases of training to assess the role of differential delayed reinforcement contingencies. They differed primarily in terms of providing differential stimulus feedback only during the second segment of the stimulus rather than during both segments. Such procedures also are of interest because of their similarity to "blocking" procedures (e.g., Williams, 1975) in which stimuli at the end of a delay-of-reinforcement interval have been shown to reduce the level of conditioning using the measure of response rate. It is uncertain whether

the occurrence of differential stimuli at the end of a delay interval will facilitate or retard discrimination acquisition relative to the no-signal conditions.

METHOD

Subjects

Four male White King pigeons, maintained at 85% of free-feeding weights, served as subjects. When necessary, supplemental feedings were given after the experimental sessions. Water and grit were freely available in the home cages. All subjects had prior experience with concurrent-chains schedules of reinforcement.

Apparatus

Each subject was assigned to one of four experimental chambers. The chambers were cubes, 32 cm on a side, housed within wooden enclosures. In each, one side panel was a Plexiglas door; the remaining sides and ceiling were aluminum. Three translucent response keys were mounted on the front panel. The keys were 2.5 cm in diameter and were evenly separated at 24 cm above the grid floor. The keys could be transilluminated with various colors and white forms. The food hopper opening was located 9 cm beneath the center key. When activated, the solenoid-operated hopper was illuminated by white light and allowed 3-s access to mixed grain. A houselight mounted in the center of the ceiling provided general chamber illumination except during the operation of the hopper. Stimuli, contingencies, and data collection were controlled by an XT-compatible computer with Turbo Pascal software.

Procedure

A simultaneous discrimination procedure was common to all conditions. Trials began with a horizontal white bar on one side key and vertical white bar on the other, with the stimulus/side assignments randomized across trials. These stimuli remained illuminated until there were 10 key pecks to either stimulus. Within each session, 10 pecks to one stimulus, S+, was designated a correct response and led to food. Ten pecks to the other stimulus, S−, was designated incorrect and led to blackout. In either case, the outcomes (food or blackout)

Table 1

Stimuli present during the two segments of the delay-of-reinforcement interval for each condition.

Condition	Response	1–4 s	5–8 s
Phase 1			
GG	correct	green	green
	incorrect	red	red
BG	correct	blue	green
	incorrect	yellow	red
RG	correct	red	green
	incorrect	green	red
--	correct	none	none
	incorrect	none	none
Phase 2			
GG, RG, BG, and -- were continued with the addition of			
-G	correct	none	green
	incorrect	none	red
Phase 3			
GG, RG, BG, and -G were continued with the addition of			
CG	correct	circle	green
	incorrect	circle	red
XG	correct	red/green	green
	incorrect	red/green	red

were delayed 8 s. A 30-s ITI followed either outcome.

Conditions differed only in terms of the stimuli on the side keys during the response-outcome delay. Both side keys were illuminated with the same color (or in one condition a white circle) or were extinguished (the no-signal and -G conditions). The delay interval was divided into two 4-s segments, and in most conditions different stimuli were illuminated in the two segments. The stimulus assignments are presented in Table 1.

Once a discrimination was acquired, as determined by a criterion of 10 successive correct responses within a session, the session was terminated, and the next session began another condition with the discrimination contingencies with respect to the horizontal and vertical bars reversed (i.e., the previous S+ became S− and the previous S− became S+). If the discrimination was not acquired by the end of a session, the session was terminated after 80 trials, and training on the same contingencies resumed the next day.

The order of conditions was random with two constraints. Within each phase of training, the conditions were sampled without replacement until all conditions studied in that phase

had an equal number of presentations; no condition was presented twice in succession. The three phases differed with respect to which conditions were included, as shown in Table 1. During Phase 1, only the GG, BG, RG, and no-signal (–) conditions were presented. For the no-signal condition, the keylights were off during both segments of the delay interval. During Phase 2, these conditions were continued with the addition of condition -G, which meant that the initial segment of the delay was like the no-signal condition and differential stimuli appeared during the second segment of the delay. During Phase 3, the no-signal condition was dropped and two new conditions were added: XG, in which red and green appeared randomly during the first segment of the delay interval independent of whether the choice was to the S+ or S–, and CG, in which a white circle appeared during the first segment of the delay after both types of choices. Different subjects received different numbers of reversals per condition across the three phases. For S-1, the number of reversals per condition was 8, 12, and 8 for Phases 1, 2, and 3, respectively. For S-2, the corresponding numbers were 12, 12, and 8; for S-3, 8, 12, and 8; and for S-4, 4, 12, and 8.

