

Acquisition with partial and continuous reinforcement in pigeon autoshaping

DANIEL A. GOTTLIEB

University of Pennsylvania, Philadelphia, Pennsylvania

Contemporary time accumulation models make the unique prediction that acquisition of a conditioned response will be equally rapid with partial and continuous reinforcement, if the time between conditioned stimuli is held constant. To investigate this, acquisition of conditioned responding was examined in pigeon autoshaping under conditions of 100% and 25% reinforcement, holding intertrial interval constant. Contrary to what was predicted, evidence for slowed acquisition in partially reinforced animals was observed with several response measures. However, asymptotic performance was superior with 25% reinforcement. A switching of reinforcement contingencies after initial acquisition did not immediately affect responding. After further sessions, partial reinforcement augmented responding, whereas continuous reinforcement did not, irrespective of an animal's reinforcement history. Subsequent training with a novel stimulus maintained the response patterns. These acquisition results generally support associative, rather than time accumulation, accounts of conditioning.

Probability of reinforcement and length of intertrial interval (ITI) have long been acknowledged as primary determinants of Pavlovian conditioned responding. It is generally agreed that both higher probabilities of reinforcement and longer ITIs promote performance. Most often, these factors are envisioned as separable and independent influences. However, Gibbon and Balsam (1981) and Gallistel and Gibbon (2000) have proposed a novel theoretical approach in which these influences are integrated in a common framework. They have proposed what have been termed *time accumulation models* (Bouton & Sunsay, 2003), in contrast to traditional associative models.

Theoretical Background

Emerging from a tradition in which timing has been studied, in time accumulation models learning has been conceived of as the joint encoding of stimulus duration and number of presented reinforcers, information that is combined to yield an estimated rate of reinforcement. In this framework, learning the relation between a conditioned stimulus (CS) and an unconditioned stimulus (US) involves learning the rate of reinforcement during both the CS and the background in which the CS is sometimes presented. Responding depends on a computation of the ratio of the rate of reinforcement during the background to that during the CS. Because the computation of rate is determined both by number of reinforcers and stimulus

duration, probability of reinforcement (dependent on number of reinforcers) and ITI (duration, by definition) covary to affect performance in a regular and highly specified way.

Consider, from the perspective of Gallistel and Gibbon (2000), a situation in which a CS is reinforced on every presentation. The rate of reinforcement during the CS is then determined by the duration of the CS. Similarly, the rate of reinforcement during the background is determined by its duration. In the typical situation, in which the US is never presented in the absence of the CS, the rate of reinforcement during the background alone is zero. Computing a ratio of CS reinforcement rate to background reinforcement rate then leads to a mathematically inconvenient denominator of 0. Gallistel and Gibbon have suggested that even if an animal has not yet been reinforced during the background, it holds some nonzero expectation that reinforcement may occur in the future. They proposed that the estimated rate of background reinforcement is not 0 but rather a constant divided by the total time experienced during the background alone. They set the value of this constant equal to 1, but its particular value is not important for most predictions, including those relevant to the present experiment.

Consider the specific continuous reinforcement situation in which the CS is presented for 30 sec and the ITI is 90 sec. According to Gallistel and Gibbon (2000) the estimated rates of reinforcement after 1 trial are 1/30 and 1/90 for the CS and the background alone, respectively. Thus, the computed ratio of CS reinforcement rate and background reinforcement rate is $(1/30)/(1/90)$, or 3. After multiple trials, the rate of reinforcement during the CS remains constant, but the estimated rate during the background decreases. For example, after 10 trials there will have been 900 sec of background time without rein-

This research was supported by National Institutes of Health Grant R01 MH 67858. I gratefully thank Robert Rescorla for help at every stage, Matt Lattal for advice and discussion, and Andre Roussin for assistance in data collection. Address correspondence to D. A. Gottlieb, Rutgers, Center for Cognitive Science, Psychology Building Addition, Room A129, 152 Frelinghuysen Rd., Piscataway, NJ 08854 (e-mail: dag@ruccs.rutgers.edu).

forcement, yielding a rate estimate of $1/900$ and a ratio of $(1/30)/(1/900)$, or 30. This particular value is compared with a fixed threshold value to determine whether responding to the CS is warranted. Therefore, responding increases in likelihood across conditioning trials.

As has been noted by Gibbon and Balsam (1981) and Gallistel and Gibbon (2000), an interesting prediction emerges when reinforcers are deleted from a continuous reinforcement schedule, resulting in a partial reinforcement schedule. By simply removing reinforcers, the time between CS presentations is held constant while the time between reinforcers is lengthened. This has the obvious effect of decreasing the probability of reinforcement, which has often been found to be detrimental to conditioned responding (see Mackintosh, 1974, pp. 72–73, for a discussion and notable exceptions). However, this simultaneously increases the time between reinforcers, which has been found to be beneficial to conditioned responding (see Barela, 1999, for a recent discussion; Calvin, 1939, reported in Hilgard & Marquis, 1940, pp. 148–149; Humphreys, 1940; Schlosberg, 1934).

According to time accumulation models, the tradeoff between the detrimental effect of nonreinforcement and the beneficial effect of lengthening the time between reinforcers is exact, and deleting reinforcers should not affect acquisition. To see this, consider the predictions of Gallistel and Gibbon (2000) for a variation of the previously discussed continuous reinforcement schedule in which a 50% schedule of reinforcement is generated by deleting the odd-numbered reinforcers from the continuous reinforcement schedule. After one reinforced trial, the CS has been presented twice for a total of 60 sec. Therefore, the rate of reinforcement during the stimulus is once every 60 sec. After one reinforced trial and two trials in total, the animal has been exposed to the background alone for a total of 180 sec without reinforcement, yielding a rate of $1/180$ sec. The ratio of rates is $(1/60)/(1/180)$, or 3, the same as it was after 1 reinforcer in the continuous reinforcement schedule. After 10 reinforcers, the rate of reinforcement during the partially reinforced stimulus remains at once every 60 sec. However, the background alone has been presented for 1800 sec without reinforcement, yielding a rate of $1/1,800$ sec. Thus, the ratio of rates is $(1/60)/(1/1,800)$, or 30, again the same as it was after 10 reinforcers in the continuous reinforcement schedule.

A different way of viewing the background reinforcement rate leads to similar predictions in another time accumulation model. Gibbon and Balsam (1981) acknowledged that the background is present *during* CS presentation, as well as in its absence. They proposed that background reinforcement rate be computed from the entire duration of background experience, rather than from the limited time in which CSs are absent. In this manner, as long as at least one reinforcer is delivered, the background rate of reinforcement will be greater than zero.

According to the model of Gibbon and Balsam (1981), a continuous reinforcement situation leads to different computations than does the model of Gallistel and Gibbon

(2000). After one trial with a 30-sec CS and a 90-sec ITI, the estimated rates of reinforcement are $1/30$ during the CS and $1/120$ during the background. Here, CS durations and ITI durations sum to determine the total background time; in this case, the summation of the 30-sec CS and the 90-sec ITI yields a total background time of 120 sec. The ratio of CS reinforcement rate and background reinforcement rate is now $(1/30)/(1/120)$, or 4. Also unlike Gallistel and Gibbon's model, Gibbon and Balsam's model computes the same rate of background reinforcement independently of the number of trials. This is because the background duration and the number of reinforcers presented during the background grow proportionately over trials. Because the CS reinforcement rate also remains constant over time, the computed ratio does not change over time. Although Gibbon and Balsam have suggested a mechanism to account for gradual learning over trials, this mechanism does not alter the essential predictions of the model important for the current situation.

