Partial reinforcement effects on discrimination learning

BEN A. WILLIAMS

University of California, San Diego, La Jolla, California

Rats were trained on a series of reversals of a successive discrimination in which the percentage of S+ trials ending in food was varied. Changes in the discrimination index occurred more slowly with 50% reinforcement than with 100% reinforcement when the number of training trials was equated across conditions, but were approximately invariant when the conditions were equated with respect to the number of obtained reinforcements. Presentation of free reinforcement during the intertrial intervals reduced the overall rate of discrimination acquisition, but left this invariance unaffected. Invariance in reinforcements necessary to attain acquisition also occurred when different discriminations correlated with different percentages of reinforcement were intermixed within experimental sessions. The failure of the invariance effect to be disrupted by either manipulation suggests that previous accounts of the invariance effect in terms of "comparator" models of conditioning (e.g., Gibbon & Balsam, 1981) are inadequate.

Percentage of reinforcement is among the most frequently investigated behavioral variables. The great majority of studies have examined its influence on response vigor or persistence during extinction (e.g., partial reinforcement acquisition and extinction effects). Far fewer studies have examined its effects during discrimination learning, in part because such effects seem intuitively obvious. That is, the less frequently that responses to a positive stimulus (S+) are followed by reinforcement, the more training will be required for the discrimination to be learned. Several different studies appear to confirm this intuition (e.g., Bowman, 1963; Williams, 1976b).

Although discrimination is clearly retarded in terms of the number of training trials required for acquisition, there appears to be little effect of percentage of reinforcement when the measure is the number of reinforced trials. Eckerman (1969) studied the effects of nonreinforced responses to an S+ in a successive discrimination procedure in which the proportion of S+ presentations ending in food was varied from 0.16 to 1.0 and reinforcement never occurred during S- presentations. As expected, more training time was needed to achieve the discrimination for the smaller reinforcement probabilities. But with respect to the measure of reinforcements to criterion, the rate of discrimination was approximately constant across all of the probability-of-reinforcement conditions. Similarly, Gibbon, Farrell, Locurto, Duncan, and Terrace (1980) varied the probability of reinforcement in an autoshaping procedure and found that the number of reinforcements required to attain acquisition was approx-

imately equal regardless of the probability of reinforcement. Those authors also reviewed the effects of reinforcement percentage on other classical condition procedures, and noted that such invariance in the number of reinforcements to acquisition seemed to apply to most such procedures.

Both of the above studies used between-group comparisons, but a similar finding also occurs for within-subject comparisons. Williams (1981) studied percentage of reinforcement in a serial reversal procedure involving a simultaneous color discrimination, with highly experienced pigeons performing at their optimal level, and found that individual subjects trained on different percentages of reinforcement for successive blocks of reversals required approximately the same number of reinforcers to learn each new reversal, regardless of the reinforcement percentage. Thus, the invariance in reinforcements to acquisition appears to be a general effect.

Constancy in the number of reinforcements to acquisition is of interest because it appears to challenge any incremental model of learning that assumes that acquisition can be understood as some combination of excitatory and inhibitory effects (e.g., Rescorla & Wagner, 1972). That is, any occurrence of nonreinforced responding should decrease the existing level of response strength, so as the frequency of nonreinforcement of correct responses is increased, additional reinforcements should be required to counteract the decrement in response strength due to nonreinforcement. Gibbon et al. (1980, see their Footnote 7) provided an analytic solution for the application of the Rescorla-Wagner model to the effects of partial reinforcement on asymptotic response strength, which shows that response strength does indeed decrease monotonically with the percentage of reinforcement, with the degree of decrease dependent on the relative size of the rate parameters for incremental versus decremental condition-

This research was supported by NIMH Research Grant 5 RO1 MH42797-02 to the University of California, San Diego. Requests for reprints may be addressed to the author, Department of Psychology (C-009), University of California, San Diego, La Jolla, CA 92093.

ing (the beta parameter of the model). Thus, such a model apparently is challenged by the invariance effect.

Gibbon et al. (1980) offered the Gibbon-Balsam (1981) "comparator" model as an alternative explanation of the invariance effect. Their account assumes that the controlling variable is the ratio of the rate of reinforcement during the training stimulus relative to the rate of reinforcement in the situation as a whole. The comparator account is consistent with the invariance effect because that ratio is unaffected by changes in the percentage of reinforcement. That is, changing the percentage of positive trials ending in reinforcement reduces both terms in the ratio proportionally, because the reinforcement paired with the S+ is also the only source of reinforcement associated with the background cues.

In the present study, the invariance effect is further explored. Experiment 1 represents an attempt to extend its generality to a different training situation. The remaining experiments provide a test of the explanation of the effect in terms of the comparator model. The crucial issue is the effect of the percentage of reinforcement when the only source of reinforcement in the conditioning situation is that associated with the S+, a situation for which the comparator model predicts the invariance effect, compared to when other sources of reinforcement [e.g., free reinforcement during the intertrial intervals (ITIs)] are added to the situation. According to the comparator model, additional sources of reinforcement should destroy the invariance effect because the ratio of time to reinforcement during the S+ relative to reinforcement during the background no longer would remain constant. Instead, reductions in the percentage of reinforcement produces a proportionally greater reduction in the rate of reinforcement during the S+ than during the background. The predicted result is that proportionally more reinforcements should be required for acquisition with smaller percentages when background reinforcement is added.

EXPERIMENT 1

Experiment 1 extends the generality of the invariance effect along several dimensions. All previous reports have used pigeons, and stimuli were projected onto response keys, which suggests that the effect might occur only in situations in which Pavlovian contingencies play an important role. In contrast, in the present experiment, rats were used in a discriminated barpress procedure. Like the earlier study of Williams (1981), a serial reversal procedure was used in which a given subject received several reversals of the stimuli serving as S+/S-, with different reversals associated with different percentages of reinforcement. Unlike the earlier study, in which a simultaneous discrimination procedure was used, in the present study, a successive procedure was used in which the S+ and the S- were alternated over trials.

Method

Subjects. Four experimentally naive, albino Sprague-Dawley rats, approximately 6 months of age at the start of training, were housed

in a colony room with a light cycle of 14 h on:10 h off. Experimental sessions were conducted from 3–6 h after onset of the illuminated portion of the cycle, depending on the subject. Food deprivation was maintained by allowing free access to standard rat chow for 2 h per day, beginning approximately 5 min after the rats' removal from the experimental chamber.

Apparatus. The interior chamber consisted of a custom-built Plexiglas shell, 19 cm high, 25 cm wide, and 25 cm deep, equipped with a grid floor. Three walls and the ceiling were of clear glass; the front wall was painted black. Mounted on the front wall were two nonretractable stainless steel levers, 9 cm apart, edge to edge, each mounted 12 cm above the grid floor. Each lever was 3.2 cm in width, and 0.3 cm in thickness, protruded 1.9 cm into the chamber, and required a minimum force of 0.3 N for operation. Directly between the levers, 1.3 cm above the floor, was a recessed steel opening into which a liquid dipper (BRS/LVE Model SLD-002) could be inserted. The dipper, which nominally contained 0.01 cc of liquid, remained protruded into the chamber until activated, at which time it dropped into a tray of Mazola corn oil for 0.25 sec, and then was returned to the up position to allow consumption of the oil. The only light source in the chamber was provided by a 28-V miniature light bulb (Sylvania No. 28PSB) mounted 3.5 cm above the left lever. The light was either "bright" (850 fc as measured from the approximate location of the rat when pressing the lever) or "dim" (87 fc), as produced by placing a 400- Ω resistor in series. For sound attenuation, the interior chamber was placed within a larger exterior chamber equipped with a ventilating fan for masking noise.