RESULTS

Figure 1 shows the mean number of trials required for each reversal to be learned to the criterion of 10 consecutive correct responses, subdivided according to the phase of training and the stimulus conditions during the delay interval. Although the absolute level of trials required varied widely across subjects, the pattern of results was generally similar for all subjects. For statistical purposes, planned comparisons were performed by first normalizing the measure of absolute trials to criterion for individual subjects to remove the large differences in rate of learning for individual subjects as a source of variance. This was done by calculating the average number of trials per reversal from all reversals learned within a given phase of training for a given subject, and then converting the trials per reversal for individual stimulus conditions into proportions of this average value. Normalizing the trials to criterion within phases, as opposed to across the entire duration of training, was done to remove any effect of changes in the rate of

learning for reasons other than the different stimulus conditions (e.g., learning-to-learn effects), which was a potential confounding effect with the stimulus variable because different stimulus conditions were used in the different phases of training. A paired t test, with three degrees of freedom and a .05 significance level, was then used for comparisons of individual means.

Phase 1. Comparing first only the conditions presented during Phase 1 (depicted by the black bars in Figure 1), there was no consistent difference between the GG and BG conditions, indicating the absence of any effect of stimulus segmentation. The mean proportions of the average number of trials were 0.59 for the GG condition and 0.63 for the BG condition. This difference did not approach statistical significance ($t < 1$). To assess the role of differential stimuli during the delay, the BG and GG conditions were averaged and compared to the no-stimulus condition. The mean proportion for the combined conditions was 0.61, whereas that of the no-signal condition was 1.76. This difference was statistically significant, $t = 7.05$. Of primary interest was the RG condition. All subjects required more trials with this condition than with the GG and BG conditions. The mean proportion of the combined BG and GG conditions was again 0.61, whereas that of the RG condition was 1.03; this difference was statistically significant, $t = 3.24$. As will be seen below, however, the size of this difference decreased over the course of training. Perhaps most importantly, the rate of learning within the RG condition was more rapid than in the no-stimulus condition for all subjects, indicating that the presence of the stimuli during the delay interval facilitated learning despite the reversal of the roles of the stimuli during the two halves of the delay interval. The difference between the respective proportions of the average number of errors (1.76 vs. 1.03) was statistically significant, $t = 5.94$.

Phase 2. The -G condition was added during Phase 2. Because it entailed no stimuli during the first 4 s of the delay, any difference between the -G and no-signal conditions would necessarily be due to the differential stimuli that occurred during the second half of the delay for the -G condition. Similarly, the stimuli during the second half of the delay interval were the same as those during the second half of the RG, BG, and GG conditions, so that
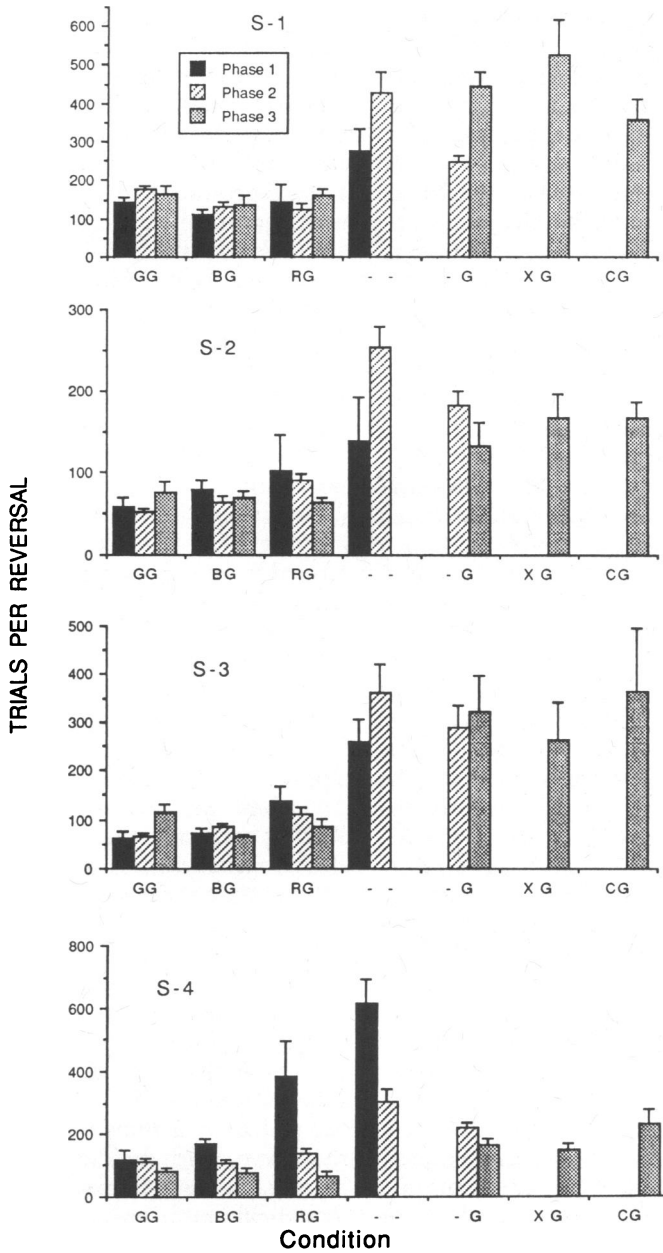
Fig. 1. Mean number of trials required for subjects to reach the learning criterion for each reversal. Results are subdivided into the three phases noted in Table 1. Stimulus conditions during the delay interval are shown below each set of data bars. Standard errors of the mean are shown by the error bars. Note that different numbers of reversals per condition occurred in the different phases.