Like Gallistel and Gibbon's (2000) model, Gibbon and Balsam (1981) predict the same ratio computation in the case of the partial (50%) reinforcement situation discussed above as that in the continuous reinforcement situation. After one reinforced trial in the partial reinforcement schedule, the CS reinforcement rate is once every 60 sec, having been presented twice and reinforced once. The total background time over the course of these two trials is 240 sec. The background rate of reinforcement is, therefore, once every 240 sec. The ratio of CS reinforcement rate and background reinforcement rate is $(1/60)/(1/240)$, or 4, the same as it is with continuous reinforcement. Again, this ratio is the same independently of the number of trials, due to the proportional growth of reinforcer number and background duration.

For both views of background reinforcement rate, removing reinforcers from a continuous reinforcement schedule proportionally increases the estimated length of time until reinforcement during both the CS and the background, leading time accumulation models to predict that learning will not be affected in this situation. This prediction emerges both if background rate of reinforcement is computed using total background time and if it is computed using time when the background is presented alone. It is important to note that this prediction is based on learning's being assessed as a function of the number of reinforcers, not the number of trials.

It is also important to note that in the above example, deleting odd-numbered reinforcers from the continuous reinforcement schedule generates a partial reinforcement schedule in which reinforced and nonreinforced trials occur in strict alternation. This particular method of generating a partial reinforcement schedule was chosen for ease of exposition. However, the predictions of time accumulation models also hold for the more typical case in which a partial reinforcement schedule involves random or pseudorandom presentations of reinforced and nonreinforced trials.

Unlike time accumulation models, associative models of learning (e.g., Mackintosh, 1975; Pearce, 1987; Res-

corla & Wagner, 1972) are based not on an organism's encoding of time, but on the development of associative strengths between a stimulus and a reinforcer. These models envision learning as the incrementing of associative strengths on reinforced trials and the decrementing of associative strengths on nonreinforced trials. They are thus well suited for analyzing differences in probability of reinforcement; however, many of these models do not explicitly take the duration of the CS or the ITI into account and are thus poorly suited for characterizing temporal relations. Even those associative models that do incorporate the duration of events (e.g., Sutton & Barto, 1990; Wagner & Brandon, 1989) do so as a separate effect, unrelated to the probability of reinforcement. Consequently, associative models do not predict the quantitative tradeoff between the effects of probability of reinforcement and ITI that time accumulation models predict. If learning were shown to be invariant under conditions in which reinforcers are deleted from some trials, that would be strong evidence in favor of time accumulation models.

Empirical Background

In a recent theoretical paper, Gallistel and Gibbon (2000) claimed that in circumstances in which the time between trials is held constant and acquisition is plotted against reinforcements, "the nonreinforcements that occur during partial reinforcement do not affect the *rate of acquisition*" (p. 298). That claim was based largely on the work of Gibbon, Farrell, Locurto, Duncan, and Terrace (1980), who reported a systematic study of reinforcement probabilities and ITIs, using a pigeon autoshaping procedure. In that study, 25 separate groups of birds were presented with daily sessions containing a constant number of reinforcers. The groups differed in time between trials and the number of nonreinforced trials, representing every combination of five ITIs and five probabilities of reinforcement ranging from 10% to 100%. When learning was assessed as the number of reinforcers to reach a criterion of 3/4 trials with a response, groups sharing the same ITI but differing in probability of reinforcement showed similar performance. There was a small numerical superiority in groups with higher probabilities of reinforcement, but it was significant only by some statistical measures. Consequently Gallistel and Gibbon argued that the removal of reinforcers does not slow the rate of acquisition. That is, the predicted tradeoff between probability of reinforcement and ITI was observed.

Although these results seem to provide strong support for time accumulation models, there are reasons to view them with a measure of caution. First, because the aim of the experiment was to explore a range of parameter values, the number of subjects in each condition was small, usually 4 per group. This may have had the effect of reducing the statistical power of any particular comparison, even though a total of 113 subjects were used.

Second, 15 of the 113 birds were replaced for failure to reach criterion in 10 sessions. Unfortunately, such failures to reach criterion were not distributed equally

among the groups. Eight of the 15 replacements occurred in the groups with the lowest (10%) probability of reinforcement. This seems likely to result in a bias toward overestimating the rate of acquisition in groups with lower probabilities of reinforcement.

Third, the dependent variable in Gibbon et al.'s (1980) experiment was the point at which the animal made at least one response on 3 out of 4 consecutive trials. However, that performance was assessed as a function of the number of reinforcers, and not trials, to reach criterion. Consequently, subjects reinforced less frequently had more trials on which they had the opportunity to reach criterion for any given number of reinforcers. Indeed, a subject reinforced every 10th trial would have 10 times as many presentations of the CS as a subject reinforced every trial. As a result, it would have approximately 10 times as many opportunities to reach this acquisition criterion. Thus, a 3/4 criterion measured across all trials increases the probability of the lower frequency groups meeting criterion for any given level of learning.

Fourth, learning was assessed according to differences in performance during the course of acquisition. As has been noted by several authors (e.g., Lattal, 1999; Rescorla, 1988), this complicates interpretation by confounding the conditions of learning with the conditions in which learning is assessed. Probability of reinforcement and ITI may exert influence over learning, or they may exert influence over the expression of learning. Because observations were made while groups differed both in their history of reinforcement and in their current schedule of reinforcement, it is unclear whether differences in conditioned responding were attributable to learning or performance effects.

Finally, the learning criterion used as the dependent variable was meant to capture the point of conditioned response (CR) emergence. In doing so, the criterion ignores the changes in responding that take place postcritically. Because time accumulation models do not endow the point of CR emergence with special properties, there is no reason to restrict an analysis of the models' behavior to a single point in training. As long as time accumulation models predict behavior equally at all points in acquisition, it seems reasonable to assess these models at all points in acquisition as well. Restricting analysis to a single point in acquisition may add an inadvertent bias toward a null result since any single arbitrarily defined point in training may be particularly insensitive to between-group differences brought about by variations in reinforcement schedule.

The results of other relevant studies in the autoshaping literature have proven inconclusive. Using pigeons, Gonzalez (1973) reported superior performance in a group reinforced on every trial, as compared with a group reinforced every fourth trial; however, on average, partially reinforced pigeons first showed a CR at an earlier reinforcement number than did continually reinforced pigeons. Because there were only 3 subjects per group, statistical information was not provided. Wasserman,

Hunter, Gutowski, and Bader (1975) and Wasserman, Deich, Hunter, and Nagamatsu (1977) compared partial (67%, 50%, and 33%) and continuous reinforcement in chicks, plotting data against trials, and not against reinforcements. If one replots their data by reinforcer, there is some evidence of superiority in the partially reinforced conditions; however, such replotting necessarily consists of a single data point with no statistical evaluation. Crawford, Steirn, and Pavlik (1985), using Japanese quail, compared 100% and 50% reinforcement, plotting their results by trial, and not by reinforcer. If one replots their data, there is no obvious difference in initial acquisition and an asymptotic difference in favor of a partially reinforced group in both rate and percent CR measures. Again, however, there are no statistics to support these observations.