Procedure. The rats were initially presented one session of magazine training in which presentations of the corn oil reinforcer were provided on a variable-time (VT) 60-sec schedule for a total of 50 reinforcements. In the next session, they were hand-shaped to press the left lever, with the corn oil as the contingent reinforcer. The right lever was never used at any time during the study and responses to it were not recorded. After shaping, the subjects received 100 continuously reinforced responses in the absence of either light intensity, at which time they were introduced to the bright light, which initially served as the S+. Over a series of four sessions, the schedule during S+ presentation was gradually increased from fixed-interval (FI) 5 sec to FI 15 sec. Reinforcement was presented for the first response after the FI requirement had elapsed. The reinforced response also terminated the S+ and produced a 1.5-sec ITI, during which the chamber was completely dark.

After initial exposure to the S+ alone, discrimination training was begun between the S+ (bright) and the S- (dim). The reinforcement contingency during the S+ remained FI 15 sec. During the S-, responses had no effect and the stimulus terminated independent of behavior after 15 sec had elapsed. On S+ trials that were nonreinforced, the S+ remained on until the FI 15-sec requirement had been met (i.e., its termination depended upon a response).

Two subjects were trained on the initial discrimination with all S+ trials ending in reinforcement (100%). Two other subjects were trained with 50% of the S+ trials ending in reinforcement and 50% in nonreinforcement. Training within a given session continued until 40 reinforcements had been presented. This meant that the 50% condition required twice as many trials as the 100% condition. Training on the initial problem continued until 90% of the total number of responses to both the S+ and the S- occurred to the S+, or for a total of 15 sessions.

After acquisition of the initial problem with bright as the S+ for all subjects, the reward values of bright and dim were reversed, and training was continued until the learning criterion was again attained. Three subsequent reversals were presented, each after the preceding reversal had been learned to criterion.

Two subjects were trained with a probability of reward of 1.0 for Reversals 1 and 2 and with a probability of 0.50 for Reversals

3 and 4. The remaining two subjects received the opposite order of probabilities of reinforcement.

Results

The results from the initial acquisition will not be presented because the use of only 2 subjects per condition prevents any meaningful comparison between the 100% and 50% conditions. The mean results for the acquisition of the subsequent four reversals, averaged over subjects and reversals, are presented in Figure 1. Note that the abscissa is the number of reinforcers, not the number of trials, since the 50% condition required twice as many trials for each block of training that is shown. Because training on some reversals was terminated before the full 15 sessions of training had been presented, whenever a subject reached a 90% discrimination criterion it was assigned a score of 90% for all remaining sessions in order to have a complete set of 15 sessions for the purpose of data analysis. This procedure had the result of producing an asymptotic discrimination performance at slightly below the 90% level. It is apparent from Figure 1 that both percentage-of-reinforcement conditions produced a regular improvement in discrimination performance and, most importantly, that the acquisition functions were virtually identical.

The similarity in the acquisition functions might partly be due to the termination of training whenever each reversal was learned to the 90% criterion and the assignment of a 90% score for the remaining sessions that were not run. That this was not the case can be seen by considering only the data prior to the point at which any subject reached the 90% criterion. The earliest this occurred for any subject was during Session 10, which is represented in Figure 1 as a part of the Block 4. Thus, Blocks 1-3 were not affected by the learning criterion that was used. Although Block 2 showed a small difference in favor of the 50% condition, an analysis of variance (ANOVA) restricted to Blocks 1-3 revealed no evidence of a signifi-

Figure 1. Discrimination accuracy during Experiment 1, averaged over Reversals 1-4. The ordinate is the number of responses to the S+ divided by the total responses to the S+ and the S-.

Figure 2. Discrimination accuracy during Experiment 1 as a function of the number of trials.

cant effect of reinforcement percentage or of the interaction between blocks and reinforcement percentage (both $F_s < 1$).

Figure 2 shows the acquisition results as a function of the number of training trials, rather than the number of reinforcers. Each block of training corresponds to one session for the 50% condition and two sessions for the 100% condition. Thus, only the first half of training is presented for the 50% condition. Unlike Figure 1, a consistent difference is apparent throughout training, as the discrimination was learned more rapidly for the 100% group. These data were tested with a two-way ANOVA (probability of reinforcement $\times$ blocks). The effect of probability of reinforcement approached, but did not attain, statistical significance [$F(1,3) = 8.56, p = .06$]. The interaction term ($F < 1$) did not approach significance, but the effect of blocks was significant [$F(6,18) = 40.3, p < .05$].

Discussion

The results of Experiment 1 extend the generality of the invariance in the number of reinforcements to acquisition by demonstrating the phenomenon with a procedure that differed from previous studies on several dimensions. Perhaps most notable is that all previous reports of the effect have reported only the number of reinforcers required to reach some particular learning criterion, whereas the present experiment shows that the effect appears to apply to the entire learning function. As shown in Figure 1, the acquisition functions were virtually identical at all points in training. Apparently, the rate of acquisition depends only on the number of reinforcers, not on the percentage of trials ending in reinforcement.

EXPERIMENT 2

The critical test of the comparator account of the invariance effect is whether it is destroyed when a source of reinforcement other than that associated with the S+ is added to the conditioning situation. Given such an ad-

ditional source, the comparator account predicts that the rate of acquisition would be retarded for all percentages of reinforcement, but more so the smaller the reinforcement percentage. This can be seen clearly by examining the parameters of reinforcement used in Experiment 2. Trial periods were 15 sec in duration, and subjects received either a 100% or a 50% reinforcement schedule during the S+. A 15-sec ITI separated trials, which meant that the average time between reinforcements in the situation as a whole was 60 sec for the 100% condition and 120 sec for the 50% condition, assuming, as in Experiment 1, all reinforcement in the situation occurred on the S+ trials. The corresponding times to food reinforcement during the S+ would be 15 sec and 30 sec. Thus, the ratio of reinforcement rate to the background during the S+ would be 4:1 in both cases. The critical change from Experiment 1 was the addition of a VT 30-sec schedule that provided free food during the ITI. Thus, the average time to food reinforcement became 30 sec for the 100% condition and 40 sec for the 50% condition. Since the rate of food during the S+ was unaffected, the corresponding ratios of S+ reinforcement rate to background rate became 2.0 and 1.33. The smaller ratio for the 50% condition implies that its rate of acquisition should be differentially retarded.

Method

Subjects and Apparatus. Four naive subjects similar to those used in Experiment 1 served in the present experiment. The apparatus was the same as in Experiment 1.

Procedure. The same pretraining regimen was used as in Experiment 1. After all subjects were responding in the presence of the S+ (bright) on an FI 15-sec schedule, the ITI was gradually extended to 15 sec. A discrimination procedure was instituted in which the S+ and the S- (dim) were randomly alternated within sessions. Training on the initial discrimination continued until 90% of the responses during the S+ and the S- were occurring to the S+. Training on a particular session continued until 30 reinforcers had been earned, and this criterion for terminating a session continued for the remainder of training. Two subjects received reinforcement on 100% of S+ trials; the remaining 2 received reinforcement on 50% of the trials.

After all subjects had reached the discrimination criterion for the original problem, three additional sessions were given in which a VT 30-sec schedule operated during the ITIs, which continued for the duration of the experiment. After these three sessions, training on the first reversal of the discrimination commenced, with the dim light now serving as S+ and the bright light as S-. Training on this problem again continued either until a 90% criterion had been reached or for a total of 20 sessions. Three subsequent reversals were then presented, each trained to the same criterion. Two subjects received the 50% schedule during Reversals 1 and 2, and then the 100% schedule during Reversals 3 and 4. The other 2 subjects received the opposite order of probabilities of reinforcement.