any difference between these conditions would necessarily be due to the role of the differential stimuli during the first segment of the interval. Figure 1 shows that the rate of learning during Phase 2 (hashed bars) was substantially better under the -G condition than under the no-

signal condition for all 4 subjects. The rate of learning under the RG condition in Phase 2 was also substantially greater than during the -G condition for all 4 subjects. Averaged over subjects, the mean absolute number of trials required to reach the learning criterion during

Phase 2 was 336 for the no-signal condition, 234 for the -G condition, and 117 for the RG condition. Thus, the stimuli in both segments of the delay interval played a significant role in determining the level of performance.

Statistical analysis similar to that used for Phase 1 was performed on the results of Phase 2. The mean proportions of the average trials to criterion for the BG and GG conditions were 0.59 and 0.60, and this difference did not approach statistical significance, again indicating no effect of the segmentation variable on the rate of acquisition. The difference between the combined BG and GG conditions versus the no-signal condition (mean proportion = 2.01) was again statistically significant, t = 12.9. The difference between the combined BG and GG conditions versus the RG condition (mean proportion = 0.70) was now not significant, t = 0.99, in contrast to the same comparison during Phase 1. The difference between the RG and no-signal conditions was again statistically significant, t = 8.27.

The combined BG and GG conditions were significantly different from the -G condition (mean proportion = 1.42), t = 5.42, as was the RG condition, t = 6.91. Also, however, the -G condition was statistically significantly different from the no-signal condition, t = 4.62, demonstrating that the differential signals during the second half of the delay did facilitate learning relative to no signals at any time during the delay interval (i.e., a delayed conditioned-reinforcement effect). It should be noted that this analysis applies only to the comparison of the two conditions within Phase 2. Figure 1 shows that the rate of learning during the no-signal condition was variable across Phases 1 and 2, although there is no consistent pattern to the differences across subjects. Similarly, performance under the -G condition varied considerably across Phases 2 and 3. The result is that a comparison of the average trials to criterion for all reversals for the two conditions (Figure 1) suggests a smaller difference between the two conditions, although a difference still did occur. Averaged over all reversals, the average trials to criterion for the no-signal condition was 317, whereas that for the -G condition was 247. This difference occurred for all 4 subjects.

Phase 3. The no-signal condition was eliminated in Phase 3, and two new conditions with nondifferential signals during the first half of the delay were added. For Condition XG, red and green appeared randomly, independent of the choice response; for Condition CG, a white circle appeared on the key immediately after both correct and incorrect choices. Performance produced by these two new conditions was essentially similar to that obtained with the -G condition, with substantially slower learning during all three of these conditions relative to that during the RG condition (and also the GG and BG conditions). For 2 subjects, S-1 and S-3, performance under the -G condition was noticeably worse in Phase 3 than in Phase 2, suggesting that the introduction of the new conditions decreased performance generally.

