

Partial reinforcement in autoshaping with pigeons

J. GIBBON

*New York State Psychiatric Institute, New York, New York 10032
and Columbia University, New York, New York 10027*

and

L. FARRELL, C. M. LOCURTO, H. J. DUNCAN, and H. S. TERRACE
Columbia University, New York, New York 10027

The acquisition, maintenance, and extinction of autoshaped responding in pigeons were studied under partial and continuous reinforcement. Five values of probability of reinforcement, ranging from .1 to 1.0, were combined factorially with five values of intertrial interval ranging from 15 to 250 sec for different groups. The number of trials required before autoshaped responding emerged varied inversely with the duration of the intertrial interval and probability of reinforcement, but partial reinforcement did not increase the number of reinforcers before acquisition. During maintained training, partial reinforcement increased the overall rate of responding. A temporal gradient of accelerated responding over the trial duration emerged during maintenance training for partial reinforcement groups, and was evident for all groups in extinction. Partial reinforcement groups responded more than continuous reinforcement groups over an equivalent number of trials in extinction. However, this partial-reinforcement extinction effect disappeared when examined in terms of the omission of "expected" reinforcers.

Since Brown and Jenkins' (1968) discovery of autoshaping, in which a skeletal contact response directed to a visual signal for food comes under associative control, several studies have compared autoshaped behavior to traditional classical conditioning preparations (see Hearst & Jenkins, 1974, for a review). For example, effects of massed vs. distributed training on autoshaped responding (Gibbon, Baldock, Locurto, Gold, & Terrace, 1977; Terrace, Gibbon, Farrell, & Baldock, 1975) have been found to parallel the traditional finding of more effective training with spaced trials. The present experiment is addressed to another traditional variable in classical conditioning, partial reinforcement. There is evidence that partial reinforcement exerts effects on maintained responding (Gonzalez, 1973, 1974; Perkins, Beavers, Hancock, Hemmendinger, Hemmendinger, & Ricci, 1975; Schwartz & Williams, 1972; Wasserman, 1974; Wasserman, Hunter,

Gutowski, & Bader, 1975). The effect of partial reinforcement on acquisition of autoshaped responding has received less parametric study. Gonzalez (1973) has reported somewhat slower approach to asymptotic response levels under 25% as opposed to continuous reinforcement, and Wasserman et al. (1975) have reported that with heat reinforcement in chicks, the number of trials to the first response is a decreasing function of the percentage of reinforced trials. Wasserman, Deich, Hunter, and Nagamatsu (1977) have amplified this result and examined its relation to the "extra" unpaired signals introduced under the random control procedure (Rescorla, 1967).

The present experiment was designed to assess parametrically the effects of partial reinforcement on acquisition, maintenance, and extinction of responding, using a broad range of the percentage variable and the intertrial interval duration, known to produce large differences in the speed of acquisition. Five values of reinforcement probability, $P = .1, .33, .5, .75, \text{ and } 1.0$, were combined factorially with five values of the intertrial interval, $ITI = 15, 30, 50, 100, \text{ and } 250$ sec, resulting in 25 groups of subjects. Maintenance of responding after acquisition was studied for all groups, followed by extinction.

Since low probabilities of reinforcement result in many more trials (per reinforcer) than high probabilities, pseudoconditioning control groups were also studied to permit assessment of nonassociative

Portions of the data from this experiment were reported by L. Farrell, at the 45th Annual Meeting of the Eastern Psychological Association, 1974. This research was supported by NSF Grants BG 34095 and BNS 76-01229 and Small Grant MH 25070 (J. Gibbon, Principal Investigator), and NSF Grant GB 30781 and NIH Grant HD 00930 (H. S. Terrace, Principal Investigator). Reprints may be obtained from J. Gibbon, Department of Psychophysiology, New York State Psychiatric Institute, 722 West 168th Street, New York, New York 10032. C. M. Locurto is now at College of the Holy Cross, Worcester, Massachusetts 01610.

responding under extended exposure to unreinforced trial signals.