Results

The acquisition of the brightness discrimination, averaged over the four reversals, is shown in Figure 3. Any subject reaching the 90% criterion level before the completion of 20 sessions was assigned a score of 90% for all remaining sessions. Figure 3 shows that the acquisition functions for the different percentages of reinforcement were highly similar throughout training. The similar-

Figure 3. Discrimination accuracy during Experiment 2 as a function of the number of reinforcers.

ity cannot be ascribed to the use of the 90% learning criterion, since the earliest any subject reached that criterion was during Training Session 11, which corresponds to Block 6 of training shown in Figure 3. The two acquisition functions were more similar prior to this point than afterwards. Regardless of the point of training, any difference between the functions was in favor of the 50% condition, contrary to the comparator hypothesis. When the functions shown in Figure 3 were submitted to a two-way ANOVA (blocks $\times$ percentage of reinforcement), the blocks variable was highly significant ($p < .01$), but the percentage of reinforcement variable was not significant [$F(1,3) = 1.79$]. The interaction term was also not significant ($F < 1$).

Figure 4 shows the same results as a function of the number of training trials. Once again, this meant that only the first half of the trials of the 50% condition is included. As in Figure 2, discrimination performance was superior for the 100% condition throughout training. The results shown in Figure 4 were submitted to a two-way ANOVA

Figure 4. Discrimination accuracy during Experiment 2 as a function of the number of trials.

(blocks $\times$ percentage of reinforcement). The effect of blocks was significant [$F(9,27) = 77.6, p < .01$], as was the effect of percentage of reinforcement [$F(1,3) = 22.6, p < .05$]. However, the interaction was not significant [$F(9,27) = 1.61, p > .10$].

Discussion

Figure 3 demonstrates that the invariance effect occurs even when free reinforcement is delivered during the ITI. This finding is directly contrary to the predictions of the comparator account of the effect (Gibbon et al., 1980), which entails that slower acquisition should have occurred with the 50% condition when the free reinforcement was added to the situation. The failure to support this prediction cannot be attributed to insensitivity in the procedure, since a significant effect, in favor of the 100% condition, was obtained for acquisition as a function of the number of training trials.

The applicability of the Gibbon-Balsam version of the comparator model to the present learning situation can, of course, be challenged. Their model was developed for Pavlovian conditioning procedures, primarily autoshaping. The present task, in contrast, involved not only an operant discrimination, but reversals of that discrimination was well. Despite these differences, there are good reasons for regarding the present procedure as a meaningful test. A large amount of research suggests that the effects of free reinforcement in operant procedures are qualitatively similar to those in Pavlovian procedures (e.g., Rachlin & Baum, 1972; Williams, 1983), since in both cases, the noncontingent reinforcement decreases the effectiveness of the conditioning contingency to maintain behavior. Since the invariance effect has been demonstrated in both Pavlovian and instrumental learning situations, it seems plausible that its determinants have generality across different types of conditioning.

The application of the comparator concept to reversal learning is more problematic. Gibbon et al. (1980) argued persuasively that the ratio of S+ reinforcement rate to background reinforcement rate determined not only the rate of acquisition in response to the S+, but also the rate of extinction. If it is assumed that reversal learning can be understood as a combination of acquisition and extinction, the implication is that the comparator concept should apply to reversal learning as well. However, there is no consensus about the determinants of reversal learning. Some accounts argue that the same mechanisms are involved as in simple acquisition and extinction, but with the additional complication of effects of proactive interference (Gonzalez, Behrend, & Bitterman, 1967; Woodard, Schoel, & Bitterman, 1971). A second, competing account assumes that an additional component is involved, consisting of a conditional discrimination based on the outcome of the preceding trial (Williams, 1976b). Still other accounts hypothesize changes in attentional factors (Mackintosh, 1974, chap. 10). All of these different accounts

assume that the basic effects of reinforcement and non-reinforcement are involved, with the implication that simple theories of associative learning should apply to the reversal learning situation, but with changes in various parameter values. Even if that assumption is denied, the demonstration that the invariance effect does occur with the reversal learning procedure, both here with a successive discrimination procedure and previously with a simultaneous discrimination (Williams, 1981), implies that any explanation of the invariance effect must encompass the serial reversal procedure.

EXPERIMENT 3

Before rejecting the comparator hypothesis as an account of the invariance effect, it is important to establish the generality of the effect. In particular, it is not clear how the addition of the VT reinforcement during the ITIs affected discrimination performance, because several procedural differences between Experiments 1 and 2 prevent a direct comparison. Consequently, Experiment 3 was a replication of both Experiments 1 and 2 within the same study. Such a comparison not only allows a determination of the robustness of the invariance effect, but also an assessment of whether the free reinforcement had any impact on discrimination performance.

Method

The subjects were the 8 rats from both Experiments 1 and 2. The apparatus was the same as in the earlier experiments. All subjects were continued on the reversal procedure used in Experiment 2. Trials consisted of 15-sec stimulus presentations separated by 15-sec ITIs during which the VT 30-sec schedule would or would not operate, depending upon the experimental condition. Subjects 1-4 (from Experiment 1) were presented Reversals 5-8 (Reversals 1-4 had been presented in Experiment 1) with the VT in the ITI, and then Reversals 9-12 with no VT schedule in effect. Subjects 5-8 (from Experiment 2) received Reversals 5-8 with no VT schedule and then reversals 9-12 with the VT. Each reversal was again trained to a 90% discrimination criterion or for 20 sessions. Each session was continued until 30 reinforcers were earned by responses to the S+.

The reinforcement percentage assigned to each reversal was varied across subjects. For 4 subjects, reversals 5-6 and 11-12 were assigned the 50% schedule, and Reversals 7-10 the 100% schedule. The opposite assignment occurred for the remaining 4 subjects. This provided a 2×2 factorial design in which reinforcement percentage was one factor and the presence/absence of the VT schedule during the ITI was the second factor.

Results

Figure 5 shows the acquisition function averaged over the four reversals of each condition. As in Experiment 1, when no VT schedule occurred in the ITI, the acquisition functions for the 50% and the 100% schedules were highly similar throughout training. As in Experiment 2, the acquisition functions when the VT schedule was added were also similar, although there was slightly better discrimination for the 50% schedule, especially during

Figure 5. Discrimination accuracy during Experiment 3 as a function of the number of reinforcers.

the later discrimination sessions. For both reinforcement percentages, acquisition was generally retarded when the VT schedule was added during the ITI.

The data shown in Figure 5 were analyzed with a three-factor within-subject ANOVA (blocks $\times$ percentage of reinforcement $\times$ ITI reinforcement). The blocks variable was significant [$F(4,28) = 113.3, p < .01$]. Percentage of reinforcement was not significant [$F(1,7) = 1.35, p > .05$] whereas the effect of the presence/absence of the VT schedule was significant [$F(1,7) = 14.70, p < .01$]. The interaction between VT and percentage of reinforcement was not significant ($F < 1$). The interaction between percentage of reinforcement and blocks of training did not attain significance [$F(4,28) = 2.53, p > .05$], nor did the interaction between blocks and the VT schedule [$F(4,28) = 2.21, p > .05$]. Thus, the statistical analysis confirms that there was no difference between the acquisition functions for the different percentages of reinforcement, and no differential effect of the added free reinforcement on acquisition under the different percentages of reinforcement. However, the addition of the VT reinforcement did cause an overall retardation of the discrimination.