Statistical analysis confirmed these observations. Performance in the -G condition continued to be significantly worse than the combined BG and GG conditions, t = 5.80, but was not statistically different from either the CG, t = 1.11, or XG conditions, t = 0.27. Thus, the nature of the stimuli during the first half of the delay, as long as they were nondifferential with respect to the choice response, appears to be unimportant. In contrast, the RG condition, with the reversed stimuli following correct versus incorrect responses, produced significantly faster learning than each of the other three conditions: RG versus -G, t = 7.32; RG versus CG, t = 6.61; RG versus XG, t = 7.23. The RG condition produced results not significantly different from those of the combined BG and GG conditions, t = 1.47. Finally, unlike Phases 1 and 2, in which the results of BG and GG conditions were virtually identical, here the BG condition produced faster learning than the GG condition for all 4 subjects, with mean proportions of 0.50 versus 0.61, although this difference did not attain statistical significance, t = 2.34. A possible cause of this change in the relative effectiveness in the two stimulus conditions was the introduction of the XG condition, which degraded the information value of the green stimulus appearing in the first half of the delay interval.

Figure 1, along with the preceding statistical analysis, shows that the effects of the RG condition changed over the course of training. For all 4 subjects the rate of learning under that condition was greater during Phase 2 than during Phase 1, and for 3 of the 4 subjects

(excepting S-1) a further improvement occurred during Phase 3. This trend cannot be ascribed to a generalized improvement in the rate of learning, because it is not evident for either the GG or the BG condition.

A closer examination of the changes in the rate of learning over the first series of reversals is shown in Figure 2, which restricts the analysis to the first 12 reversals of the GG, BG, and RG conditions. For all 4 subjects the number of trials needed to attain the learning criterion during the first two reversals was greater for the RG condition than for the GG and BG conditions, which were never systematically different throughout training. The rate at which performance under the RG condition then approached that under the other conditions varied for the different subjects. For S-1 and S-4, no difference was evident by the second pair of reversals; for S-2, the difference persisted until the third pair of reversals; and for S-3, it persisted until the sixth pair of reversals. These observations were assessed statistically with a two-way ANOVA (Blocks × Stimulus Condition). The main effect of Stimulus Condition was not significant, $F(2, 6) = 4.20$, $.05 < p < .10$; the main effect of Blocks was also not significant, $F(5, 15) = 1.30$, $p > .10$. The interaction between the two variables was significant, $F(10, 30) = 2.23$, $p < .05$.

The pattern of within-reversal improvement was also examined for the RG condition. Unfortunately, the data were not recorded on a trial-by-trial basis, so this analysis was restricted to the changes in the accuracy of the discrimination across the different sessions, in cases when more than one session was required to attain the learning criterion of 10 consecutive correct trials. Our expectation was that the RG condition would maintain below-chance performance for the first sessions of reversals early in training, because of the conditioned reinforcement of incorrect choices from the green color being immediately contingent on such responses (cf. Cronin, 1980). Although this pattern did occur in some cases, the pattern was not consistent over subjects or over reversals for the same subject. The pattern also occurred for some reversals during the no-signal condition, presumably because the use of the reversal procedure entailed that the S+ of the preceding reversal now became the S−, so that below-chance performance could be the
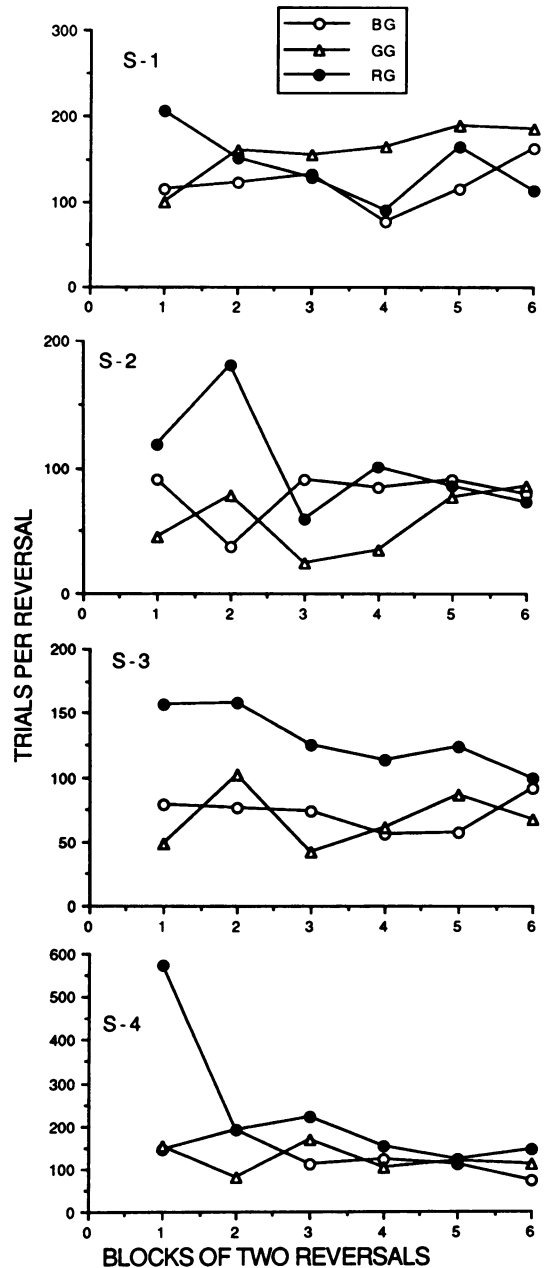