METHOD

Subjects

The subjects were 166 naive male and female White Carneaux pigeons, aged 6 months to 1 year. They were housed with water continuously available in individual cages in a continuously illuminated vivarium. Previous observations in our laboratory revealed no sex differences either in the rate of acquisition or in the frequency of occurrence of an autoshaped response. Accordingly, the subject's sex was ignored in assembling experimental groups. Each subject was maintained at 77%-83% of its free-feeding weight. Grain, presented during and (if necessary) following the experimental sessions, was composed of 40% vetch, 50% kaffir, and 10% hemp.

Apparatus

The experiment was conducted in four operant conditioning chambers of the type described by Ferster and Skinner (1957). The subject's portion of the chamber was 30.5 cm wide, 33 cm deep, and 30 cm high. Three walls and the ceiling were of white plastic; the fourth wall, an aluminum panel painted flat black, contained a single translucent response key and the hopper aperture. Chamber illumination was provided by an overhead 10-W house lamp.

The response key was centered on the panel 20.32 cm above the floor. In three of the chambers, the diameter of the response key was 2.54 cm; in the fourth chamber, the diameter was 2.80 cm. A force of approximately .18 N was required to interrupt the normally closed contact at the base of the key. Each contact closure provided a brief "feedback" click. Homogeneous illumination of the response keys was accomplished with an IEE stimulus display unit which projected the output of a 28-V bulb (GE 1829) through a green Wratten filter onto the rear of the response key. The luminance of the homogeneous field of green light was 1.80 cd/m².

The hopper aperture (5.08 cm on each side) was 8 cm beneath the response key. When the hopper was raised, it was illuminated from above by a 28-V bulb (GE 1829) of luminance 14.96 cd/m². A second bulb projected a beam of light from the right side of the hopper onto a photocell mounted on the left side to detect entries into the hopper when the food magazine was raised.

Each of the chambers was ventilated and contained a speaker which transmitted white noise at 80 dB re 20 μ N/m². Additional masking was provided by housing the chambers in a separate room which attenuated sound entry by approximately 35 dB.

Programming and data collection were carried out by a PDP-8/e computer system. One special use of the computer was recording response durations with a Schmitt trigger. Durations were recorded in 8-msec time bins.

Hopper Training

Prior to autoshaping, each subject was required to satisfy an eating criterion. At the start of hopper training, a naive subject

was placed in an experimental chamber in which the response key had been covered. The only source of illumination in the chamber was the hopper light. Prior to placing the bird in the chamber, a few grains of food were scattered on the floor immediately in front of the hopper aperture. Typically, subjects ate this grain almost immediately. Within approximately 30 sec, they then ate from the food hopper. Whenever the subject inserted its head into the hopper, it interrupted the beam of light directed at the photocell. Three and one-half seconds later, the hopper was lowered and the houselight illuminated the chamber. Following an interval averaging 20 sec, the hopper was raised again.

After completion of 10 eating trials, the first session was terminated upon the first sequence of 5 trials with eat latencies less than or equal to 1.5 sec. If no such sequence occurred, the session terminated after 50 eating trials. The session was also terminated if the subject did not eat within 45 min after the hopper was raised. Every subject was required to satisfy the latency criterion twice in no more than four hopper training sessions. Subjects which did not satisfy the latency criterion were excluded from the experiment. This screening resulted in an attrition rate of about 20%.

Autoshaping

Twenty-five experimental groups, each representing a different combination of average ITI and P value were studied. The ITI was variable, with values selected from a geometric series of 25 intervals (Catania & Reynolds, 1968). Three different series were derived for each average ITI and stored in the computer. During successive sessions, these series were varied irregularly. Sessions lasted until 25 reinforcers had been delivered.

