Signaling functions of the second-order CS: Partial reinforcement during second-order conditioning of the pigeon's keypeck

MICHAEL E. RASHOTTE, BEVERLY S. MARSHALL, and JEFFREY M. O'CONNELL
Florida State University, Tallahassee, Florida 32306

The signaling function of the second-order CS (S2) was manipulated in second-order autoshaping by arranging a partial reinforcement schedule. S2 was paired with a well-conditioned first-order CS (S1) on a continuous reinforcement or a 25% reinforcement schedule in different groups. Schedule of reinforcement did not influence the number of S2-S1 pairings required to establish keypecking to S2. However, in the postacquisition sessions, responding to S2 was initially weaker but persisted for many more sessions on the 25% schedule than on the 100% schedule. The data indicate that S2-S1 pairings are responsible both for the acquisition of second-order keypecking to S2 and for the subsequent conversion of S2 into an inhibitory stimulus.

The second-order CS (S2) has two signaling functions in the traditional procedure for second-order conditioning. It signals impending presentations of a well-conditioned first-order CS (S1). It also acts as an explicit signal for US omission, since S1 reliably predicts US presentation except when S1 occurs in conjunction with S2. The S2-S1 signaling function is responsible for S2's acquiring second-order excitatory strength (e.g., Pavlov, 1927; Rescorla, 1980). Presumably, the S2-US omission signaling function is responsible for S2's eventually losing its effectiveness as a second-order CS during extended training, and for its ultimate status as a conditioned inhibitor (Rashotte, 1981; Rescorla, 1973).

How do the associative products of these two signaling relations interact to influence performance in second-order conditioning? There is currently no satisfactory theoretical account that answers this question, and it is probable that one will not be forthcoming until these signaling relations are examined in greater detail in the laboratory. In the available experimental work, the signaling relation between S2 and S1 has been manipulated by varying the interstimulus interval on second-order trials. These experiments indicate that in some preparations, at least, S2 will become inhibitory more easily in second-order conditioning if it overlaps S1 than if it occurs strictly prior to S1 (Pavlov, 1927; Rescorla, 1973). This result implies that S2's function as a signal for US

omission is enhanced by having both S2 and S1 present at the time US is omitted.

Another manipulation of the signaling function of S2 that has several interesting features involves imposing a partial reinforcement schedule in second-order conditioning. On such a schedule, S2-S1 pairings would occur on only a percentage of the second-order trials and S2 would be presented alone on the remaining trials. This manipulation degrades the signaling relation between S2 and S1 and, therefore, might be expected to hamper the development of S2's second-order excitatory strength. The effect of partial reinforcement on the eventual development of inhibitory strength to S2 is less certain. On the one hand, S2 might become inhibitory rapidly on a partial reinforcement schedule because the S2-alone trials could promote learning that S2 is never followed by US. On the other hand, there is some evidence that the rate at which S2 becomes inhibitory may be little affected, or even retarded, by the occurrence of S2-alone trials on the partial schedule. Zimmer-Hart and Rescorla (1974, Experiment 3) have reported that the rate at which a stimulus becomes inhibitory in a traditional conditioned inhibition procedure is *not* altered when nonreinforced presentations of that stimulus alone are intermixed among the usual nonreinforced trials on which the stimulus occurs in simultaneous compound with an excitatory CS. In a second-order conditioning experiment in our laboratory, Marshall (1976; Marshall & Rashotte, Note 1) found preliminary evidence that S2 might become inhibitory at a *slower* rate when partial reinforcement is employed in second-order conditioning.

Marshall employed the second-order autoshaping preparation (Rashotte, Griffin, & Sisk, 1977), and she set the percentage of S2-S1 pairings at 100%,

This research was supported by NSF Grant BNS-16844 (Michael E. Rashotte, principal investigator) and by NIMH Training Grant MH-11218. The authors thank Dianne L. Beidler, Steven M. Friedman, and Eileen M. Weilson for help at various stages of the work. Beverly Marshall is now at the Department of Psychology, University of Iowa. Reprints may be obtained from Michael E. Rashotte, Department of Psychology, Florida State University, Tallahassee, Florida 32306.

50%, or 25% for different groups of pigeons. Her work yielded two interesting results. First, second-order conditioning was achieved in all three groups, although a larger number of second-order trials was required to establish keypecking to S2 in the 50% and 25% groups than in the 100% groups. Second, the 25% group was most persistent in responding to S2 throughout extended second-order training (i.e., 480 trials), which suggests that degrading the signaling relation between S2 and S1 by partial reinforcement retarded the growth of inhibition to S2.