Fig. 2. Mean number of trials required for subject to reach the learning criterion across the first 12 reversals of training for each condition. Separate functions correspond to the three different stimulus conditions during the delay interval.

Table 2

Mean responses per trial and discrimination ratios during the two halves of the delay-of-reinforcement interval. Data are averages of all reversals from Phase 3. Note that the discriminaation indices were calculated as the mean of the discrimination indices of individual reversals and not from the mean response rates across reversals.

Sub-ject	Condi-tion	1-4 (S+)	1-4 (S−)	Dis-crim-ina-tion	5-8 (S+)	5-8 (S−)	Dis-crim-ina-tion
S-1	GG	12.2	7.1	.64	11.8	2.5	.82
	BG	13.7	3.4	.81	12.9	1.6	.91
	RG	13.8	10.4	.58	12.9	4.5	.74
	-G	11.5	11.3	.51	13.6	2.8	.85
	XG	12.6	12.1	.51	12.0	7.3	.64
	CG	15.5	15.5	.52	14.3	3.3	.84
S-2	GG	8.5	4.9	.70	19.6	1.8	.87
	BG	9.1	1.0	.93	11.0	1.6	.92
	RG	10.8	8.2	.57	12.0	1.8	.89
	-G	9.4	8.8	.51	11.6	1.1	.92
	XG	9.3	9.3	.51	11.6	3.4	.77
	CG	11.2	10.1	.50	11.0	1.3	.88
S-3	GG	11.0	3.3	.77	19.6	1.8	.92
	BG	6.2	2.3	.74	21.2	0.1	.99
	RG	8.4	12.2	.41	21.6	2.4	.92
	-G	7.0	7.8	.45	19.1	1.0	.95
	XG	11.0	10.5	.50	19.6	3.3	.86
	CG	17.2	17.5	.50	18.7	1.4	.93
S-4	GG	18.6	15.0	.57	14.0	6.9	.68
	BG	19.6	12.8	.63	15.1	3.9	.80
	RG	17.8	17.4	.52	15.8	6.3	.75
	-G	16.4	17.2	.49	15.2	3.3	.86
	XG	18.2	18.5	.51	16.0	5.9	.73
	CG	21.0	21.5	.49	15.6	2.4	.88
M	GG	12.6	7.6	.67	13.9	3.2	.82
	BG	12.1	4.9	.78	15.0	1.8	.91
	RG	12.8	11.8	.53	15.6	3.8	.82
	-G	11.1	11.3	.49	14.9	2.1	.89
	XG	12.8	12.6	.50	14.8	5.0	.75
	CG	16.2	16.2	.50	14.9	2.1	.88

result of perseveration of what was learned on the preceding reversal rather than a reflection of the conditioned-reinforcement properties of the green color. Because of this interpretative difficulty, the pattern of within-reversal improvement will not be discussed further.