Because a large number of subjects reached the learning criterion between 10 and 15 sessions of training, at least during some reversals, the similarity between the two functions for the last two points of the functions shown in Figure 5 may be due partly to a criterion artifact. To remove this influence, a second ANOVA was performed on only the data from the first three blocks of training (the first 12 sessions). The effect of the blocks variable was again significant [$F(2,14) = 113.5, p < .01$], as was the effect of the presence/absence of the VT reinforcement [$F(1,7) = 9.91, p < .05$]. But the effect of probability of reinforcement was not significant ($F < 1$). The only interaction term that was significant was that between probability of reinforcement and blocks [$F(2,14) = 4.53, p < .05$], which indicates that the change in discrimination performance was greater for the 50% condition. In-

spection of Figure 5 suggests that this was due to poorer discrimination performance during the first few sessions of training, followed by better discrimination performance during the middle sessions, although a test for simple effects revealed that there was no difference between the two reinforcement probabilities for any of the first three blocks of training.

Figure 6 shows the results of Experiment 3 as a function of the number of training trials (and also the number of S- trials, which were always half the number of training trials). Consequently, each block of 120 trials shown in Figure 6 corresponds to two sessions for the 100% conditions and only one session for the 50% conditions. Data from the last several sessions of training are excluded, because of the role of the criterion artifact noted above (for the 100% condition). Unlike Figure 5, there is considerable overlap between the two middle functions, but with the relative ordering of the 100% and 50% conditions reversed. Both reinforcement percentages produced slower acquisition when the VT reinforcement was added, but for both levels of that variable, acquisition was slower for the 50% condition. An ANOVA on the data shown in Figure 6 revealed that the effect of probability of reinforcement was significant [$F(1,7) = 24.0, p < .01$], as was the effect of blocks [$F(6,42) = 104.2, p < .01$]. The effect of the presence/absence of the VT reinforcement approached, but, did not attain, significance [$F(1,7) = 5.49, .06 > p > .05$], but the interaction between the VT variable and blocks was significant [$F(6,42) = 2.69, p < .05$]. No other interactions were significant.

Discussion

The results of Experiment 3 provide still stronger evidence against the comparator account of the invariance effect. Contrary to that account, the invariance effect was not abolished by the addition of free reinforcement during the ITI, despite the fact that the free reinforcement did retard the rate of discrimination acquisition in general. As in Experiment 2, any differential effect of the free reinforcement was in the opposite direction from that

Figure 6. Discrimination accuracy during Experiment 3 as a function of the number of trials.

predicted, since the rate of discrimination was slightly (but nonsignificantly) greater for the 50% condition when the free reinforcement was added.

The general retardation of discrimination acquisition by the addition of the ITI reinforcement is worth consideration in its own right. One possible interpretation is suggested by the previous results of Williams (1977), who demonstrated that the rate of acquisition of a simultaneous discrimination was inversely related to the rate of reinforcement produced by a second, already learned discrimination that was interspersed over trials. Such a contrast effect was interpreted as showing that relative rate of reinforcement applies to discrimination acquisition in a fashion similar to measures of response rate more typically used to study contrast. Such an interpretation could be applied here as well, on the assumption that the addition of free reinforcement during the ITI increased the overall context of reinforcement, thus causing the effectiveness of the contingent reinforcers to be diminished. The major difficulty with such an interpretation is that it, like the comparator hypothesis, entails that the addition of the VT schedule should have differentially retarded the 50% reinforcement condition. That is, the contrast interpretation assumes that the relative rate of reinforcement contingent on the S+ is the controlling variable, and the relative rate of reinforcement is clearly lower with the 50% condition when the VT schedule was superimposed. The question posed is how the present data are to be reconciled with the previous results of Williams (1977), which suggest that relative rate of reinforcement was the controlling variable.

An alternative interpretation of the present VT data is in terms of the noncontingent nature of any delayed reinforcement effect caused by the ITI reinforcers. Such reinforcers could occur at any time during the ITI, including shortly after the offset of the trial. Responses near the end of a trial would thus be followed by these ITI reinforcers, and any delayed reinforcement effect could strengthen behavior during the S-, thereby reducing the level of discrimination. However, such an interpretation, although plausible, relies heavily upon the concept of superstitious conditioning, the validity of which remains controversial (e.g., Staddon & Simmelhag, 1971; Timberlake & Lucas, 1985). It is also unclear why such superstitious reinforcement effects should produce the invariance effect with the added VT reinforcers.

EXPERIMENT 4

The results of Experiments 1-3 extend the generality of the finding of invariance in the number of reinforcements necessary to attain acquisition to still another situation. The effect has thus been demonstrated in successive discrimination learning (Eckerman, 1969), autoshaping and other classical conditioning preparations (Gibbon et al., 1980), and serial reversal learning using both successive and simultaneous discrimination procedures. This generality is challenged, however, by the results of Papini and Overmier (1985), who reported the effect in auto-

shaping, but only under circumscribed conditions. They compared the acquisition rates of pigeons trained on either 100% or 25% reinforcement schedules, while holding either the number of trials per session (resulting in fewer reinforcers/session for the 25% group) or the number of reinforcers per session (thus producing more trials/session for the 25% group) constant across groups. They found that the invariance effect occurred only when the number trials per session was held constant. The apparent reason for the failure of the effect when the number of reinforcers per session was held constant is that acquisition was generally retarded with more trials per session. This trials effect occurred for both the 100% and the 25% conditions. Papini and Overmier noted that their results were contrary both to the comparator theory of Gibbon and Balsam (1981) and to incremental learning theories such as that of Rescorla and Wagner (1972).

The implications of the results of Papini and Overmier (1985) for the present study are uncertain because the invariance effect was obtained here under conditions in which they failed to obtain the effect. That is, the 50% versus 100% conditions in Experiments 1-3 were equated with respect to the number of contingent reinforcers presented per session, not the number of trials. Nevertheless, because their results suggest that there may be an interaction between the invariance effect and the trials per session variable, Experiment 4 was conducted in order to determine whether the invariance effect still occurred with the present procedure when the different groups of subjects were equated with respect to the number of trials per session.

A second purpose of Experiment 4 was to test more rigorously whether the invariance effect could be abolished by the addition of background reinforcement. In Experiments 2 and 3, the additional reinforcement supplied by the VT schedule during the ITI did produce a disparity between the comparator values (see introduction of Experiment 2), but the disparity was not large (2.0 vs. 1.33). To increase that disparity, the ITI in Experiment 4 was extended from 15 to 30 sec, so that the free reinforcers presented during the ITI would constitute a larger fraction of the total number of reinforcers.

Method

Subjects. Four experimentally naive subjects, maintained in all respects like those in Experiments 1-3, served as subjects in Experiment 4.

Apparatus. The apparatus was the same as in Experiment 1, with two modifications. A noise stimulus, 10 dB above the ambient noise level in the chamber (approximately 75 dB SPL) was added to the presentation of the dim light. The noise was presented through an overhead speaker attached to the center portion of the chamber ceiling. The second change was that the dipper presenting the corn oil was replaced with a food-pellet dispenser. Pellets (45-mg Noyes pellets) were presented into a pellet receptacle located directly in the center of the intelligence panel, approximately 2 cm above the grid floor.

Procedure. All subjects were pretrained on a continuous reinforcement schedule in the presence of the initial S+ (the bright light), and the reinforcement schedule during that stimulus was gradually

extended to VI 15 sec. After each reinforcer was delivered, an ITI, initially 3 sec in duration, was gradually extended over the course of three sessions to 30 sec. On the following session, all subjects were begun on the discrimination procedure involving the interspersal of the S+ and the S- (the dim light in combination with the noise). The ITI continued to be 30 sec (in contrast to the 15 sec of Experiments 1-3) and the average trial duration continued to be 15 sec. However, unlike Experiment 1, in the present experiment, the trial duration was variable in length, defined by the VI reinforcement schedule. This was also true for S- trials, which terminated automatically after a VI interval had elapsed.