Generally, sessions were scheduled on successive days, 7 days a week, at the same time each day. This could not be done in the case of the ITI = 250 sec, P = .1 group. On the average, the duration of these sessions was 18 h. Accordingly, these subjects were kept in their cages for 1 day between sessions. They were weighed and fed, if necessary, at the time their session was scheduled to begin on the following day.

Acquisition of responding was defined as the first sequence of four successive trials during which a peck occurred on at least three trials. If a subject did not satisfy the acquisition criterion during 10 sessions, it was replaced. The resulting composition of each group is shown in Table 1. Seventeen of the 25 groups consisted of 4 subjects each. The number of subjects in the remaining groups ranged from 5 to 14.

The autoshaping procedure began 1 or 2 days after each subject's final hopper training session. At various intervals, the green light was projected onto the response key for 10 sec. At the offset of the trial, either the food hopper was raised (during which time the houselight was turned off) or the intertrial interval began. This sequence, ITI, trial, and reinforcement (with probability P) was repeated until 25 reinforcements had been delivered. When P was less than 1.0, reinforcement was programmed by a probability generator. For any given session, the proportion of trials ending in reinforcement varied somewhat, but the mean probability of reinforcement over several sessions closely matched the programmed value. The hopper duration was 3.5 sec for all groups, except ITI = 250 sec, P = .1. For this group, pilot data indicated that 3.5-sec access to food was not sufficient to maintain eating

Table 1
Number of Subjects in Each Group that Satisfied Acquisition Criterion/Total Group Size

ITI (in Seconds)	P(Trials/Session)				
	.1(250)	.33(76)	.50(50)	.75(34)	1.00(25)
15	6/ 8(105)	4/4(33)	4/5(22)	5/9(18)	4/4(12)
30	8/14(168)	4/4(52)	4/4(35)	4/4(24)	4/4(18)
50	4/ 4(252)	4/4(78)	4/4(52)	4/6(36)	4/4(27)
100	4/ 4(460)	4/4(141)	4/4(93)	4/4(64)	6/6(47)
250	4/ 4(1085)*	4/4(331)	4/4(218)	6/6(149)	6/6(110)

Note—Average trials per session are indicated beside each P value in parentheses, and average session durations (in minutes) are indicated by each cell entry in parentheses. *Hopper duration = 10 sec.

during the 18-h session. Therefore, this group was given 10-sec access to food.¹

In contrast to the procedure followed during the hopper training sessions, the duration of each of the 25 reinforcements was independent of the subject's behavior. Because subjects did not always eat when the hopper was raised, the actual number of reinforcements consumed was monitored via the photocell apparatus. If a subject did not eat on at least 12 of the 25 reinforcement occasions for two successive sessions before satisfying the acquisition criterion, it was replaced.

Subjects were exposed to the autoshaping procedure for 15 maintenance sessions following the first session on which the acquisition criterion was met. Deviations from this plan were occasionally necessitated by the requirement that subjects of a particular experimental group begin each phase of the experiment at the same time. In 14 instances, subjects received either 14 ($N = 2$) or 16 ($N = 12$) maintenance sessions.

Extinction

Following autoshaping, each subject was given 10 sessions of extinction. Extinction sessions differed from the autoshaping sessions only in that reinforcement and the visual stimuli associated with reinforcement were omitted during the 25 reinforcement periods selected by the probability generator. Thus, at the offset of the trial with probability P , there occurred a brief interval which was indistinguishable from the ITI. Each extinction session terminated after 25 reinforcements would have occurred. Consequently, the average duration of an extinction session was the same as that of an autoshaping session.

Pseudoconditioning Control Groups

Pseudoconditioning control groups were studied at each of the five values of the ITI ($N = 5$ in each case). Each subject was first required to satisfy the hopper-training criterion described earlier. Pseudoconditioning sessions were identical to extinction sessions programmed with a P value of .10. Thus, approximately 250 trials occurred during each pseudoconditioning session. The number of sessions given each group approximated the number of sessions that occurred before a peck was observed in the case of the corresponding autoshaping groups.