Aside from Gibbon et al.'s (1980) report, the most systematic autoshaping data come from work by Papini and Overmier (1984, 1985), using pigeons. In early acquisition, they found no difference in performance between 25% and 100% groups when the groups were equated for number of trials per session. In this situation, the number of reinforcers per session necessarily differs between the partially and the continuously reinforced groups. On the other hand, they did see a difference in favor of the 100% animals when numbers of reinforcers per session were equated. These findings are at odds with those of Gibbon et al. (1980), who also equated groups by the number of reinforcers, and not by the number of trials.

For present purposes, it is important to note that Papini and Overmier (1984, 1985) did not present rate data, so it is not known how this measure would correspond with the percent CR measure that they used. Also, no common test was administered to separate out the influences of learning and performance. Furthermore, there were only 4 subjects per group, although a significant difference was still detected in one of the conditions with this number of subjects. It would be hasty to make any firm conclusions on the basis of the available data, but Papini and Overmier's (1984, 1985) data suggest that the regularity predicted by time accumulation models may not consistently hold in autoshaping preparations.

The Present Experiment

The main intention of the present experiment was to document the full course of acquisition in the autoshaping procedure under conditions of partial and continuous reinforcement. Autoshaping was used for two reasons. First, it allows more direct comparisons to be made with prior work in which this preparation was used. Second, a critical examination of the autoshaping preparation itself might give insight as to its potential usefulness in evaluating quantitative claims pertaining to partial reinforcement. To this end, several dependent measures were examined, and learning was assessed during acquisition, a common test, a reversal of contingency, and novel stimulus training. The reversal of contingency was useful for examining the influence of a

change in reinforcement schedule on an already developed response. Novel stimulus training was useful in determining whether the pattern of responding developed over the course of acquisition and reversal training was stimulus or animal specific. It was further hoped that additional information could be provided about the trade-off between probability of reinforcement and ITI, as well as about the correspondence between rate of responding and the probabilistic percent CR and criterion measures.

In order to optimize the chance of observing differences, this experiment deviated from previous ones in several ways. First, as compared with previous autoshaping work on this topic, a relatively large number of subjects was concentrated in only two groups that substantially differed in their probability of reinforcement. Groups receiving 100% and 25% reinforcement contained 15 and 16 subjects, respectively.

Second, data were taken from the same number of corresponding trials in each group. Groups receiving 100% and 25% reinforcement differ not only in the average recency of reinforcement, but also in the total number of possible trials to reach criterion. These confounds were eliminated by plotting the data only from trials that immediately followed a reinforced trial. This served two functions. First, by equating trial number, the type of bias due to differential opportunities to reach a 3/4 learning criterion could be avoided. Second, plotting data only from trials following reinforced trials controlled for any local performance effects due to differential recency of reinforcement.

Third, to best characterize the data, three different measures were taken on those trials: response rate, probability of at least one response, and number of trials to reach a 3/4 criterion. Response rate is a commonly reported measure in autoshaping studies. Probability of responding is an important measure because it allows for direct tests of the time accumulation models. The response rule of these models is one in which the ratio of CS reinforcement rate and background reinforcement rate is compared with some threshold value. Responding should occur if the computed ratio is higher than the threshold. As such, the models explicitly predict *whether* an organism should respond. The natural measure to use in this situation is probability of responding.

The last measure, number of trials to a 3/4 criterion, is one example of a probabilistic measure and is useful in comparing results with those of Gibbon et al. (1980). The 3/4 criterion is meant to capture the first point at which probability of responding is appreciably greater than 0. It is important to note that although this point may represent the emergence of conditioned responding, time accumulation models do *not* treat this point as otherwise special. That is, as long as the computed ratio is the same across groups, there should be equivalent probability of responding. If this were not the case, there would be no reason for equivalent ratio values to promote equivalent responding at *any* point. In other words, time accumulation models in their present forms must neces-

sarily predict probabilistic responding at all points in acquisition, and not solely at the point of CR emergence.

Finally, this experiment included a test to determine whether group differences were attributable to learning or to local performance effects. This assessment involved observing performance under conditions of both 25% and 100% reinforcement in subjects that had received either 25% or 100% reinforcement. Without an assessment of this type, it is not possible to separate out the influences of learning and performance. It is conceivable, for instance, that current exposure to partial reinforcement evokes motivational states that may influence performance without affecting learning. Indeed, the frustration brought about by nonreinforcement in a partial reinforcement schedule has been theorized to underlie the finding that running speeds in instrumental runway studies are sometimes faster for partially, rather than continuously, reinforced animals (Amsel, 1967; Mackintosh, 1974, p. 160; Spence, 1960, pp. 98–100). A common test such as that used in the present study can help to control for the general effects of motivational factors, such as frustration; however, it is important to note that the common test used in this study does not control for motivational effects that are under contextual or stimulus control, such as those effects thought to underlie the partial reinforcement extinction effect (e.g., Amsel, 1958).

The first phase of this experiment involved training pigeons with either 100% or 25% reinforcement. Each animal received 12 reinforcers/day, with the 25% subjects receiving additional nonreinforced trials. After 12 days of acquisition, half the animals receiving 25% reinforcement were switched to 100% reinforcement, and half the animals receiving 100% reinforcement were switched to 25% reinforcement. After 24 sessions in this new condition, all the animals were presented a novel stimulus that was reinforced 75% of the time. This final phase of the experiment was intended to examine the stimulus specificity of response differences that emerged in training. Throughout, the primary data were the probability and rate of responding on trials following reinforcement.

METHOD

Subjects

The subjects were 31 experimentally naive female White Carneau pigeons purchased from Palmetto Pigeon Plant (Sumter, SC). They were housed in pairs, with water continuously available. The subjects were kept at 80% of their free-feeding weight and were fed immediately after their experimental sessions.

Apparatus

The subjects received training in eight identical operant chambers measuring $27 \times 27 \times 35$ cm. Each chamber was placed in a wooden sound- and light-attenuating box. A fan attached to a side wall provided ventilation and raised the ambient noise level to approximately 62 dB re 20 μ N/m. During training, a 6-W bulb attached to the back wall was illuminated, except during the 5-sec presentation of the grain. The front wall of the operant chamber was aluminum, whereas the side walls and ceiling were made of clear Plexiglas. The floor consisted of 1-cm wire mesh screening.

The food magazine was located on the front wall, 5 cm from the floor, with an opening of 5×5 cm. When the food hopper was raised, Purina Pigeon grain appeared in this opening. A 6-W light attached to the top of the food magazine was illuminated when grain was available.