The present paper provides a more complete analysis of the effects of 100% and 25% reinforcement schedules during extended training in the second-order autoshaping preparation. Acquisition and maintenance of responding to S2 by a group trained with 25% S2-S1 pairings on second-order trials were compared with the performance of two groups trained with different variations on a 100% schedule. One of the 100% groups comprised two subgroups that received only the reinforced (i.e., S2-S1) trials that were scheduled for the 25% group; the subgroups differed in the presence or absence of additional nonreinforced presentations of a neutral stimulus at the times S2-alone trials occurred in the 25% group. The other 100% group received the same total number of second-order trials as the 25% group, but, of course, they were all S2-S1 trials.

METHOD

Subjects

Forty-five experimentally naive White Carneaux pigeons, 6 months to 1 year old, were maintained in individual cages at 80% of their free-feeding weights. Water and grit were always available to the pigeons, and the colony room was on a 12-h/12-h light-dark cycle. Training sessions were held at approximately the same time in the light portion of each cycle.

Apparatus

The experiment was conducted in eight test chambers. The animal's portion of the chamber measured 34.3 cm long, 31.1 cm wide, and 34.9 cm high. Three walls and the ceiling were painted flat black. The fourth wall was a buffed aluminum panel that included openings for a clear plastic response key (2.5-cm diameter; located 23.5 cm above the floor and 8.3 cm to the left of center) and for a recessed food cup (4.5 × 5.7 cm, centered on the wall, with bottom edge 10.2 cm above the floor). The key could be lighted different colors by applying 6-V ac current to No. 47 bulbs in a 12-cell in-line projector mounted behind the key. Keypecks with a force of at least .01 N activated the counting circuits. The food cup was the housing for a standard pigeon grain hopper modified so that a retractable metal floor covered the opening in the bottom through which grain becomes available in the usual configuration. Food pellets (45-mg, Noyes Formula C) could be dispensed onto the floor of the food cup by activating a Gerbrands pellet dispenser, and after a fixed period of time the floor was retracted to dump uneaten pellets into an inaccessible receptacle. This arrangement prevented accumulation of pellets in the food cup across trials. Throughout the experiment, the food cup was lighted during pellet delivery, and for 5 additional seconds, by a No. 1891 bulb operated by 6 V ac. Retraction of the food-cup floor was always coincident with offset of the bulb in the food cup. Whenever multiple pellet deliveries occurred in one reinforcement presentation,

pellets were delivered at a rate of 4 pellets/sec. When the session was in progress, the chamber was illuminated by a ceiling-mounted houselight in the center of the chamber (110 V ac, 7.5 W). An exhaust fan in each chamber and a white-noise source (80 dB re 20 N/m²) in the experimental room masked extraneous sounds. Programming and data collection were by a PDP/8-E computer equipped with SUPERSKED software (Snapper, van Haaren, & Inglis, 1978).

Procedure

Preliminary training. The pigeons were trained to eat reliably from the food cup. In the first session, the pigeons found food pellets in the lighted hopper when they first entered the chambers. When these pellets were eaten, an observer controlled spaced deliveries of a variable number of pellets (range 1-10). By the end of the session, the pigeons were usually eating all the pellets delivered on each occasion. The next two sessions included 20 presentations of 5 pellets each; presentations occurred at variable intervals, averaging 90 sec.

First-order conditioning. On the day following completion of preliminary training, first-order conditioning sessions began. Each session consisted of 30 trials on which a 6-sec colored keylight (S1) was followed immediately by presentation of 5 food pellets. S1 was a white keylight for half the birds and a red keylight for the other half. Throughout first-order autoshaping, and the entire experiment, the ITI was variable, averaging 90 sec.

First-order conditioning sessions continued for 25 days (750 trials). Two pigeons failed to reach the acquisition criterion of 4 trials with at least one peck to S1 in a series of 5 successive trials. The remaining 43 pigeons were assigned to three groups (Ns = 15, 15, 13) matched on the basis of asymptotic response rate to S1.

Second-order conditioning: Baseline. The groups received four to six sessions in which unreinforced presentations of additional keylight stimuli to be used in second-order conditioning were intermixed among the reinforced S1-US trials. These sessions were intended to reduce the level of generalized keypecking to the stimuli before second-order conditioning began. All groups received 30 trials in each session. For Group 100%-12 (N = 15), 10 presentations of the 6-sec green keylight that served as S2 for all groups were unsystematically intermixed among 20 S1-US trials in each session. For Group 100%-3 (N = 15) and Group 25% (N = 13), 10 presentations each of S2 and of S3 were intermixed among 10 S1-US trials. S3 was a 6-sec keycolor for these groups, red for birds trained with a white S1 and white for birds trained with a red S1.