RESULTS

Acquisition

Pseudoconditioning. Little pecking occurred in the pseudoconditioning control groups. Of the 20 subjects, 9 showed some responding, and of these, 6 showed some trial responding. In all cases, responding to the key was sporadic and never achieved consistent levels. None of the subjects met the 3/4 acquisition criterion. When trial responding did occur, it tended to occur in the first session, and so the number of trials prior to the first trial response was relatively low. Responding was somewhat more frequent with the longer ITI groups, but not reliably so. These data therefore suggest nonzero, but very low, levels of nonassociative keypecking engendered by food previously given in the context of the experimental chamber.

Autoshaping. The results of principal interest are presented in Figure 1. The median number of trials required to satisfy the acquisition criterion is shown as a function of the probability of reinforcement in the panel on the left. The panel on the right shows

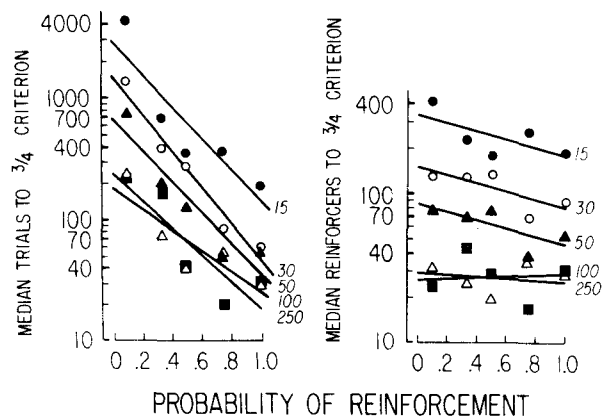


Figure 1. Log median number of trials (left panel) and reinforcements (right panel) to satisfy the acquisition criterion as a function of probability of reinforcement. The linear functions are least-squares regressions, and the parameter is ITI duration.

the median number of reinforcers prior to satisfying the acquisition criterion as a function of probability of reinforcement. The ordinate is logged and the linear functions are least squares best fits for each ITI value. For both measures of acquisition, ITI has a clear effect on the speed with which responding emerges. The linear functions for increasing ITI are displaced downward, indicating more rapid overall acquisition levels, consonant with earlier findings (Gibbon et al., 1977; Terrace et al., 1975). The data in the left panel show that the number of trials to acquisition decreased with increasing probability of reinforcement per trial. When examined in terms of the number of reinforcers required before responding emerges, however, the effect of reinforcement probability is greatly reduced. Regression analysis of the logged medians displayed in the right panel revealed that the slopes did not differ from each other and the combined slope differed (marginally) from zero [$F(1/15) = 5.14, .025 < p < .05$]. The comparable analysis of log median trials to acquisition showed no difference between slopes for different ITIs, but the combined slope differed very substantially from zero [$F(1/15) = 26.77, p < .001$].

The acquisition data for reinforcers to acquisition were subjected to an analysis of variance. A log transform was used to reduce an otherwise substantial covariation of standard deviation and mean. The results showed a large ITI effect [$F(4/87) = 27.64, p < .001$] but no effect for probability of reinforcement or interaction. When all the data from individuals were analyzed, variability within groups was sufficient to obscure the marginal probability of reinforcement effect. Evidently, only the ITI variable is a significant factor in the speed of acquisition when acquisition is examined in terms of the number of reinforcers rather than the number of trials before

responding emerges.² This point will come in for more discussion later.

Maintenance and Extinction

Response rate. Responding after the first response in a trial (running rate) for each group is shown for the 15 postacquisition sessions in Figure 2. Each panel represents a different ITI, and the five functions within each panel represent distinct probability groups. The first point to note is that responding is acquired more slowly when the ITI is short, as with the acquisition measures. This is particularly true for the 15-sec ITI groups. Responding also generally achieves higher terminal levels with the longer ITI groups, with the exception of the 250-sec ITI.