The percentage of S+ trials ending in reinforcement for the initial discrimination and the first three reversals was 50% for 2 subjects, followed by Reversals 4-5 with a 100% schedule, Reversals 6-7 with a 100% schedule with the VT schedule during the ITI, reversals 8-9 with the 50% schedule and the VT in the ITI, Reversals 10-11 with the 50% without the VT, and Reversals 12-13 with the 100% schedule without the VT. The order of presentation for the remaining 2 subjects was reversed with respect to the 50% versus the 100% schedules, but these also received the VT schedule during the ITI during Reversals 6-9. As in Experiments 1-3, a non-reinforced S+ trial ended only after a barpress occurred after the VI interval had elapsed.

Individual sessions terminated after 80 trials. This meant that the 50% schedule produced 20 reinforcers per session on average, whereas the 100% schedule produced 40 reinforcers. Training on a particular reversal continued until a 90% discrimination accuracy was attained, or for 20 sessions of training.

Results

Figure 7 shows the acquisition rates for the four different conditions presented to each subject: the two percentage-of-reinforcement conditions with and without the VT schedule added in the ITI. The data shown are the averages of all reversals presented with each condition, with the exclusion of the initial acquisition of the discrimination and the first reversal. Analysis of the differences in acquisition rate across reversal training (e.g., Reversals 2-3 vs. Reversals 12-13) showed no systematic differences.

The acquisition functions shown in Figure 7 are less clearly separated than those in Figure 5 from Experiment 3, despite the designs of the studies being highly similar. This was because there was considerably greater

Figure 7. Discrimination accuracy during Experiment 4 as a function of the number of reinforcers.

Figure 8. Discrimination accuracy during Experiment 4 as a function of the number of trials.

difference between the 50% and 100% schedules, both without the VT (top two functions) and when the VT schedule was added (bottom two functions). The result was that the 100% condition without the VT overlapped substantially with the 50% condition with the VT. Despite this overlap, the results of the quantitative analysis, by a three-factor within-subject ANOVA (blocks $\times$ percentage of reinforcement $\times$ ITI schedule) were generally similar. The effect of the blocks variable was significant [$F(4,12) = 227.8, p < .01$], and the effect of the presence/absence of the VT schedule in the ITI was also significant [$F(1,3) = 41.6, p < .01$]. The effect of probability of reinforcement was not significant [$F(1,3) = 5.26, p > .10$], but the interaction between the probability of reinforcement and blocks of training was significant [$F(4,12) = 3.78, p < .05$]. This interaction was not due to an increasing difference between the two conditions over training, since an analysis of the simple effects composing this interaction showed that only during Block 2 was there a significant difference between the two probability-of-reinforcement conditions.

The results for the first half of training were analyzed with a separate ANOVA in order to reduce the influence of the training terminating after differing numbers of sessions because of the learning criterion that was used. The effect of blocks was again significant [$F(4,12) = 77.75, p < .01$], as was the effect of the VT schedule [$F(1,3) = 29.61, p < .05$]. In addition, the effect of probability of reinforcement was also significant [$F(1,3) = 10.9, p < .05$], as was the interaction between probability of reinforcement and blocks [$F(4,12) = 3.68, p < .05$]. None of the other interactions approached significance. Thus, when the results are restricted to the first half of the data shown in Figure 7, the invariance effect was no longer obtained. Instead, faster acquisition occurred with the 50% reinforcement probability, but this difference was not systematically affected by the addition of the VT schedule.

Figure 8 shows the results as a function of the number of training trials. The overlap was similar to that in

Figure 7, but the order of the 50% and 100% conditions was reversed. Faster acquisition occurred for the 100% than for the 50% probability of reinforcement, and for both probabilities, the rate of acquisition was reduced by the addition of the VT schedule. An ANOVA of these data showed that the effect of blocks was significant [$F(3,12) = 269.3, p < .05$], as was the effect of VT schedule [$F(1,3) = 122.3, p < .01$] and probability of reinforcement [$F(1,3) = 23.69, p < .05$]. The interaction between probability of reinforcement and the VT schedule did not approach significance ($F < 1$), nor did that between probability of reinforcement and blocks ($F < 1$), or any of the remaining interactions.

All of the results presented heretofore have been in terms of discrimination accuracy. This measure comprises two separate components, response rate during the S+ and response rate during the S-. It is of some interest, therefore, to determine how the percentage of reinforced S+ trials affected these measures separately. Figure 9 shows the response rates per session to each stimulus. The top panel shows the behavior in response to the S+. Note that each datum point for both 50% reinforcement conditions represents the mean behavior for two sessions, whereas those for the 100% conditions represent the behavior for single sessions. This was because only 20 contingent reinforcers per session occurred for the former conditions, whereas 40 reinforcers per session occurred for the latter conditions. Figure 9 shows that there was a small increase in responding to the S+ with extended training for all conditions, although for the conditions without the VT schedule, most of the increase occurred during the first three blocks of training. The effect of the VT schedule was to suppress responding for both percentage-of-reinforcement conditions at all stages of training. In addition, Figure 9 shows that the 100% condition produced more behavior in response to the S+ than did the 50% condition, both with and without the addition of the VT schedules. These observations were tested statistically with a three-factor ANOVA (blocks $\times$ percentage of reinforcement $\times$ presence/absence of VT). The effect of the VT schedule was significant [$F(1,3) = 67.9, p < .01$], as were the effects of blocks of training [$F(9,27) = 4.17, p < .01$]. But the effect of probability of reinforcement was not significant [$F(1,3) = 2.84, p > .10$]. None of the interactions approached significance.

The middle panel of Figure 9 shows the behavior in response to the S-. As in the top panel, the blocks of training correspond to number of reinforcers presented, not to the number of trials. Thus, each datum point for the 50% conditions corresponds to the average of twice the number of trials as for the 100% conditions. In general, there was a regular decrease in responding across blocks of training, and the presence of the VT schedule generally suppressed behavior. In addition, response rate to the S- was generally lower with the 50% conditions than with the 100% condition. These observations were also subjected to a three-way ANOVA. The effect of the

Figure 9. Response totals from Experiment 4 as a function of the duration of training. The top graph shows response rates during the S+; the middle graph shows response rates during the S- when the conditions were equated with respect to the number of reinforcers; the bottom graph shows response rates during S- when the conditions were equated with respect to the number of trials.

VT schedule was significant [$F(1,3) = 29.4, p < .02$], as was the effect of blocks of training [$F(9,27) = 21.9, p < .01$]. The interaction of VT schedule with blocks of training was also significant [$F(9,27) = 22.3, p < .01$]. However, the effect of percentage of reinforcement was not significant [$F(1,3) = 3.89, p > .10$]. None of the other interactions approached significance.

The bottom portion of Figure 9 again shows behavior in response to the S-, but here with the conditions equated with respect to the number of S- trials rather than to the number of S+ reinforcers. The result of this change is that the 100% and the 50% conditions were now more similar, both with and without the VT schedule, although there is some evidence that the 50% condition produced higher response rates in the later stages of training. The

results of the three-way ANOVA were that the VT schedule was significant [$F(1,3) = 60.1, p < .01$], as was the effect of blocks of training [$F(4,12) = 22.0, p < .01$] and the interaction between blocks and the VT schedule [$F(4,12) = 20.7, p < .01$]. No other effects approached significance (all remaining F values < 1).