The bottom of the response key was located 9 cm above the food magazine. Made of clear Lucite acrylic, it measured 11.8×14.5 cm and was exposed by a 10×8 cm opening in the front wall. Any peck of sufficient force to displace the response key forward activated a relay that recorded that response. Approximately 5 mm behind the response key was a small Magnavox CK3923 monitor on which stimuli were presented. Opaque black tape blocked the top half of the monitor from view in four chambers and the bottom half of the monitor in the other four chambers. The stimuli appeared in the center of the visible portion of the monitor. The stimuli consisted of a white dot measuring 0.8 cm in diameter and a blue star measuring 1.8 cm in diameter.

Experimental events were controlled and recorded by a computer located in an adjacent room.

Procedure

Magazine training. The subjects were initially trained to retrieve food from the magazine when the hopper was activated. The doors to the sound-attenuating boxes were kept open, and the birds were monitored with a video camera. At variable intervals, the hopper was activated and kept active for variable lengths of time. Both these parameters were under the control of the experimenter. Session lengths changed according to the progress of the birds but were generally 45–60 min. The subjects received one session per day for 5 days, at the end of which they all approached the magazine and ate grain when the hopper was active.

Initial acquisition. The subjects were divided into two groups, both of which received excitatory conditioning for the next 12 days. Group 100% contained 15 subjects that received 12 5-sec presentations of the small white dot, each followed immediately by a 5-sec access to grain. Group 25% contained 16 subjects that also received 12 reinforced presentations of the white dot. In addition, Group 25% received 36 nonreinforced presentations of the white dot interspersed in pseudorandom order with reinforced presentations. The 25% schedule was constrained so that no more than two reinforced and six nonreinforced trials could be presented in succession. The ITI for this and all the following stages averaged 60 sec, with a range of 30–90 sec. Total session length for Group 100% was 12 min, and session length for Group 25% was 48 min.

Reversal training. Following initial acquisition, Group 100% and Group 25% were each divided into two groups matched on their final performance. Group 100→100 ($n = 8$) continued to receive 12 reinforced presentations of the white dot. Group 100→25 ($n = 7$) received 12 reinforced presentations of the dot, as well as 36 intermixed nonreinforced presentations. Group 25% was likewise divided into two groups. Group 25→25 ($n = 8$) continued to receive 12 reinforced and 36 nonreinforced presentations of the white dot. Group 25→100 ($n = 8$) received only 12 reinforced presentations of the dot. This phase of conditioning lasted 24 days.

Novel stimulus training. After Phase 2 conditioning, all the animals received 5 days of training with a novel stimulus, the solid blue star. All the subjects received 12 reinforced presentations of the star intermixed with 4 nonreinforced presentations. Thus, all the subjects were reinforced 75% of the time. The ITI continued to be variable around a 60-sec mean, ranging from 30 to 90 sec.

All the data were analyzed using a two-tailed significance level of .05. Nonparametric statistics were used throughout. With nonparametric statistics, as compared with parametric statistics, fewer assumptions need be made about the manner in which an underlying learning process maps into behavior. Consequently, the inferences made from behavior to learning are applicable to a broader range of models—generally, any that assume at least a monotonic mapping from learning into behavior.

RESULTS

Initial Acquisition

Figure 1 shows the course of acquisition for Group 100% and Group 25%. The top panel displays percent trials with a peck (percent CR), and the bottom panel shows response rates plotted for each session (12 reinforcers), with data points calculated from trials following reinforcement. Both panels depict initially greater responding in Group 100% than in Group 25%. This effect reversed on Session 5, when Group 25% began to show greater responding; Group 25% responded more often than Group 100% for the remainder of acquisition.

Statistical analysis confirmed that Group 100% initially acquired responding more quickly than did Group 25%. Over the first two sessions (24 reinforcers), this difference was reliable both for percent CR [Mann–Whitney $U(16,15) = 64$] and for response rate [$U(16,15) = 65$]. Beginning with Session 5, Group 25% showed a higher rate and higher probability of responding than did Group 100%. This difference was first significant for rate of responding during Sessions 7 and 8 [$U(16,15) = 66$] and was also reliable during the final two sessions of acquisition [$U(16,15) = 49$]. This higher responding in Group 25% was also significant, using percent CR when the numerical difference was greatest, during Sessions 8 and 9




Figure 1. Responding during initial acquisition training in animals receiving 100% or 25% reinforcement. Data points represent the 12 trials following reinforcement over one session. The top panel shows the percentage of trials with at least one keypeck. The bottom panel shows mean responses per minute. Error bars represent SEMs.

[$U(16,15) = 58$], but not during the final two sessions [$U(16,15) = 75, p = .06$].

The data were also analyzed in a manner similar to that in Gibbon et al. (1980), using a criterion of responding on three out of four consecutive trials. Following the procedure in Gibbon et al., subjects that did not reach criterion by the end of acquisition were omitted from this analysis. This led to the removal of 3 animals from Group 100% and 0 animals from Group 25%. (These animals were removed from the criterial analysis in order to make close comparisons possible with the analysis of Gibbon et al.; however, these animals were not excluded from any other analyses presented in this article.) Median reinforcers to criterion were 21.5 and 37 for Group 100% and Group 25%, respectively, a difference that was not reliable [$U(16,12) = 60, p = .1$].

In order to reduce the biases in opportunity to reach criterion and to control for recency of reinforcement, a more detailed three out of four analysis was done including only trials immediately following reinforcement. Here, median reinforcers to criterion were 21.5 and 39.5 for Group 100% and Group 25%, respectively, a reliable difference [$U(16,12) = 49$]. The difference between this restricted analysis and the one above in which all trials were used was due to 9 out of 16 animals needing at least two more reinforcers to reach criterion in the restricted case. On average, the subjects in Group 100% reached criterion during the second session whereas, on average, the subjects in Group 25% reached criterion during the fourth session.

To gain insight as to the relative point in acquisition at which the learning criterion was reached, responding immediately following the reaching of criterion was compared with responding at the end of initial acquisition. The subjects in Group 100% responded, on average, 66.5 times per minute during the first full session after the session in which they reached criterion. These same subjects increased their responding to an average of 112.9 times per minute in their final session of initial acquisition. (The 3 subjects in Group 100% that did not reach criterion could not be included in this analysis. In addition, a 4th subject in Group 100% could not be included, because it reached criterion on the last acquisition session.) The subjects in Group 25% responded, on average, 58 times per minute in the first full session after the session in which they reached criterion. These same subjects increased their responding to an average of 221.9 times per minute in their final session of initial acquisition. These increases in responding after reaching criterion were reliable for both Group 100% [Wilcoxon $T(11) = 10$] and Group 25% [$T(16) = 1$].

Overall, the pattern of data suggests a close correspondence between percent CR, response rate, and trials to criterion at the beginning of acquisition. All the measures converge to the conclusion that acquisition originally proceeds more rapidly in a group reinforced on every trial than in a group reinforced, on average, every fourth trial.

With more extended training, it is also clear that performance in Group 25% grows to be superior to that in Group 100%. This is most apparent in the response rates, which continued to rise even when the animals were responding on nearly every trial.