Second-order conditioning. Every group received four first-order S1-US trials in each session and three reinforced second-order trials on which an S2-S1 pairing occurred. US was never presented on a trial when S2 occurred. The groups differed with respect to additional trials scheduled in each session. Group 25% received an additional nine trials on which S2 was presented alone, thereby making it the partial reinforcement group for second-order conditioning. For the two comparison groups, the additional trial treatments ensured that they received continuously reinforced S2 presentations. Group 100%-12 received nine additional S2-S1 trials that equated it and Group 25% for total number of S2 presentations but, of course, not for reinforced second-order trials. Group 100%-3 was divided into two subgroups. Group 100%-3A (N = 7) received nine additional trials on which the S3 stimulus was presented alone on trials when S2 occurred alone for Group 25%; Groups 100%-3A and 25% were thereby equated for number of reinforced and nonreinforced trials, but differed with respect to whether the nonreinforced trials were signaled by S2 (yielding partial reinforcement) or by another stimulus, S3. Group 100%-3B (N = 8) received no additional trials, so it was simply equated with Group 25% for number of reinforced second-order trials; for this group, the ITI continued through the period that additional trials were scheduled for the other groups. Sequencing of the trials common to all groups, and of the additional trials, varied irregularly from session to session.

Each group continued training until it received a total of 168

reinforced second-order trials. This required 14 sessions for Group 100%-12 and 56 sessions for Groups 25% and 100%-3. In the second-order conditioning phase of the experiment, the two 100% groups received 168 presentations of S2, all of them followed immediately by S1; Group 100%-3A also received 404 presentations of S3. Group 25% received 672 S2 presentations, 168 of them followed by S1 and 404 comprising S2-alone.

RESULTS

First-Order Conditioning and Second-Order Baseline

First-order autoshaping proceeded uneventfully, and stable levels of responding were achieved after a few sessions. At the end of training, the groups were matched for asymptotic performance, and, as planned, there were no reliable differences among groups in rate of keypecking during S1 presentations or in the percent of trials with at least one keypeck. Across all pigeons, mean keypeck rate was 93.2 pecks/min and mean percent trials with a response was 87.75% in the final three first-order sessions.

In the second-order baseline phase, generalized keypecking to S2 and S3 was virtually eliminated in all groups in two to four sessions.

Second-Order Conditioning

Statistical analyses indicated that, within each group, performance was not influenced by the counter-balanced colors of the S1 stimuli. Also, the two subgroups comprising Group 100%-3 did not differ in any performance measure. Accordingly, in the following presentation of results, the data were collapsed across S1 colors for each group and the data of Groups 100%-3A and 100%-3B were combined for presentation as Group 100%-3.

Acquisition. The acquisition of keypecking to S2 was quantified by examining the consistency of pecking in blocks of five successive S2 presentations. The number of trials prior to the first five-trial block in which at least one peck occurred on four or more trials was used as the index of acquisition. A total of seven pigeons spread unsystematically across the three groups failed to reach this criterion (4, 2, and 3 pigeons in Groups 25%, 100%-12 and 100%-3, respectively). Data from these pigeons are eliminated from further consideration here, leaving final group sizes of 9, 13, and 12 pigeons for the 25%, 100%-12, and 100%-3 groups, respectively.

Table 1 presents the median number of trials required by each group to reach the acquisition criterion. Group 25% required significantly more S2 trials to reach criterion than either of the 100% groups (Mann-Whitney $U_s \leq 13$, $p_s \leq .05$), and the difference between the 100% groups was not statistically reliable ($U = 50$). The larger number of trials to acquisition by the 25% group would be expected if reinforced (i.e., S2-S1) trials control acquisition of responding to S2. Table 1 also presents the median

Table 1
Median Number of Trials and Reinforced Trials to Reach Acquisition Criterion on S2

Group	Trials		Reinforced Trials	
	Median	Range	Median	Range
25%	91.0	21-204	23.0	6-57
100%-12	13.0	2-58	13.0	2-58
100%-3	7.5	4-15	7.5	4-15

number of reinforced trials to acquisition, and statistical comparison indicated no significant differences between any pair of groups in this measure ($U_s \geq 48.5$). This latter result suggests that S2-alone trials intermixed among the reinforced S2-S1 trials had little effect on the acquisition of keypecking to S2. This finding implies a commonality in the effects of partial reinforcement in first- and second-order autoshaping, since experiments in first-order autoshaping have shown that more trials, but not more reinforced trials, are required to establish keypecking to a partially reinforced than are to a continuously reinforced S1 (Gibbon, Farrell, Locurto, Duncan, & Terrace, 1980).