Discussion

Experiment 4 differed in several respects from Experiment 3, including the nature of the reinforcer, the difficulty of the discrimination, the length of the ITI, and the number of trials per session. The major difference in outcome was that Experiment 4 produced a small but reliable superiority in the acquisition rate for the 50% condition, a superiority that occurred regardless of whether the VT schedule of free reinforcers operated during the ITI. The results are thus partially consistent with those of Papini and Overmier (1985), who found faster conditioning when fewer reinforcers were presented per session. However, the present results differ from those of Papini and Overmier in showing an invariance effect when the different probabilities of reinforcement were equated with respect to the number of reinforcers per session (Experiment 3), while failing to show an invariance effect when equated with respect to the number of trials per session (Experiment 4). Papini and Overmier obtained the opposite pattern. But because of the various procedural changes, it remains unclear whether the difference between Experiments 3 and 4 was due solely to the number of reinforcers per session.

Despite the difference in acquisition rates obtained for the 50% versus the 100% condition, the results of Experiment 4 provide still stronger evidence against the comparator account. As discussed above, that hypothesis entails that additional reinforcers during the ITI should affect the 50% reinforcement condition to a greater degree, because the added reinforcement destroys the proportional change in the time to reinforcement during the S+ versus the background when the percentage of reinforcement is varied. As applied to the present data, the ratio of reinforcement rate to the background during the S+ was 6:1 for the 50% and 100% conditions without the VT schedule in the ITI, but was reduced to 2:1 for the 100% condition and 1.2:1 for the 50% condition when the VT schedule was added. The result should have been a greater decrement in discrimination acquisition for the 50% condition. Contrary to that prediction, the decrement caused by the VT schedule was essentially similar for both percentage-of-reinforcement conditions.

The reasons for the general retardation of discrimination acquisition by the VT schedule remain unclear. Figure 9 reveals that the VT schedule suppressed responding to both the S+ and the S-. In order for discrimination accuracy also to be reduced, the percentage of suppression had to be greater during the S+ than during the S-. However, such an effect is contrary to a substantial body of literature on "resistance to change." Nevin and his colleagues (Nevin, 1974, 1979; Nevin, Mandell, &

Atak, 1983; Nevin, Mandell, & Yarensky, 1981) have shown that response rate is less reduced by a variety of manipulations, including schedules of free reinforcement, when the response is maintained by higher rates of reinforcement. Assuming that the S- in the present case is functionally correlated with reward at a lower rate than is the S+, the implication is that the VT schedule should have decremented the response rate more in the S-, with the result that the discrimination ratio should have increased. The fact that a decrease occurred in discrimination performance suggests that the notion of resistance to change may not easily be applied to discrimination procedures.

EXPERIMENT 5

The use of VT schedules during the ITI as a test of the comparator explanation of the invariance effect assumes that the subject "calculates" the background rate of reinforcement by simply averaging over the reinforcers received on S+ trials and those received during the ITI. However, there are now several studies that demonstrate that such an assumption is simplistic, even with Pavlovian conditioning procedures. With autoshaping, for example, extra reinforcers during the ITI have a differential effect depending upon whether they are signaled by a different stimulus than the CS. When unsignaled, such reinforcers typically prevent the acquisition of keypecking; when signaled, their effects on acquisition are much weaker (Durlach, 1983; Goddard & Jenkins, 1987; Grau & Rescorla, 1984). Such results, and others (Brandon, 1981; Farley, 1980; Reilly & Schachtman, 1987; Schachtman & Reilly, 1987; Williams, 1976a) suggest that either the comparator hypothesis is wrong, or the method of calculating background reinforcement rate is more complex than simply averaging over all sources of reinforcement in an experimental session.

Regardless of how background reinforcement is calculated, the comparator hypothesis can be tested by intermixed presentations of the different percentages of reinforcement within the same experimental session. Given that the same background conditions operate throughout a session, and that the reinforcers obtained on 50% and 100% trials both contribute to that background, the impact of the response-contingent reinforcers during the S+s, for both percentages of reinforcement, should be determined relative to that background rate. And since the rate of reinforcement for the S+ of the 50% condition is half that of the S+ for the 100% condition, the acquisition of the 50% condition should be slower. In Experiment 5, this prediction was tested by intermixing two different discriminations in a procedure similar to that used in Experiments 1-4.

The two different discriminations were a visual discrimination between two light sources (which also differed in intensity) and an auditory discrimination involving a tone versus a white noise. On some reversals, the visual discrimination was assigned the 50% reinforcement con-

tingency while the auditory discrimination was assigned the 100% contingency; on other reversals, the assignments were reversed. Within an experimental session, the 100% discrimination was presented every third trial, which meant that twice as many trials occurred for the 50% condition, while the number of reinforcers per session was equated. The issue is whether the invariance effect would still occur under these intermixed conditions.

Method

The subjects were 4 male albino rats similar to those used in Experiments 1-4. The apparatus was also the same, with two modifications. An unshielded 28-V houselight (General Instrument No. 1820 bulb) was mounted on one of the outside walls of the clear Plexiglas housing, 3 cm from the rear panel. This new light source was used with the dim panel light from Experiments 1-4 for the visual discrimination. A second sound source was also added, in addition to the previously used white noise in Experiment 4. The sound was produced by a Radio Shack oscillator unit, which produced a complex tone with a major frequency of 845 Hz of approximately 85 dB SPL intensity.

After all rats learned to barpress under a continuous reinforcement schedule in the absence of any of the four stimuli, they were given 25 reinforced presentations of each stimulus, separated by 1.5-sec ITIs. The discrimination procedure then began with trial durations that initially were 5 sec and an ITI that was 10 sec. Over the next three sessions, the trial durations and the ITIs were extended to 15 sec, at which level they remained for the duration of the experiment. The visual and auditory problems were interspersed across trials on a regular basis. Whichever discrimination was assigned the 100% probability of reinforcement was presented every third trial; the problem with the 50% probability of reinforcement occurred on the remaining two trials. Given that a particular discrimination was scheduled, whether an S+ or S- was presented was randomly determined, as was whether an S+ trial for the 50% condition was reinforced. All reinforced trials, for both percentages of reinforcement, involved an FI 15-sec schedule; that is, the first barpress after 15 sec had elapsed produced a 45-mg Noyes pellet. S- trials terminated automatically after 15 sec had elapsed. Unlike Experiments 1-4, nonreinforced S+ trials for the 50% condition did not require a response to terminate the trial, so they too terminated automatically after 15 sec had elapsed.

Different subjects were initially assigned different reinforcement probabilities for different stimuli. For S1, the auditory discrimination was assigned the 100% probability and the visual discrimination was assigned the 50% probability, with their respective S+ stimuli being the noise and the houselight. The same assignment of reinforcement probabilities were used for S2, but the S+ stimuli were the noise and the dim light. For S3 and S4, the 50% probability was initially assigned to the auditory discrimination and the 100% probability to the visual discrimination, and for both subjects the noise was the S+ for the auditory discrimination. However, the houselight was initially the other S+ for S3, whereas the dim light was the second S+ for S4. Training on a given set of contingencies continued until either a total of 10 sessions had been presented or both problems had been acquired to a criterion of 90% correct on a given session. Sessions continued for a total of 120 trials, which meant that, on average, 20 reinforcers were collected for each type of discrimination.

After the criterion had been reached for a given set of contingencies, the S+/S- status for each stimulus was reversed and training began on the new contingencies in the next session. The same assignment of probability of reinforcement to the auditory/visual discrimination remained in effect for two reversals, and then the assignment was reversed. Thus, for each block of four reversals,

each of the four stimuli served as both the S+ and the S- for both the 100% and the 50% probabilities of reinforcement. A total of eight reversals was presented.