Response rates to the different stimuli occurring during the delay intervals were recorded in Phase 3 for the different conditions and are shown in Table 2. Also shown is a discrimination index, which was calculated separately for the S+ versus S− trials for each segment of the delay interval. The discrimination measure was calculated as the response rate following a choice of the S+ relative to the sum of the response rates after an S+ choice and after an S− choice. A value of .50 indicates that the animals were responding nondifferentially during the delay interval with respect to the nature of their preceding choice response, whereas a value of 1.0 indicates that they discriminated perfectly during the delay interval as a function of the preceding choice response.

The 4 subjects varied widely in the pattern of their response rates for the two halves of the delay interval. For S-1, response rates during the first 4 s and last 4 s of the delay were similar; for S-2 and S-3, substantially lower rates occurred during the first 4 s; and for S-4, higher rates occurred during the first 4 s of the delay interval. Much less variation occurred across stimulus conditions within individual subjects, because subjects that responded at low rates during the first half of the delay under one condition also tended to respond slower during the first halves of other conditions as well (e.g., S-3). The major exception to this generalization is that for all 4 subjects, response rates were higher during the first half of the delay of the CG condition than for any other stimulus condition. In general, there appears to be little relation between the level of the absolute response rates and the degree of discrimination between the S+ and the S− trials.

Little evidence of differential responding occurred during the first segment of the delay interval for the RG, -G, XG, and CG conditions; the average discrimination index was near .50 for all of these conditions. For the last three of these conditions, performances of the 4 subjects were generally similar, whereas for the RG condition, there was considerable variability, as will be described below. The greatest degree of discrimination during the first delay segment occurred for the BG condition, and slightly less discrimination occurred for Condition GG. Discrimination under both of these conditions was considerably greater than under the remaining four conditions.

The degree of differential responding during the second segment of the delay interval was much more similar across conditions. The best discrimination again occurred for Condition BG, but comparable levels of performance occurred for -G and CG, with GG and RG producing only slightly less discrimination. The lowest level of discrimination oc-

curred for Condition XG, in which the second delay segment occurred after random presentations of red or green during the first segment. Note that -G and CG were similar to XG in the absence of any consistent relation between the stimulus during the first delay segment and the type of choice response, so that the use of random stimuli seems to have played a significant role in its own right.

Of greatest interest from Table 2 is the pattern of discrimination during the RG condition and its relation to the overall level of learning proficiency. A modest degree of differential performance during the delay interval did occur for 2 subjects under the RG condition (S-1 and S-2). For S-3, however, the level of discrimination was below .50, indicating that S-3 responded faster during the first 4 s of the delay after an S− choice than after an S+ choice, presumably because the stimulus presented after an S− choice (green) occurred just prior to food on S+ trials. The range of discrimination values for the individual reversals was .31 to .57, with four of the eight values below .40. For S-4, the mean discrimination value of .52 indicated no discrimination as a function of the preceding choice; this value was typical of the behavior on individual reversals, in which individual discrimination values clustered closely on either side of the .50 value.

Given that some variability in the degree of discrimination occurred across individual reversals, we attempted to discover any relation, across reversals, between the degree of discrimination and the overall learning proficiency during the RG conditions, as assessed by the number of trials required to reach the criterion of learning. Correlational analyses using Pearson's r were conducted for individual subjects, with the result that none approached significance. For S-1 through S-4, the obtained r values were −.23, +.34, 0, and −.01, respectively.

We also attempted to isolate any change in the discrimination performance as a result of training within a reversal. Because data were not recorded on a trial-by-trial basis within reversals, the analysis was restricted to comparing discrimination performance on the 1st day of a reversal to the remaining days, for only those reversals that required more than one session to attain the learning criterion. There was no consistent difference in the discrimination indices between these different days

of reversal training. Note, however, that all of these data are from Phase 3, so that considerable opportunity for learning the values of the different stimuli in each condition had already occurred.

DISCUSSION

The present results confirm the findings of previous studies in that differential stimuli presented during a delay-of-reinforcement interval facilitated the rate of learning. The results further indicate that some facilitatory effect occurred even when the differential stimuli occurred only later in the delay interval and were not immediately contingent on the choice responses. That is, learning occurred more readily in Condition -G than in the no-signal condition. This effect is perhaps surprising, given the findings in other studies that have examined the effects of delay of reinforcement on simple response rate, in that stimuli occurring only at the end of a delay interval reduced response rate relative to a no-stimulus condition, an effect that has been interpreted as showing "blocking" of the response–reinforcer association (Williams, 1975, 1978; Williams, Preston, & DeKervor, 1990). Apparently, such blocking effects do not occur in simultaneous discrimination procedures, a finding we have replicated in several other unpublished studies.