However, it may be noted that there was no significant difference in acquisition between the two 100% groups despite the fact that there was a large difference in the length of interval between reinforced second-order trials. Ignoring the four S1-US trials common to both groups, the interval between S2-S1 trials averaged 120 and 480 sec in Groups 100%-12 and 100%-3, respectively. In first-order autoshaping, lengthening the intertrial interval on a continuously reinforced schedule facilitates the acquisition of keypecking to S1, particularly at the short end of interval values (Terrace, Gibbon, Farrell, & Baldock, 1975). It is not clear whether the present data represent a difference in the effects of intertrial interval in first- and second-order autoshaping or whether, in the present case, the difference between the interval durations was not sufficiently large to yield a behavioral effect.

Extended training. Responding to S2 during extended second-order training is summarized for each group in the left-hand panel of Figure 1. Rate of responding is plotted for successive blocks of 12 reinforced second-order trials; the means are based on responding during all S2 presentations in each block (12 presentations for the 100% groups, 48 presentations for the 25% group). Of course, the 100%-12 and 100%-3 groups differed in the temporal distribution of those trials (12/session vs. 3/session, respectively), and the 25% group is distinguished from the others in having 36 presentations of S2-alone intermixed among the 12 S2-S1 trials in each block.

There are two main features of the S2 data. In the 100% groups, responding to S2 increased to a maximum by the second block of trials and then declined

Figure 1. Rate of keypecking to S2 and to S1 in second-order sessions (see text for explanation).

precipitously during the remaining blocks. Second, responding to S2 increased more slowly and reached a lower maximum level in the 25% group, but did not show the steep decline characteristic of the 100% groups' performance. In fact, the 100% curves crossed over the 25% curve at about the fifth block.

Statistical analyses supported this summary of the data. The apparent changes in performance across trials and the differential effects of the reinforcement schedules evident in Figure 1 were indicated to be statistically significant in an overall analysis of the data. A repeated-measures analysis of variance with groups (100%-12, 100%-3, 25%) and blocks (14) as factors indicated both that there was a reliable change in performance across blocks [$F(13,403) = 14.2$, $p < .01$] and that the change occurred differentially across groups (Group by Block interaction [$F(26, 403) = 2.37$, $p < .01$]; the main effect of groups was not significant ($F < 1$). A separate analysis yielded no reliable difference between the two 100% groups: There was neither a main effect of groups nor a Group by Blocks interaction ($F_s < 1.06$); there was only a highly reliable effect of blocks. When the 25% group's data were compared with the combined 100% groups' data, both the Group by Block interaction and the block effect were significant [$F_s(13,416) = 3.4$ and 8.94 , respectively, $p_s < .01$]. Accordingly, the Group by Block interaction indicated by the overall analysis lies principally in the differential effects of the 100% and 25% reinforcement schedules.

The amount of training needed to establish keypecking to S2 varied across individual pigeons (see ranges in Table 1), although the acquisition data indicated that the groups did not differ reliably in number of reinforced trials required to meet the criterion for acquisition to S2. A separate analysis determined

whether individual differences in rate of acquisition influenced the results shown in Figure 1. The data were first corrected for differences in acquisition rate by designating that block in which each pigeon first met the acquisition criterion as its first block. In the resulting data, there were 10 successive postacquisition blocks common to all pigeons in the experiment. The analysis across these 10 blocks compared responding to S2 by the 25% group and the combined 100% groups. The outcome was similar to that obtained when the data from all blocks were used for each pigeon: The Group by Block interaction [$F(9,288) = 3.67$, $p < .01$] and the block effect [$F(9,288) = 2.49$, $p < .05$] were significant.

The apparent differential rate of decay in responding to S2 by the 25% and 100% groups during extended training was evaluated further in an analysis of slopes of regression lines fitted to the data of individual birds across Blocks 3 to 14. In computing these slopes, the rate of keypecking in each block was subjected to a transformation of the form $\log(\text{keypecks/min} + 1)$, which improved the linear fit. The difference between groups was statistically significant [$F(1,32) = 7.37$, $p < .05$], thereby confirming the more rapid loss of responding to S2 by animals trained on the 100% schedule.