Results

Figure 10 shows the results in terms of the percentage of responses to the S+ stimuli, averaged over the 4 subjects. The top two panels combine the results across both types of discriminations, but subdivide them according to the first four versus the second four reversals, and hence show the effects of continued training. The bottom two panels show results averaged over all eight reversals but subdivided according to the type of discrimination.

Considering first the initial four reversals (top left panel), the two acquisition functions are virtually identical throughout training, with the exception of a slightly higher accuracy for the 100% condition during the first two blocks. For Reversals 5-8, the functions are still quite similar, but with the function for the 50% reinforcement probability slightly higher throughout training, with the exception of the first session. The results in the top two panels were analyzed with a three-factor ANOVA (percentage of reinforcement $\times$ reversal number $\times$ blocks). The main effects of percentage of reinforcement and reversal number did not approach significance (both F s < 1), whereas the effect of blocks was significant [$F(9,27) = 29.43$, $p < .01$]. The interaction between reversal number and blocks was also significant [$F(9,27) = 2.73$, $p < .05$], which reflects that the performance during the first few blocks was lower during the first few reversals, but the acquisition rate, as shown by the slope of the function, and terminal performance, was considerably higher. Such effects as a function of continued reversal training have been interpreted as the effects of proactive interference (see Woodard et al., 1971). The interaction between blocks and percentage of reinforcement was also significant [$F(9,27) = 2.36$, $p < .05$]. An analysis of the simple effects for this interaction revealed that it was due to a significant difference during Block 1 of training [$F(1,3) = 12.76$, $p < .05$], with no significant differences at any other point in training.

The bottom two panels of Figure 10 show that the acquisition performance was better with the visual discrimination than with the auditory discrimination, although this effect occurred for only 3 of the 4 subjects and was not statistically significant. Figure 10 also shows that the difference during the first session, noted above, was due primarily to the auditory discrimination, although the three-way interaction among type of discrimination, blocks, and percentage of reinforcement was also not significant.

The top portion of Figure 11 shows the same results as Figure 10, but averaged over all reversals and both types of discrimination problems. These averaged data reveal that the acquisition functions were quite similar except for the slightly lower performance during the first session for the 50% condition. The bottom portion of Figure 11 shows the same data, plotted in terms of the

Figure 10. Discrimination accuracy during Experiment 5 as a function of the number of reinforcers. The top two panels separate the results according to early (Reversals 1-4) versus late (Reversals 5-8) in training. The bottom panels separate the results according to the type of discrimination.

number of training trials (for each separate problem) rather than the number of reinforcers. Thus, the results for the 100% condition correspond to blocks of two sessions, whereas those for the 50% condition correspond to one-session blocks. Unlike the results as a function of the number of reinforcers, here there is a clear effect of percentage of reinforcement, as discrimination accuracy was higher for the 100% condition throughout training. The effects seen in the bottom portion of Figure 11 were tested with a two-way ANOVA (blocks $\times$ percentage of reinforcement). The main effect of percentage of reinforcement was significant [$F(1,3) = 41.2, p < .01$], as was the effect of blocks [$F(4,12) = 25.8, p < .01$]. The interaction term was not significant ($F < 1$).

Figure 12 plots the average response rates to the S+ and S- stimuli, again averaged over all stimuli. The top portion shows the results as a function of the number of reinforcers; the bottom portion shows the results as a function of the number of trials (which equates the number of S- presentations). The pattern for S+ responding was similar for both panels: response rate to the S+ during the 100% condition was slightly higher than during the 50% condition. In addition, there was little change in responding across sessions. The greater S+ response rate for the 100% condition was consistent for 3 of the 4 subjects for all four stimuli, but no difference occurred for the 4th subject. Separate ANOVAs for S+ responding for the two portions of the graph showed no significant effect of blocks ($F < 1$ in both cases), whereas the effect

of percentage of reinforcement approached, but did not attain, significance [for the top portion, $F(1,3) = 6.71, .10 > p > .05$; for the bottom portion, $F(1,3) = 5.03$]. For neither portion did the interaction term approach significance.

The pattern for S- responding differed as a function of whether the results are plotted in terms of the number of reinforcers or the number of trials. With the number of reinforcers as the independent variable (which meant that the 50% condition had twice the number of S- presentations), lower response rates occurred in the 50% condition. An ANOVA of the S- results in the top panel revealed that this effect of percentage of reinforcement approached, but did not attain, significance [$F(1,3) = 7.26, .10 > p > .05$], whereas the effect of blocks was significant [$F(9,27) = 41.7, p < .01$]. The interaction term was not significant ($F < 1$). With the number of trials as the independent variable (which meant that number of S- presentations was equated), response rates to the S- were lower in the 100% condition, and this effect of percentage of reinforcement was significant [$F(1,3) = 11.1, p < .05$], as was the effect of blocks [$F(4,12) = 96.4, p < .01$]. However, the interaction was not significant [$F(4,12) = 1.05, p > .10$].

Discussion

Experiment 5 differed from Experiments 1-4 in that the different percentages of reinforcement were intermixed together rather than being presented as separate condi-

tions. The effect of this change was, presumably, to cause the background cues to be shared for the different discriminations, so that a comparator process would predict that the rate of acquisition should be slower for the smaller percentage of reinforcement. Contrary to this prediction, the invariance effect was at least as closely approximated as in Experiments 1-4. The only deviation was during the first session of reversal training, when the 50% condition produced slightly lower accuracy. Evidence for a difference during the early sessions of the reversal also occurred in Experiments 1 and 3. Despite this difference, the overlap of the functions seen in Figure 10, and in the top portion of Figure 11, make a strong case that acquisition is approximately invariant as a function of the number of reinforcers presented, regardless of the number of nonreinforced S+ presentations that occur in addition. Given that the same background cues were present for both reinforcement percentages, a comparison of the rate of reinforcement during the S+ to the background rate does not appear to underlie the invariance effect.

The similarity in the functions seen in Figure 10 cannot be ascribed to the insensitivity of the procedure, or to the use of a small number of subjects, because a clear

Figure 11. Discrimination accuracy during Experiment 5, averaged over all reversals and both discriminations, as a function of the number of reinforcers (top panel) and as a function of the number of trials (bottom panel).

Figure 12. Response rates to the S+ and the S- of each reinforcement condition in Experiment 5, shown separately as a function of the number of reinforcers (top) and of the number of trials (bottom).

effect of percentage of reinforcement occurred when the results were plotted in terms of the number of training trials. The results are thus similar to those obtained in Experiments 1-4.

The apparent reason why a significant effect occurred with respect to the number of trials but not the number of reinforcers is a change in the relative degree of S- responding. When the two conditions were equated with respect to the number of reinforcers, S- response rates were slightly lower with the 50% condition (although this difference was surprisingly small considering that twice the number of S- trials was presented). Responding to the S+ was also reduced with the 50% reinforcement, and the degree of reduction for both the S+ and the S- was proportionally the same, resulting in the invariance effect in terms of the measures of percentage of correct responses. When the two measures were equated with respect to the number of trials, which meant that the two reinforcement percentages were equated with respect to the number of S- presentations, S- responding was less with the 100% reinforcement probability. And since greater response rates occurred to the S+ with the 100% probability, the result was faster acquisition for the 100% condition. The greater response rate to the S+ condition with 100% reinforcement is not surprising. But there is no obvious rationale for why significantly less respond-

ing should occur to the S- under the 100% condition, when the number of S- presentations was equated for the two conditions.