The results further demonstrate that the facilitatory effects of delay stimuli do not require that the stimuli contingent on the choice behavior are themselves temporally contiguous with food at the end of the delay interval. There was no consistent difference between the BG and GG conditions throughout training, despite the blue color never being paired contiguously with food. Thus, the present results stand in contrast to the effects of stimulus segmentation that have been reported in concurrent-chains studies of preference (Leung & Winton, 1988). This difference may reflect only that the preference measure is more sensitive than the discrimination-acquisition measure, or that the delay intervals used here (8 s) were substantially shorter than those in concurrent-chains studies that have shown strong stimulus-segmentation effects. Numerous other differences in the two procedures also exist (e.g., differential vs. nondifferential probabilities of reinforcement at the end of the delay) that

could account for the different effects of the segmentation variable.

The most important finding of the present study is that stimuli in the delay interval may facilitate learning even when the reward values of the stimuli immediately contingent on the choice response are inversely related to the delayed contingencies of primary reinforcement. Not only did the RG condition facilitate learning relative to the no-signal condition, it also produced a level of performance comparable to that of the BG and GG conditions. However, this proficient level of performance developed only with training; learning under RG was below that under the other two conditions early in training. Thus, any account of the facilitatory effects of stimuli in delay-of-reinforcement intervals must provide an explanation for how such learning occurred.

One interpretation of how such stimulus effects changed with training was provided by Astley and Perkins (1985), who argued that the conditioned-reinforcement function of a stimulus is dependent on the context in which the stimulus is paired with the reinforcer. Thus, it should be expected that the presentation of green during the first 4 s after a choice of the S− could be discriminated as a different event than the presentation of green during the last 4 s of the delay prior to food delivery. Consequently, green after an S− choice should not strengthen the incorrect behavior and thus retard the acquisition of the discrimination. Discrimination of the different presentations of red and green as different events would convert the red-green sequence following a correct choice to a different sequence (e.g., A-B) and the green-red sequence following an incorrect choice into a different sequence (e.g., C-D), in effect making it analogous to the BG condition that produced highly proficient learning. The reason that the RG condition required considerable training to produce the same level of proficiency would then be due to the difficulty of learning the context discrimination.

Although the preceding account can explain the facilitation obtained with the RG data, it is not supported by the discrimination data presented in Table 2. Little evidence of discrimination between the first segment of the delay interval as a function of the S+ versus S− choice is evident for the RG condition, in contrast to the considerable discrimination during the GG and BG conditions. Instead,

the RG condition appeared to be similar to the other conditions in which the stimulus conditions during the first delay segment were nondifferential as a function of the preceding choice. Yet discrimination performance under the RG condition clustered with the GG and BG conditions, with all performances considerably better than those in all of the remaining conditions. The ability of the initial-segment stimuli of the RG condition to facilitate learning thus appears to be at best weakly related to their ability to serve as discriminative cues for responding to the stimuli themselves. Perhaps this dissociation of stimulus functions reflects only differences in sensitivity of the different measures. However, this possibility is challenged by the finding that the 1 subject with below-chance discrimination during the first segment of the RG condition (S-3) was similar to the remaining subjects in the pattern of discrimination acquisition. Such a pattern cannot easily be ascribed to poor sensitivity in the discrimination measure, and suggests the need for entertaining some more fundamental basis for the dissociation between the two measures. One possibility may be that the functions of a cue as a conditioned reinforcer rather than as a discriminative stimulus for differential responding may be controlled by separable variables.