The nature of responding to S2 on the two reinforcement schedules during extended training is described in greater detail in Figure 2. In the upper panel, the overall rate of keypecking to S2 has been broken down into rate of responding in successive 1-sec intervals of S2. The lower panel shows responding in each bin expressed as a proportion of the total; proportions were computed separately for each animal and then averaged. These within-CS distributions of responding are shown for representative

Figure 2. Representative distributions of keypecking during S2 presentations at various stages of second-order conditioning. Data are shown for the combined 100% groups and for the 25% group.

blocks throughout training for the combined 100% groups and for the 25% group. Two features of these data are notable. First, in the early blocks, the shapes of the distributions for the 100% and 25% groups were similar to those reported for continuous and partial reinforcement schedules in first-order auto-shaping (Gibbon et al., 1980). That is, response rate accelerated across the entire S2 period in the 25% group, but responding was relatively slowed as the end of the trial approached in the 100% group. Second, in the later blocks, the characteristic shape of the 25% distribution was reasonably well preserved, but the 100% distribution changed. This is seen most clearly in the proportion measure, in which, in the later blocks, the 100% groups' distribution tends to approximate the pattern of increasing acceleration shown by the 25% group. Other data collected in this laboratory indicate that an increasing gradient of response strength across an auto-shaped CS is often indicative of a relatively weak reinforcer for that CS (O'Connell & Rashotte, in press; Marshall, O'Connell, & Rashotte, Note 2). Consequently, one implication of these present data is that the change of shape in the 100% group's distribution results from the gradual weakening of the reinforcer for S2 as training proceeds.

Four other aspects of the data from this experiment should be noted. First, a subset of animals in the 100%-3 group had nonreinforced presentation of

a neutral stimulus, S3, presented in place of the S2-alone trials scheduled for the 25% group. There was an increase in the strength of responding to S3 in all these pigeons once S2-S1 pairings began, and two of the seven pigeons in this subgroup actually met the criterion for acquisition of keypecking (four or more trials with at least one peck in five consecutive S3 presentations). The maximum rate of responding to S3 occurred in the fourth second-order session (mean rate, 23.6 keypecks/min). Responding to S3 is attributable to generalization from S2 once it became a second-order excitator. In fact, responding was always stronger to S2 than to S3 in every pigeon and, after reaching its maximum rate in Session 4, responding to S3 quickly weakened to a low level for the rest of training (e.g., by Session 6 the mean rate had dropped to 2.9 keypecks/min). The S3-alone presentations had no consequences for responding to S2 that could be detected in statistical comparisons with the other 100%-3 subgroup.

Second, it is worth noting that the remarkable durability of second-order conditioning on the 25% schedule is not immediately obvious in the data presented in Figure 1. Those data were blocked so that 12 S2-S1 presentations comprise each data point. For the 25% group, the 14 blocks of second-order data plotted in Figure 1 represent 56 sessions of second-order conditioning, during which there were 672 S2 presentations. By the end of this training period, the

animals in Group 25% were still responding at about 20 keypecks/min. The durability of second-order conditioning in this group is exceptional by the standards of previous second-order conditioning data.

Third, the results discussed above are based entirely on the rate of keypecking measure. No important conclusions would be changed by considering the other measure that is commonly used in second-order autoshaping experiments, the percent of trials on which at least one keypeck occurred.

Finally, the right-hand panel of Figure 1 compares the strength of keypecking evoked by S1 when it occurred on the first-order S1-US trials and on the second-order S2-S1 trials in each block. It is clear that responding to S1 was suppressed sooner and to a greater degree on second-order trials in the combined 100% groups than in the 25% group. This pattern of results was confirmed in an analysis of variance of discrimination ratios based on responding to S1 on the two types of trials in each block. The ratio had the form $A/A + B$, where A represents responding to S1 on S1-US trials and B represents responding to S1 on S2-S1 trials. The analysis of ratios computed for each animal on every block yielded a significant effect of blocks [$F(13,416) = 9.38$, $p < .01$] and a significant Group by Block interaction [$F(13, 416) = 1.78$, $p < .05$]. The ability of S2 to suppress responding to S1 on S2-S1 trials provides one index of the inhibitory strength of S2. Using this measure, the present finding suggests that S2 became inhibitory more readily on the 100% schedule than in the 25% schedule.