Reversal Training

Figure 2 shows responding over blocks of two sessions (24 reinforcements), including only trials following reinforcement, during both initial acquisition and reversal training. Initial acquisition has been replotted so as to display the full course of acquisition. The top panel shows percent CR, and the bottom panel shows response rates. The first data point after the switching of contingencies can be considered a common test. During these sessions, the animals that differed in their initial acquisition conditions were tested under either common 25% reinforcement or common 100% reinforcement. During this test, Group 100→100 and Group 100→25 behaved similarly, as did Group 25→100 and Group 25→25. Thus, in the first two postshift sessions, the animals showed no signs that responding was governed by the postreversal conditions. Rather, responding immediately after the shift was similar to responding before the shift. This suggests that differences observed during acquisition were not attributable to the local effects determined by schedule of reinforcement, effects that include the temporal proximity of nonreinforcement in partially reinforced animals, as well as the differential time in the session from which observations were taken in partially and continuously reinforced animals.

Over the course of reversal training, continuously reinforced groups did not appear to increase either their response probability or their response rate. There was no reliable difference between the first 2 and last 2 days of reversal training for the combined Groups 100→100 and 25→100 on either probability [Wilcoxon $T(9) = 13$] or rate of responding [$T(16) = 37$]. Groups given 25% conditioning during reversal training increased rate of responding, a difference that was reliable for both 25% groups combined [$T(15) = 0$] and individually for Group 25→25 [$T(8) = 0$] and Group 100→25 [$T(7) = 0$]. The combined 25% groups also showed an increase in probability of responding [$T(10) = 4$] from the first 2 to the last 2 days of reversal training. However, neither group alone showed a reliable increase in probability of responding. The animals in Group 25→25 were responding close to maximally at the beginning of this phase. The probability of responding appeared to increase in Group 100→25, but this was not significant, because 2 of the 7 animals in the group responded maximally at both the beginning and the end of Phase 2 conditioning, effectively reducing the sample size to 5. The remaining 5 animals all responded more frequently at the end of this stage than at the beginning. These observations suggest that in reversal training, reinforcing on every fourth trial increased responding, whereas reinforcing on every trial did not.




Figure 2. Responding during initial acquisition training and reversal training. Data points represent the 24 trials following reinforcement over two sessions. After initial acquisition training, half the animals in Group 100% and in Group 25% had their contingencies switched. The top panel shows the percentage of trials with at least one keypeck. The bottom panel shows mean responses per minute. Error bars represent *SEMs*.

Another observation is that groups that received the same schedule of reinforcement during reversal but differed in their schedule of reinforcement during initial acquisition showed different response patterns at the end of reversal training. This is most obvious when one looks at rates of responding. Group 25→100 responded at more than twice the rate of Group 100→100 over the last two sessions of reversal training, but this difference did not quite reach reliability [$U(8,8) = 14$, $p = .06$]. Group 25→25 responded at almost twice the rate of Group 100→25 over the last two sessions, a reliable difference [$U(8,7) = 10$].

To better characterize responding over the 5-sec CS, second-by-second data were taken from sessions at the end of reversal training. Figure 3 shows the proportion of total responding during each second of the 5-sec CS, averaged across the last four sessions of reversal training. It appears that the animals receiving 100% reinforcement in reversal training showed proportionately greater responding early in the stimulus than did animals receiving 25% reinforcement. In support of this, groups differed significantly during the first 2 sec of the stimulus (Kruskal-Wallis $H = 11.29$) and during the final 2 sec of the stimulus ($H = 11.60$). Further analysis showed that the animals receiving

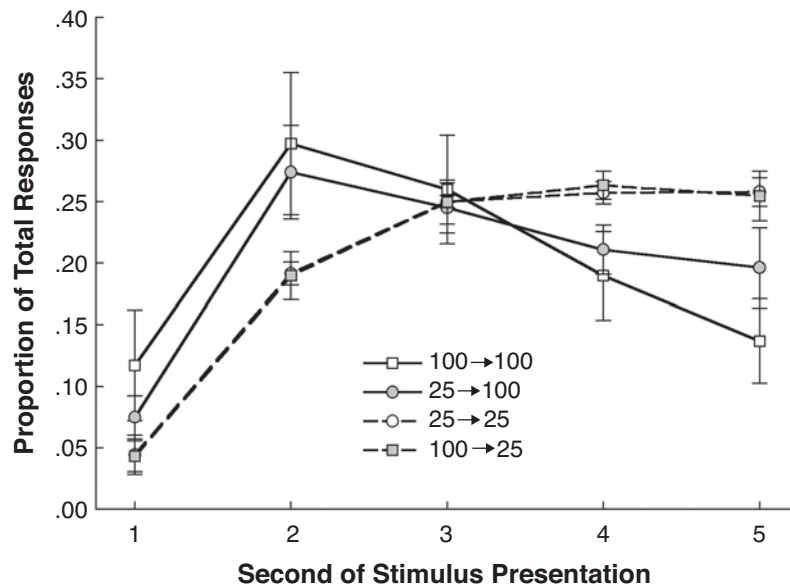


Figure 3. Proportion of total responding as calculated for each second of the 5-sec conditioned stimulus. Data are from the last 4 days of reversal training. Error bars represent *SEMs*.

100% reinforcement during reversal training responded proportionately more during the first 2 sec of the CS than did the animals receiving 25% reinforcement [$U(16,15) = 36$]. The animals receiving 100% reinforcement in reversal training also responded proportionately less during the final 2 sec of the CS than did the animals receiving 25% reinforcement [$U(16,15) = 39$]. Thus, the current schedule of reinforcement appeared to at least partially determine the pattern of responding during the CS. A continuous reinforcement schedule led to proportionately greater responding during the initial part of the CS and to proportionately less responding during the final part of the CS than did a partial reinforcement schedule.

Novel Stimulus Training

Figure 4 shows responding during the five sessions of novel stimulus training. The first data point shows responding over the first 4 trials. The other data points show responding blocked by 16-trial sessions. Responding to the novel stimulus was originally low across all groups. Over the first 4 trials, responding did not differ significantly among groups ($H = 3.47$). Differences in responding emerged rapidly with continued training and were reliable over the final two sessions ($H = 10.30$). Furthermore, the pattern of responding was similar to terminal reversal training responding. Group 100→100 responded the least, Group 25→25 responded the most, and Group 25→100 and Group 100→25 were in the middle. In support of this observation, the difference between Group 100→100 and the two middle groups, Group 25→100 and Group 100→25, was significant [$U(8,15) = 26$], as was the difference between Group 25→25 and the two middle groups [$U(8,15) = 28$].

The low responding during the first few trials suggests that the subjects were able to discriminate the novel stimulus from the original stimulus. Response rates increased rapidly and, by the end of training performance to the novel stimulus, took on the characteristics of responding to the original stimulus.

DISCUSSION

The present data suggest that even in situations in which the time between CS presentations is held constant, partial reinforcement slows acquisition of a CR relative to continuous reinforcement, when learning is assessed across reinforcements and not across trials. The data further suggest that in the pigeon autoshaping procedure, partial reinforcement can lead to higher rates and probabilities of responding than does continuous reinforcement. This differential responding was shown to be a product of a subject's history of reinforcement, and not of the current schedule of reinforcement. Differential responding was resistant to a change in reinforcement schedule, although partial reinforcement eventually augmented responding, whereas continuous reinforcement did not after an asymptotic level had been reached. Acquisition with a novel stimulus suggested that the differences in responding were not tied to a particular stimulus.