The present results are consistent with previous findings showing that the facilitatory effect of stimuli during delay-of-reinforcement intervals cannot be due simply to the history of pairings of those stimuli with primary reinforcement. Instead, the "value" of a stimulus is the result of the complex determinants of conditional discrimination, such that the nominally same stimulus is viewed as a different event when it occurs in different contexts. This realization may be crucial to interpreting various studies that have been taken as evidence against the concept of conditioned reinforcement because differential choice responding did not correspond to differential conditioned-reinforcement contingencies (e.g., Schuster, 1969). The results of Rescorla (1982) provide an illustration. In his study, two different keylights were followed on 50% of the trials by delayed food and on 50% of the trials by an immediate conditioned reinforcer. The critical difference between the stimuli was that for one of them (A), the conditioned reinforcer filled the delay interval between the stimulus and

delayed reinforcer, whereas for the other (B), the immediate stimulus and delayed food occurred on different trials. Rescorla argued that because the same stimulus was immediately contingent on both Stimuli A and B equally often, they should be equated in the "value" of the conditioned-reinforcement contingency. However, substantially stronger responding occurred to Stimulus A than to Stimulus B, a finding which Rescorla interpreted as requiring the concept of "bridging" rather than conditioned reinforcement. Subsequent research (Honey, Schachtman, & Hall, 1987; Thomas, Robertson, & Cunniffe, 1989) has challenged this interpretation, because response rate to the intervening stimulus was higher when presented following Stimulus A than following Stimulus B. This finding suggests that the two different locations of presentations of the same nominal stimulus were functionally different stimulus events, controlled by the conditional cue of which keylight had preceded it. The implication is that unambiguous demonstrations of the power of conditioned reinforcement require procedures that minimize the possibility of such conditional discrimination. Recent data from this laboratory support the validity of this conclusion (Williams & Dunn, 1991a, 1991b).

REFERENCES

Astley, S. L., & Perkins, C. C. (1985). Stimulus duration and conditioned reinforcing value measured by a learning-tests procedure. *Animal Learning & Behavior, 13*, 18–24.

Cronin, P. B. (1980). Reinstatement of postresponse stimuli prior to reward in delayed-reward discrimination learning by pigeons. *Animal Learning & Behavior, 8*, 352–358.

Grice, G. R. (1948). The relation of secondary reinforcement to delayed reward in visual discrimination learning. *Journal of Experimental Psychology, 38*, 1–16.

Honey, R. C., Schachtman, T. R., & Hall, G. (1987). Partial reinforcement in serial autoshaping: The role of attentional and associative factors. *Learning and Motivation, 18*, 288–300.

Leung, J. P., & Winton, A. S. W. (1988). Preference for simple interval schedules of reinforcement in concurrent chains: Effects of segmentation ratio. *Journal of the Experimental Analysis of Behavior, 49*, 9–20.

Lieberman, D. A., McIntosh, D. C., & Thomas, G. V. (1979). Learning when reward is delayed: A marking hypothesis. *Journal of Experimental Psychology: Animal Behavior Processes, 5*, 224–242.

Rescorla, R. A. (1982). Effect of a stimulus intervening between CS and US in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes, 8*, 131–141.

Schuster, R. (1969). A functional analysis of conditioned reinforcement. In D. P. Hendry (Ed.), *Conditioned reinforcement* (pp. 192–234). Homewood, IL: Dorsey Press.

Thomas, G. V., Robertson, D., & Cunniffe, G. (1989). Facilitation of responding in a filled-delay trace autoshaping procedure: An occasion-setting hypothesis. *Learning and Motivation, 20*, 225–241.

Williams, B. A. (1975). The blocking of reinforcement control. *Journal of the Experimental Analysis of Behavior, 24*, 215–225.

Williams, B. A. (1978). Informational effects on the response–reinforcer association. *Animal Learning & Behavior, 6*, 371–379.

Williams, B. A. (1981). Invariance in reinforcements to acquisition, with implications for the theory of inhibition. *Behaviour Analysis Letters, 1*, 73–80.

Williams, B. A., & Dunn, R. (1991a). Preference for conditioned reinforcement. *Journal of the Experimental Analysis of Behavior, 55*, 37–46.

Williams, B. A., & Dunn, R. (1991b). Substitutability between conditioned and primary reinforcers in discrimination acquisition. *Journal of the Experimental Analysis of Behavior, 55*, 21–35.

Williams, B. A., Preston, R. A., & DeKervor, D. E. (1990). Blocking of the response–reinforcer association: Additional evidence. *Learning and Motivation, 21*, 379–398.

Winter, J., & Perkins, C. C. (1982). Immediate reinforcement in delayed reward learning in pigeons. *Journal of the Experimental Analysis of Behavior, 38*, 169–179.