DISCUSSION

The groups that were trained with a 100% reinforcement schedule in second-order autoshaping responded strongly to S2 early in training but subsequently lost responsiveness to S2 during extended training. A similar result has been reported in other conditioning preparations when a 100% reinforcement schedule is employed (Herendeen & Anderson, 1968; Pavlov, 1927, 1928; Rescorla, 1973). It has been demonstrated that S2 is actually inhibitory late in training in both autoshaping (Gokey & Collins, 1980; see Rashotte, 1981) and conditioned suppression (Rescorla, 1973). This pattern of results could imply that the S2-S1 signaling relation initially yields increments in the excitatory strength of S2 but later contributes to the loss in excitatory strength of S2 and to its ultimate inhibitory status. Such an implication is supported by the performance of the 25% reinforcement group which had the S2-S1 signaling relation degraded. With one exception, both the excitatory effect of the S2-S1 pairings early in training and the apparent inhibitory effect of the pairings late in training were diluted for that group. The exception was that the excitatory effect was not

weakened at the very outset of training when keypecking was established to S2 in approximately the same number of pairings in the 100% and 25% groups. Once keypecking had been acquired to S2, however, the pattern was clear: The 25% group at first responded more weakly to S2 than did the 100% groups; later in training, it responded more persistently to S2 and showed less suppression of responding to S1 on the S2-S1 trials than did the 100% groups.

Certain aspects of the data deserve particular comment. First, the failure of the 25% schedule to dilute the excitatory effect of the S2-S1 pairings at the very outset of training may simply be attributable to variability inherent in second-order acquisition data. On the other hand, the present finding parallels the effect of partial reinforcement on the acquisition of keypecking in first-order autoshaping (Terrace et al., 1975), as noted earlier, and suggests a commonality in the two orders of conditioning. It may be noted that in Marshall's (1976; Marshall & Rashotte, Note 1) original experiment there was even a slight enhancement of keypeck acquisition in the 25% group compared with the 100% group, but the small number of animals in that experiment makes the result problematical.

Second, the fact that S2 suppressed responding to S1 more readily in the 100% groups on the second-order trials possibly implies that S2 was more inhibitory in those groups, but this implication should be viewed with some caution. An alternative possibility is that the between-group differences in responding to S1 on S2-S1 trials could simply be a by-product of differences in responding to S2. That is, the 25% group is demonstrably more likely to respond to S2 during extended training and, therefore, is more likely to be at the response key when S1 is presented. The relatively high levels of responding to S1 by that group might be a simple consequence of "run-through" responding carried over from S2 or of increased response probability caused by the favorable positioning of the pigeon at S1 onset. A more complete assessment of the inhibitory status of S2 after extended training on partial reinforcement would be desirable. It is worth noting, however, that the relatively persistent responding to S2 in the 25% group may, itself, indicate that S2 is less inhibitory when partial reinforcement is employed.

Third, our interpretation of the within-CS keypeck distributions is based on other findings in first- and second-order autoshaping. There are now several reports that the within-CS distribution of keypecks tends to assume an increasing linear shape in first-order autoshaping when reinforcer strength is relatively low, for example, when the percentage of reinforced trials is low (Gibbon et al., 1980) or when reinforcement occurs on 100% of trials but US magnitude is low (O'Connell, 1980; O'Connell & Rashotte, in press). Comparable distributions are found in second-order autoshaping when S2 and S1 are paired on a

100% schedule, but when the second-order reinforcing effectiveness of S1 has been attenuated by pairings with low probabilities of US in first-order conditioning (Marshall, O'Connell, & Rashotte, Note 2). However, the within-CS keypeck distribution for relatively high-valued reinforcers tends to have a different shape: keypecking increases in the early segments of the CS, but shows no increase, or even a decrease, later in the CS, nearest the time of reinforcer presentation. These previous findings suggest the following interpretation of the keypeck distributions shown in Figure 2. Early in training, the shape of the 100% group's keypeck distribution was characteristic of a high-valued reinforcer, but as training progressed, the changing shape indicated a weakening of S1's reinforcement strength. The 25% group's distribution, on the other hand, displayed the shape characteristic of a low-valued reinforcer throughout training. This molecular analysis of second-order keypecking supports the idea that partial reinforcement during second-order autoshaping dilutes both the excitatory and the inhibitory consequences of the S2-S1 pairings. Incidentally, the differential shapes of the within-CS distributions that correlate with differences in reinforcer value probably reflect the differential involvement of a competing tendency to direct responses towards the US site. When reinforcer value is high, the tendency to direct responses to the food hopper rather than towards the CS is presumably strong late in the CS period (Boakes, 1977, 1979), and competition from this response tendency is probably responsible for the slowing of keypecking (Gibbon et al., 1980; O'Connell & Rashotte, in press).