Initial Acquisition

The data from initial acquisition replicate the findings of Papini and Overmier (1984, 1985) that under conditions in which both the number of reinforcers and the time between trials is constant, birds learn faster to a continuously reinforced stimulus than to a partially rein-

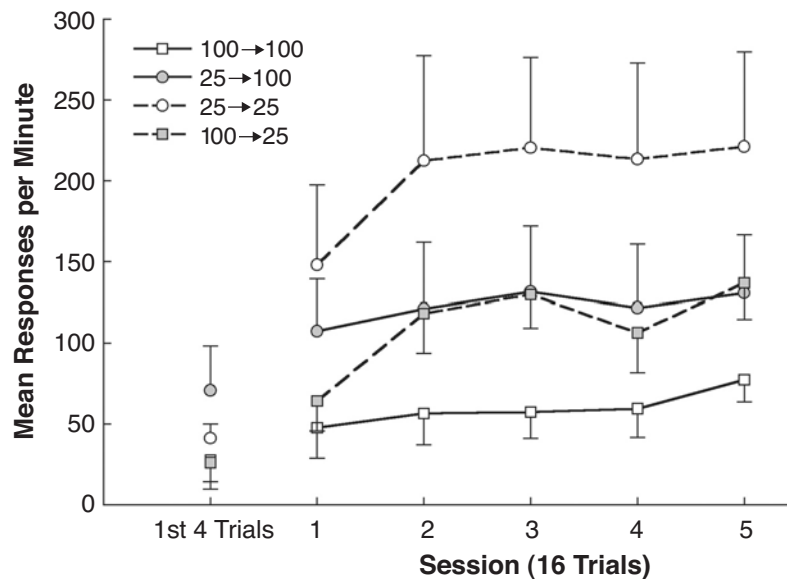


Figure 4. Responding during novel stimulus training. The data point on the left side of the graph represents responding over the first 4 trials. Other data points represent responding blocked by 16-trial sessions. Because all the animals received the same reinforcement schedule in the novel stimulus training, all the trials, and not only the trials following reinforcement, are included. Error bars represent *SEMs*, plotted unidirectionally.

forced stimulus. These results are at odds with those results of Gibbon et al. (1980), who found no convincing relationship between faster acquisition and higher probabilities of reinforcement.

As was discussed previously, the existence of this regularity is a necessary prediction of time accumulation models that assume that animals learn rates of reinforcement. According to these models, when a partial reinforcement schedule is generated by deleting reinforcers from a continuous reinforcement schedule, the effect is to lengthen proportionately the time until reinforcement during both presentation of a stimulus and the absence of the stimulus. Learning about the predictive value of a stimulus should not be affected by this manipulation. On the other hand, associative models of learning do not predict any regular covariation of probability of reinforcement and ITI. Many associative models strongly emphasize the role of nonreinforcement in learning while placing less emphasis on the role of ITI.

Although nonreinforcement is often strongly emphasized in associative models, the predicted effects of nonreinforcement in a partial reinforcement schedule can be surprisingly small early in acquisition, because, for many such models, learning on each trial is governed by the discrepancy between an actual outcome and an expected outcome. If a reinforcer is not strongly expected, nonreinforcement will lead to a relatively small discrepancy, and there will be little learning. For instance, at a point in learning at which the associative strength is one fifth of asymptotic value, error correction models calculate

the discrepancy on reinforced trials to be four times that on nonreinforced trials. The consequence is that a nonreinforcement has substantially less impact than a reinforcement early in learning. Therefore, such models anticipate that effects of nonreinforcement during early acquisition might be difficult to detect.

Figure 5 shows simulated learning curves generated from one of the most prominent trial-based error correction models, the Rescorla–Wagner (RW) model. Each curve represents the associative strength generated by one of the frequencies of reinforcement used by Gibbon et al. (1980), plotted against reinforcements. The learning rate parameters for reinforcement and nonreinforcement were assumed to be equal and set to a value of .01. At points early in acquisition, only the group receiving the lowest (10%) schedule of reinforcement shows strong signs of slowed learning, relative to 100% reinforcement. To illustrate this, consider the last point in which all the groups, except for Group 10%, differ in associative strength by less than .05. This is represented by the vertical line in Figure 5. At this point, Group 100% is already predicted to have associative strength that is 23% of asymptotic. Group 33% is even closer to maximal predicted learning, having associative strength that is 55% of asymptotic. Although the specific values are parameter dependent, this general pattern is stable. For example, with a learning rate parameter for nonreinforcement that is half as large, Group 100% and Group 33% are predicted to be at 32% and 53% of asymptotic learning at the same point of divergence as that employed above.




Figure 5. Simulated Rescorla–Wagner model acquisition curves for the five probabilities of reinforcement used by Gibbon, Farrell, Locurto, Duncan, and Terrace (1980). Associative strength is plotted against reinforcements, and not against trials. The simulated learning rate parameters for reinforcement and nonreinforcement are both .01. The vertical line represents the last point in which all groups except Group 10% are separated by less than .05.

Because nonreinforcement has such a small effect early in acquisition, differences among groups receiving different probabilities of reinforcement might not be evident if assessment is made before the groups are predicted to diverge.

The criterion employed by Gibbon et al. (1980), in which learning is said to have taken place when responding occurs on three out of four consecutive trials, may well assess learning at a point in acquisition at which associative models predict negligible between-groups differences. Consistent with this notion, the subjects in Gibbon et al.'s experiment appeared to be responding at a very low rate at the point of reaching criterion and to increase their rates substantially in subsequent sessions. In the present experiment, the subjects in Group 100% increased their response rates 70% from the session following criterial acquisition to the end of initial acquisition. The subjects in Group 25% increased their response rates 380% from the session following criterial acquisition to the end of initial acquisition.

Proponents of time accumulation models emphasize the point of conditioned response emergence, and not asymptotic performance. However, proponents of associative models, who assume a monotonic relationship between associative strength and response measures, would infer that the acquisition criterion was reached before maximal responding and, therefore, before maximal associative strength had been achieved. Therefore, it is possible that the subjects reached this particular criterion at a point in acquisition at which variations in groups receiving different frequencies of reinforcement might be hard to detect.

In addition, even the small differences in associative strength produced by varying reinforcement percentage might be masked by a 3/4 criterion that allows differential opportunities to reach criterion. Applying the criterion to *all* the observations in the present data failed to reveal a significant difference in the rates of acquisition of continuously and partially reinforced animals. To remove the bias that results from increased opportunity to reach criterion in partially reinforced animals, observations were equated by considering only a subset of the trials for Group 25%—those immediately following reinforced trials. As was expected, when the criterion was applied to this equated number of observations, it took the animals in Group 25% additional reinforcers to reach criterion. Furthermore, the difference in median number of reinforcers to reach criterion between the 100% and the 25% group was now significant. It appears that a bias in opportunities to reach criterion can and does alter the conclusions that can be made from the data.