GENERAL DISCUSSION

In all of the present experiments, the acquisition rate of a successive discrimination, across successive reversals of the reward values of the two stimuli, was approximately invariant as a function of the number of obtained reinforcers. Some deviation from this invariance did occur, as Experiment 4, in which equated the number of trials per session was equated, rather than the number of reinforcers, showed a slightly faster overall acquisition rate with the 50% condition. Similarly, in several experiments, there was some evidence of a slightly lower discrimination accuracy during the initial sessions of reversal training for the 50% condition. Despite these perturbations, the similarity between the acquisition functions for the 100% and 50% reinforcement conditions is impressive, especially given that clear difference consistently occurred in favor of the 100% condition when acquisition was equated with respect to the number of training trials.

The approximation to the invariance effect was not affected by adding free reinforcers to the ITIs, since whatever pattern occurred in the absence of the free reinforcers also occurred when they were added. The invariance effect also was unaffected by presenting both reinforcement percentages for different discriminations intermixed within an experimental session. Both of these manipulations should have disrupted the effect if the controlling variable were the comparison of the reinforcement rate during the S+ with the background reinforcement rate, as predicted by comparator models of conditioning (e.g., Gibbon & Balsam, 1981). That is, the ratio of S+ to background, which remained constant when S+ reinforcers were the only source of reinforcers in the situation, changed differentially for the different percentages of reinforcement when the ITI reinforcers were added. Similarly, when the same background was present during the two discriminations with different reinforcement percentages, the S+ to-background ratio was necessarily lower with the lower percentage of reinforcement. Apparently, acquisition rate is not controlled by the S+ to-background ratio.

Given that comparator models of conditioning cannot explain the invariance effects, at least for operant tasks like that used here, what are the alternatives? In the introduction, it was noted that the effect appears to challenge incremental models such as that of Rescorla and Wagner (1972), because the occurrence of nonreinforced S+ trials with the partial reinforcement conditions would seem to require additional reinforcements to offset the loss of response strength. Consequently, such models would appear to predict that smaller percentages of reinforcement would require larger numbers of reinforcements to reach a given learning criterion. Upon closer examination, however, such a prediction applies only to the growth of

excitatory strength of the S+. It says nothing necessarily about the rate of discrimination acquisition, which depends upon the *relative* response strength to the cues on S+ versus S- trials. Moreover, the response strength on such trials is due not only to the excitation/inhibition to the S+/S- themselves, but to the response strength conditioned to the common background cues as well. Given the complexity of the variables determining discrimination accuracy, it remains possible that an appropriate choice of parameter values might allow incremental models to predict the present effects. In fact, initial simulations of the Rescorla-Wagner model with respect to the present task did predict a pattern of results similar to those obtained here under some circumstances. However, it seems unlikely that any given set of assumptions would be sufficiently general to account for the invariance effect in all of the circumstances under which it has been shown to occur, including both initial acquisition of Pavlovian conditioning (Gibbon et al., 1980), of successive discrimination learning (Eckerman, 1969), and of serial reversal learning using both successive stimulus presentations (the present results) and simultaneous discriminations (perhaps the most difficult situation of incremental models to explain; see Williams, 1981). Given its generality, the invariance effect continues to challenge existing theoretical formulations and should provide a rich source of testing whether simple conditioning formulations can be extended to the more complex realm of discrimination learning.

REFERENCES

- BOWMAN, R. E. (1963). Discrimination learning-set performance under intermittent and secondary reinforcement. *Journal of Comparative & Physiological Psychology*, **56**, 429-434.
- BRANDON, S. E. (1981). Key-light-specific associations and factors determining keypecking in noncontingent schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, **7**, 348-361.
- DURLACH, P. J. (1983). Effect of signaling intertrial unconditioned stimuli in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, **9**, 374-389.
- ECKERMAN, C. O. (1969). Probability of reinforcement and the development of stimulus control. *Journal of the Experimental Analysis of Behavior*, **12**, 551-559.
- FARLEY, J. (1980). Automaintenance, contrast and contingencies: Effects of local vs. overall and prior vs. impending reinforcement context. *Learning & Motivation*, **11**, 19-48.
- GIBBON, J., & BALSAM, P. (1981). Spreading association in time. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 219-253). New York: Academic Press.
- GIBBON, J., FARRELL, L., LOCURTO, C. M., DUNCAN, H. J., & TERRACE, H. S. (1980). Partial reinforcement in autoshaping with pigeons. *Animal Learning & Behavior*, **8**, 45-59.
- GODDARD, M. J., & JENKINS, H. M. (1987). Effect of signaling extra unconditioned stimuli on autoshaping. *Animal Learning & Behavior*, **15**, 40-46.
- GONZALEZ, R. C., BEHREND, E. R., & BITTERMAN, M. E. (1967). Reversal learning and forgetting in bird and fish. *Science*, **158**, 519-521.
- GRAU, J. W., & RESCORLA, R. A. (1984). Role of context in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 324-332.
- MACKINTOSH, N. J. (1974). *The psychology of animal learning*. New York: Academic Press.

- NEVIN, J. A. (1974). Response strength in multiple schedules. *Journal of the Experimental Analysis of Behavior*, **21**, 389-408.
- NEVIN, J. A. (1979). Reinforcement schedules and response strength. In M. D. Zeiler & P. Harzem (Eds.), *Advances in the analysis of behaviour: (Vol. 1). Reinforcement and the organization of behaviour* (pp. 117-158). New York: Wiley.
- NEVIN, J. A., MANDELL, C., & ATAK, J. R. (1983). The analysis of behavioral momentum. *Journal of the Experimental Analysis of Behavior*, **39**, 49-59.
- NEVIN, J. A., MANDELL, C., & YARENSKY, P. (1981). Response rate and resistance to change in chained schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, **7**, 278-294.
- PAPINI, M. R., & OVERMIER, J. B. (1985). Partial reinforcement and autoshaping of the pigeon's keypeck behavior. *Learning & Motivation*, **16**, 109-123.
- RACHLIN, H., & BAUM, W. M. (1972). Effects of alternative reinforcement: Does the source matter? *Journal of the Experimental Analysis of Behavior*, **18**, 231-241.
- REILLY, S., & SCHACHTMAN, T. R. (1987). The effects of ITI fillers in autoshaping. *Learning & Motivation*, **18**, 202-219.
- RESCORLA, R. A., & WAGNER, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- SCHACHTMAN, T. R., & REILLY, S. (1987). The role of local context in autoshaping. *Learning & Motivation*, **18**, 343-355.
- STADDON, J. E. R., & SIMMELHAG, V. L. (1971). The "superstition" experiment: A reexamination of its implications for the principles of adaptive behavior. *Psychological Review*, **78**, 3-43.
- TIMBERLAKE, W., & LUCAS, G. A. (1985). The basis of superstitious behavior: Chance contingency, stimulus substitution, or appetitive behavior? *Journal of the Experimental Analysis of Behavior*, **44**, 279-299.
- WILLIAMS, B. A. (1976a). Elicited responding to signals for reinforcement: The effects of overall versus local changes in reinforcement probability. *Journal of the Experimental Analysis of Behavior*, **26**, 213-222.
- WILLIAMS, B. A. (1976b). Short-term retention of response outcome as a determinant of serial reversal learning. *Learning & Motivation*, **7**, 418-430.
- WILLIAMS, B. A. (1977). Contrast effects in simultaneous discrimination learning. *Animal Learning & Behavior*, **5**, 47-50.
- 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. (1983). Revising the principle of reinforcement. *Behaviorism*, **11**, 63-88.
- WOODARD, W. T., SCHOEL, W. M., & BITTERMAN, M. E. (1971). Reversal learning with singly presented stimuli in pigeons and goldfish. *Journal of Comparative & Physiological Psychology*, **76**, 460-467.

(Manuscript received August 16, 1988;
revision accepted for publication February 9, 1989.)