Fourth, it is worth emphasizing that the 100% groups in this experiment provided comparisons in which both total trials and reinforced trials were equated with the 25% group. These comparisons indicated that the obtained effects of partial reinforcement are attributable to nonreinforced S2-alone trials.

Finally, the present experiment describes performance at only one level of partial reinforcement in the autoshaping preparation. We have preliminary unpublished data (see also Marshall, 1976) in which performance on a 50% schedule falls intermediate to the 100% and 25% cases reported here.

The present findings have implications for theoretical accounts of higher order conditioning. In particular, this experiment indicates that when the S2-S1 signaling relation is degraded by partial reinforcement, the growth *and* the subsequent decline in excitatory strength of S2 are retarded relative to the case in which S2 signals S1 on a 100% reinforcement schedule. Furthermore, the experiment identified nonreinforced S2 trials as the principal source of the retardation effects generated by the partial schedule. It seems, therefore, that a successful theory of second-

order conditioning will allow for these dual retardation effects of S2-alone trials.

There are currently two theoretical frameworks for second-order conditioning in which such a theoretical account could be developed. One, sketched briefly by Rescorla (1973), proposes that S2 will acquire separate excitatory *and* inhibitory tendencies during second-order training. The excitatory tendency results from the S2-S1 signaling relation; the inhibitory tendency results from the fact that S2 and S1, jointly, signal US omission. In this account, the strength of excitatory tendency determines the strength of response evoked by S2. The inhibitory tendency acts on S1 and, as the strength of that tendency increases, the principal consequence is that the reinforcing power of S1 is suppressed. This, in turn, causes the excitatory strength of S2 to extinguish. The final result is that S2 retains only its inhibitory tendency.

In the other framework, described by Rashotte (1981), a second-order trial of the sort employed in this experiment is conceptualized by the notation, $S2 \rightarrow (S2' + S1)$, where S2 represents a presentation of the second-order CS and the remaining symbols represent the immediate occurrence of the reinforcer, a simultaneous stimulus compound comprising S1 *and* the decaying sensory trace of S2 (i.e., S2'). In this framework, associative changes in S2 are principally determined by the reinforcement potential of the $S2' + S1$ compound, which is assumed to vary as training proceeds. Initially, that compound is assumed to have a high positive reinforcing value (which is responsible for S2's acquiring second-order excitatory strength), but gradually its reinforcement potential is reduced until finally it assumes a negative value (which is responsible for the ultimate inhibitory status of S2). The reinforcement potential changes principally through the contribution of S2' to the compound. At the outset of training, S2' is assumed to be neutral, but as training proceeds it becomes increasingly inhibitory. (This process can be understood as an instance of conditioned inhibition: S2', in conjunction with S1, signals that US will be omitted; S1 presented alone signals US. Under these circumstances, S2' should become a conditioned inhibitor.) The framework includes a unique rule for estimating the net reinforcement potential of the $S2' + S1$ compound. It is assumed that a weighted-sum rule operates such that the associative strengths of stimuli in the compound are summed, and inhibitory stimuli (i.e., S2') are always weighted more heavily than excitatory stimuli (i.e., S1) in the summation. This latter assumption allows the $S2' + S1$ compound to have the requisite positive value early in training and to achieve inhibitory reinforcing potential later in training.

In either of these theoretical frameworks, the retarding effect of partial reinforcement on second-order excitatory conditioning early in training is rea-

sonably easy to understand as a loss in excitatory strength resulting from nonreinforced S2 trials. A more substantive problem is posed by the apparent retardation of loss in excitatory strength by S2 later in training. It would seem that the nonreinforced S2 trials would have to weaken the inhibitory process identified in each framework as responsible for the loss in reinforcement potential (see Rashotte, 1981, for a more detailed account in that framework). Unfortunately, the small amount of data available indicates that inhibition is not weakened by nonreinforced presentations of an inhibitory CS (Zimmer-Hart & Rescorla, 1974). From these theoretical viewpoints, then, the present results point to the importance of further data on how nonreinforcement influences the associative status of an inhibitory CS.

It must be noted that it remains to be seen whether the present effects of partial reinforcement can be demonstrated in other second-order conditioning preparations. A basis for concern about the generality of the effects reported here is that partial reinforcement has an apparently unique enhancement effect on performance in first-order autoshaping in comparison to other first-order conditioning preparations (Gibbon et al., 1980).