It should be mentioned here that Papini and Overmier (1984, 1985) did find equivalent acquisition when the number of trials (but not the number of reinforcers) per session was held constant. When the number of reinforcers per session is held constant, partially reinforced groups differ from continuously reinforced groups in overall session length. When the number of trials per session is held constant, partially reinforced groups differ from continuously reinforced groups in the number of reinforcements per session. The finding of Papini and Overmier implies either that training with a relatively few reinforcements per session is beneficial for learning or that lengthening of a session is detrimental for learning.

Although it is unclear why lengthening a session might be detrimental for learning, there are reasons to believe that partially reinforced pigeons might benefit from receiving fewer reinforcements per session than do continuously reinforced animals. Consistent with this notion, McSweeney, Swindell, and Weatherly (1996) presented evidence for a within-session decrease in responding in a pigeon autoshaping procedure. Within-session decreases in responding in instrumental paradigms are often attributed to a decrease in the efficacy of the reinforcer, through either habituation or satiation (Killeen, 1995; McSweeney & Roll, 1998). Thus, reinforcers that occur later in a session might not be as potent as reinforcers that occur earlier. Holding number of trials per session constant, and not reinforcements per session, might lead to a situation in which the average reinforcement is more powerful in partially, rather than continuously, reinforced subjects. Furthermore, when trials per session are held constant, the comparison between partially and continuously reinforced animals is plagued more than in the current situation by a lack of an obvious common test during acquisition, since a given reinforcer might occur not only on a different trial number, but also on a different day. Overall, holding reinforcements per session constant, and not trials per session, seems a better controlled procedure for testing the predictions of time accumulation models pertaining to partial reinforcement.

Asymptotic Performance

The results from the first two training sessions of the present experiment suggest that acquisition proceeds more rapidly with continuous than with partial reinforcement, when acquisition is assessed across reinforcements, and not across trials. However, the pattern of responding over subsequent sessions calls into question whether the autoshaping procedure itself is appropriate for the quantitative study of partial reinforcement. As has been documented here and elsewhere, subjects in autoshaping procedures often come to peck substantially more often to a partially reinforced stimulus than to a continually reinforced stimulus when time between trials is constant. In the present experiment, this difference was apparent using both a percent CR and a response rate measure. In fact, at the end of the initial acquisition, partially reinforced animals responded at a rate close to five times as great as did continuously reinforced animals. The results of the second part of this experiment suggest that this effect is not due to differences in the local reinforcement contingencies, because it persists even when those conditions are matched. Furthermore, the results of the novel stimulus training suggest that this effect is not specific to a stimulus.

One possibility is that this differential performance is due, in part, to summation with contexts of different value. Contexts with excitatory value have been shown to modulate responding to conditioned stimuli. However, this seems an unlikely explanation for the present results, because studies with pigeons and autoshaping have found

that *higher* context values generate greater responding (Grau & Rescorla, 1985; Rescorla & Durlach, 1987). There is no evidence that the lower density of reinforcement brought about by deleting reinforcers from a continuous reinforcement schedule leads to a more excitatory context in autoshaping. Furthermore, the contextual modulation that has been seen is typically small in magnitude, as compared with the differences presented here.

Another possibility, suggested by Gibbon et al. (1980), is that subjects receiving continuous reinforcement come to approach the food hopper (goal track) near the end of the CS more often than do partially reinforced subjects. They observed that responding decreased for continuously reinforced subjects near the end of the stimulus. This could lead to continuously reinforced subjects' responding at a lower rate, when averaged across the entire stimulus. As can be seen in Figure 3, a similar effect was observed in the present experiment. Animals receiving 100% reinforcement in reversal training did show proportionately greater responding early in the stimulus than did animals receiving 25% reinforcement and proportionately less responding late in the stimulus. However, this effect does not appear large enough to explain a fivefold difference in response rate. Furthermore, groups with similar terminal response patterns (i.e., 25→25 and 100→25) still show marked differences in response rates. Still, the decline in proportional responding in groups receiving 100% reinforcement in reversal training is consistent with the possibility of increased goal tracking, and this would have the effect of decreasing the overall response rate.

Another explanation for the superiority of responding with a partial reinforcement schedule is that birds come to develop different response forms when first trained with different reinforcement schedules. Indirect evidence for a difference in response form comes from acquisition of the novel stimulus. On the first trial with the new stimulus, the subjects in all groups were disrupted, and responding was low. Rapidly, however, the groups began to differ in the same manner as they had after reversal training. Even though the initial disruption made it clear that the birds could distinguish the novel stimulus from the old stimulus, they still were found to peck at different rates with a common treatment. This is not a surprising pattern of data if it is assumed that the animals had developed different response forms in the face of the different schedules of reinforcement and that these response forms were resistant to change.

Although the partial reinforcement procedure employed here and in previous autoshaping studies has shown greater asymptotic performance with partial rather than continuous, reinforcement, the specific aspect of the partial reinforcement schedule that is critical is unknown. Deleting reinforcers to create a partial reinforcement not only decreases the probability of reinforcement on a given trial, but also lengthens the time between reinforcers. It is unclear which of these factors led to the higher response rate in partially reinforced animals in this experiment.

Regardless of the particular reason, birds are often found to peck at higher, and sometimes much higher, rates when subject to a partial reinforcement procedure than when continually reinforced. Furthermore, rate measures and probabilistic measures often correspond; when partial reinforcement leads to a higher rate of responding than does continuous reinforcement, it also leads to a higher percentage of trials with a conditioned response. This correspondence does not hold as well when probabilistic responding is close to maximal, as was the case for three of the four groups at the end of reversal training. Here, groups with similar probabilities of responding differed in terms of their rates of responding. It appears that probability of responding is not as sensitive a measure of conditioned responding as is rate of responding when animals are responding on nearly every trial. This is particularly bothersome because it means that time accumulation models predict a null result with a measure that is relatively insensitive.

Gallistel and Gibbon (2000) argued that the rate of asymptotic performance is a poor measure of learning. Rather, they suggested that the point in training at which conditioned responding emerges is the best assessment of acquisition. Determining this point of acquisition is the function of the particular learning criterion that is chosen. For a learning criterion to be a better assessment of acquisition than asymptotic response rate is, it must be assumed that there are influences affecting asymptotic response rate that do not affect responding at the point of acquisition. However, the choice of a learning criterion is made by the experimenter and is necessarily arbitrary; it may be unwise to endow this point with special properties. Rather, it seems reasonable to assume that the influences that affect asymptotic performance also affect preasymptotic performance, including the point in acquisition that the learning criterion assesses. Furthermore, the response decision rules of time accumulation models do not explicitly treat the point of CR emergence as different from any other. Time accumulation models have the same mechanism for predicting behavior at any point during training.