Finally, unless the present effects of partial reinforcement prove to be unique to the autoshaping preparation, the data reported here imply that the exclusive use of continuous reinforcement schedules in previous research has contributed to the stereotype that second-order conditioning is a short-lived phenomenon. Even should it turn out that the present result is specific to autoshaping, however, the fact that a Pavlovian second-order conditioning contingency can maintain such durable control over directed motor actions would still have important implications for interpretations of performance in many situations. That is, because partial reinforcement schedules are a fact of life in most settings (e.g., Skinner, 1938), the possibility raised by these data is that second-order conditioning is a far more important factor in behavior control than has generally been acknowledged. In fact, the present 25% group's data probably constitute the most durable second-order conditioning that has been reported in the literature.

REFERENCE NOTES

1. Marshall, B. S., & Rashotte, M. E. *Partial reinforcement during second-order conditioning of the pigeon's keypeck*. Paper presented at the annual meeting of the Midwestern Psychological Association, Chicago, Illinois, May 1977.
2. Marshall, B. S., O'Connell, J. O., & Rashotte, M. E. Manuscript in preparation.

REFERENCES

- BOAKES, R. A. Performance on learning to associate a stimulus with positive reinforcement. In H. Davis & H. M. B. Hurwitz (Eds.), *Operant-Pavlovian interactions*. Hillsdale, N.J.: Erlbaum, 1977.
- BOAKES, R. A. Interactions between Type I and Type II processes involving positive reinforcement. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation: A memorial volume to Jerzy Konorski*. Hillsdale, N.J.: Erlbaum, 1979.
- GIBBON, J., FARRELL, L., LOCURTO, C. M., DUNCAN, H. J., & TERRACE, H. S. Partial reinforcement in autoshaping with pigeons. *Animal Learning & Behavior*, 1980, 8, 45-59.
- GOKEY, D. S., & COLLINS, R. L. Conditioned inhibition in feature negative discrimination learning with pigeons. *Animal Learning & Behavior*, 1980, 8, 231-236.
- HERENDEEN, D., & ANDERSON, D. C. Dual effects of a second-order conditioned stimulus: Excitation and inhibition. *Psychonomic Science*, 1968, 13, 15-16.
- MARSHALL, B. S. *Factors influencing acquisition and maintenance in second-order conditioning of the pigeon's key peck: Partial reinforcement*. Unpublished master's thesis, Florida State University, 1976.
- O'CONNELL, J. O. *Reinforcement magnitude effects in first- and second-order autoshaping*. Unpublished master's thesis, Florida State University, 1980.
- O'CONNELL, J. O., & RASHOTTE, M. E. Reinforcement magnitude effects in first- and second-order conditioning of directed action. *Learning & Motivation*, in press.
- PAVLOV, I. P. *Conditioned reflexes: An investigation of the cerebral cortex*. London: Oxford University Press, 1927.
- PAVLOV, I. P. *Lectures on conditioned reflexes: Twenty-five years of objective study of the higher nervous activity (behavior) of animals*. New York: Liverwright, 1928.
- RASHOTTE, M. E. Second-order autoshaping: Contributions to the research and theory of Pavlovian reinforcement by conditioned stimuli. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory*. New York: Academic Press, 1981.
- RASHOTTE, M. E., GRIFFIN, R. W., & SISK, C. L. Second-order conditioning of the pigeon's keypeck. *Animal Learning & Behavior*, 1977, 5, 25-38.
- RESCORLA, R. A. Second-order conditioning: Implications for theories of learning. In F. J. McGuigan & D. B. Lumsden (Eds.), *Contemporary approaches to conditioning and learning*. New York: Wiley, 1973.
- RESCORLA, R. A. *Pavlovian second-order conditioning*. Hillsdale, N.J.: Erlbaum, 1980.
- SKINNER, B. F. *The behavior of organisms: An experimental analysis*. New York: Appleton-Century, 1938.
- SNAPPER, A. G., VAN HAAREN, F., & INGLIS, G. B. *The SKED software system*. Kalamazoo, Mich: State Systems, 1978.
- TERRACE, H. S., GIBBON, J., FARRELL, L., & BALDOCK, M. D. Temporal factors influencing the acquisition and maintenance of an autoshaped keypeck. *Animal Learning & Behavior*, 1975, 3, 53-62.
- ZIMMER-HART, C. L., & RESCORLA, R. A. Extinction of Pavlovian conditioned inhibition. *Journal of Comparative and Physiological Psychology*, 1974, 88, 837-845.

(Received for publication August 21, 1980;
revision accepted November 19, 1980.)