It should not be ignored that response topography has been shown to change over the course of training in an autoshaping procedure (e.g., Gibbon et al., 1980). Animals receiving 100% reinforcement come to respond proportionately less at the end of the conditioned stimulus than do animals receiving partial reinforcement. Furthermore, this pattern may be less evident at the point of CR emergence than later in training. It may be tempting to conclude from this that responding earlier in training is somehow a purer measure of learning than is responding later in training. However, there are two reasons to be cautious about making this conclusion. First, it is possible that whatever is influencing the differential response topography at asymptote is influencing responding at earlier points in acquisition in a different manner. That is, the influences brought about by schedule of reinforcement might be equally strong at all points in training but may manifest themselves differently. Second, it

does not appear that the within-trial timing of responding can fully account for the different between-group rates of responding seen in the present experiment. Therefore, there appears to be some other factor leading to the differential conditioned responding seen in this experiment. There is no reason to believe that this influence, whatever it may be, is confined to particular points in the learning of a CR. The change in response topography seen over time does not seem sufficient justification for the conclusion that earlier points in training are immune to factors that influence later performance.

A separate issue is whether rate of responding is useful in assessing learning at any point in acquisition. It is certainly true that time accumulation models do not make predictions about rate of responding. Rather, they predict whether or not the animal will respond. Thus, it makes sense to assess time accumulation models by using probabilistic response measures. However, in the present experiment, probability and rate of responding corresponded, except when probability of responding was an insensitive measure due to a ceiling effect. Both rate and probability of asymptotic responding were higher in partially than in continuously reinforced animals. It appears that whatever influenced the rate of asymptotic responding also influenced whether or not to respond. Theoretical considerations aside, there is nothing in the present data to suggest that probabilistic performance is a superior measure of learning than is rate of responding. Indeed, probability of responding appears to be less sensitive than rate of responding, at least in detecting asymptotic between-group differences. In any case, to the degree that both response measures correspond, it is particularly striking that a slowing of acquisition by partial reinforcement can be seen in these circumstances that lead to superior asymptotic performance.

The present report has presented an analysis of the full course of acquisition in the pigeon autoshaping procedure under conditions of continuous and partial reinforcement. According to rate, probability, and criterial measures, conditioned responding emerged more rapidly under a schedule of continuous, rather than partial, reinforcement, when responding was plotted against reinforcements and not against trials. This pattern was evident despite superior asymptotic performance generated by the partial reinforcement procedure. These results generally accord with predictions of traditional associative models and are at odds with those of time accumulation models. The particular pattern of responding in which a partial reinforcement procedure produces long-lasting, non-stimulus-specific response superiority suggests caution in making quantitative comparisons involving partial reinforcement in the pigeon autoshaping procedure.

REFERENCES

- AMSEL, A. (1958). The role of frustrative nonreward in noncontinuous reward situations. *Psychological Bulletin*, *55*, 102-119.
- AMSEL, A. (1967). Partial reinforcement effects on vigor and persistence. In K. W. Spence & J. T. Spence (Eds.), *The psychology of*

- learning and motivation* (Vol. 1, pp. 1-65). New York: Academic Press.
- BARELA, P. B. (1999). Theoretical mechanisms underlying the trial-spacing effect in Pavlovian fear conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, **25**, 177-193.
- BOUTON, M. E., & SUNSAY, C. (2003). Importance of trials versus accumulating time across trials in partially reinforced appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, **29**, 62-77.
- CRAWFORD, L. L., STEIRN, J. N., & PAVLIK, W. B. (1985). Within- and between-subjects partial reinforcement effects with an autoshaped response using Japanese quail (*Coturnix coturnix japonica*). *Animal Learning & Behavior*, **13**, 85-92.
- GALLISTEL, C. R., & GIBBON, J. (2000). Time, rate, and conditioning. *Psychological Review*, **107**, 289-344.
- GIBBON, J., & BALSAM, P. (1981). Spreading associations 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.
- GONZALEZ, F. A. (1973). Effects of partial reinforcement (25%) in an autoshaping procedure. *Bulletin of the Psychonomic Society*, **2**, 299-301.
- GRAU, J. W., & RESCORLA, R. A. (1985). Role of context in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 324-332.
- HILGARD, E. R., & MARQUIS, D. G. (1940). *Conditioning and learning*. New York: Appleton-Century.
- HUMPHREYS, L. G. (1940). Distributed practice in the development of the conditioned eyelid reaction. *Journal of General Psychology*, **22**, 379-385.
- KILLEEN, P. R. (1995). Economics, ecologies, and mechanics: The dynamics of responding under conditions of varying motivation. *Journal of the Experimental Analysis of Behavior*, **64**, 405-431.
- LATTAL, K. M. (1999). Trial and intertrial durations in Pavlovian conditioning: Issues of learning and performance. *Journal of Experimental Psychology: Animal Behavior Processes*, **25**, 433-450.
- MACKINTOSH, N. J. (1974). *The psychology of animal learning*. New York: Academic Press.
- MACKINTOSH, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, **82**, 276-298.
- MC SWEENEY, F. K., & ROLL, J. M. (1998). Do animals satiate or habituate to repeatedly presented reinforcers? *Psychonomic Bulletin & Review*, **5**, 428-442.
- MC SWEENEY, F. K., SWINDELL, S., & WEATHERLY, J. N. (1996). Within-session changes in responding during autoshaping and automaintenance procedures. *Journal of the Experimental Analysis of Behavior*, **66**, 51-61.
- PAPINI, M. R., & OVERMIER, J. B. (1984). Autoshaping in pigeons: Effects of partial reinforcement on acquisition and extinction. *Revista Interamericana de Psicología*, **18**, 75-86.
- PAPINI, M. R., & OVERMIER, J. B. (1985). Partial reinforcement and autoshaping of the pigeon's key-peck behavior. *Learning & Motivation*, **16**, 109-123.
- PEARCE, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review*, **94**, 61-73.
- RESCORLA, R. A. (1988). Behavioral studies of Pavlovian conditioning. *Annual Review of Neuroscience*, **11**, 329-352.
- RESCORLA, R. A., & DURLACH, P. J. (1987). The role of context in intertrial interval effects in autoshaping. *Quarterly Journal of Experimental Psychology*, **39B**, 35-48.
- 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.
- SCHLOSBERG, H. (1934). Conditioned responses in the white rat. *Journal of Genetic Psychology*, **45**, 303-335.
- SPENCE, K. W. (1960). *Behavior theory and learning*. Englewood Cliffs, NJ: Prentice-Hall.
- SUTTON, R. S., & BARTO, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497-537). Cambridge, MA: MIT Press, Bradford Books.
- WAGNER, A. R., & BRANDON, S. E. (1989). Evolution of a structured connectionist model of Pavlovian conditioning (AESOP). In S. B. Klein & R. R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 149-189). Hillsdale, NJ: Erlbaum.
- WASSERMAN, E. A., DEICH, J. D., HUNTER, N. B., & NAGAMATSU, L. S. (1977). Analyzing the random control procedure: Effects of paired and unpaired CSs and USs on autoshaping the chick's key peck with heat reinforcement. *Learning & Motivation*, **8**, 467-487.
- WASSERMAN, E. A., HUNTER, N. B., GUTOWSKI, K. A., & BADER, S. A. (1975). Autoshaping chicks with heat reinforcement: The role of stimulus-reinforcer and response-reinforcer relations. *Journal of Experimental Psychology: Animal Behavior Processes*, **104**, 158-169.

(Manuscript received May 1, 2003;
revision accepted for publication March 4, 2004.)