Reinforcement probability less than 1.0 results in elevated response rates within each ITI. The continuous reinforcement groups (open triangles) show substantially lower rate functions. However, within the partial groups, a graded effect is not evident. Analysis of variance of the average of the median rates during the last three sessions of maintenance training confirmed the large effect of ITI [$F(4/92) = 5.48$, $p < .001$] and a smaller effect of probability of reinforcement [$F(4/92) = 3.65$, $p < .01$].³

The relative frequency of trials with at least one response (response probability) during maintenance showed a similar trend, with a somewhat more rapid increase in response probability with long intertrial intervals. However, response probability attained generally high levels during maintenance training for all groups, and a graded effect for P was not present. Continuous reinforcement groups showed response

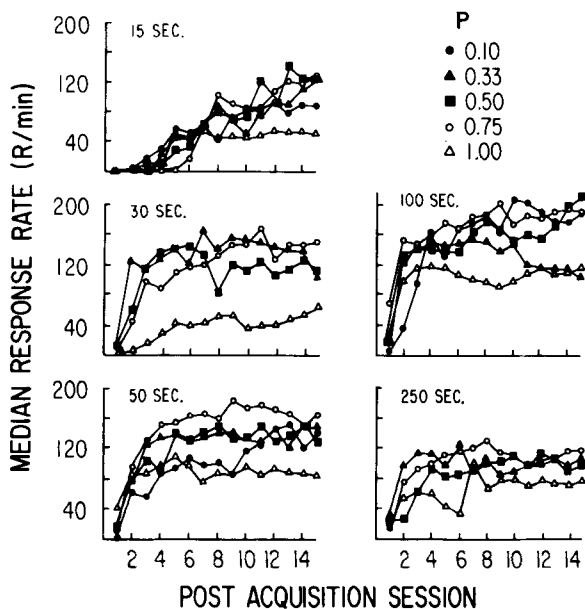


Figure 2. Median response rate over successive sessions beyond the first in which responding met the acquisition criterion. Each panel represents a different ITI, and the functions within each panel represent different probability of reinforcement groups.

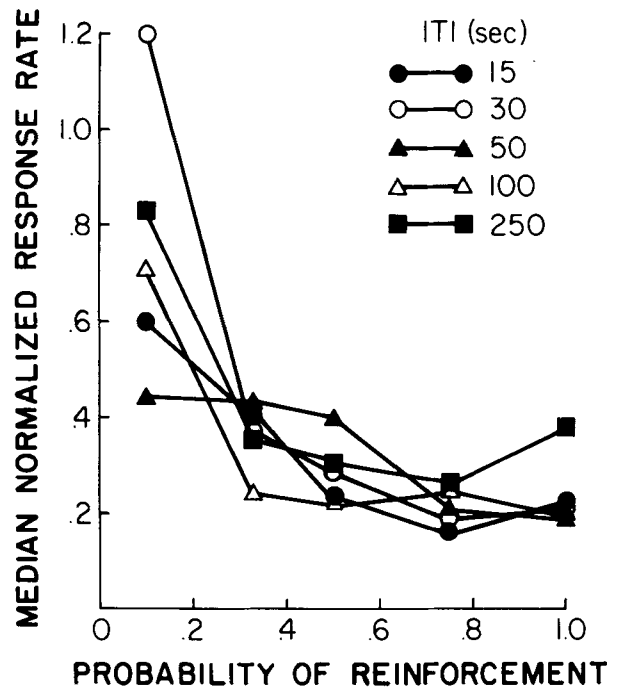


Figure 3. Median normalized response rate during the first 250 trials of extinction as a function of probability of reinforcement. Each point represents response rate as a proportion of rate at the end of maintenance training. A value of 1.0, therefore, represents a rate equal to maintenance performance. Each curve represents a different intertrial interval, with points within each curve representing different probability groups. Greater resistance to extinction is indicated by higher function levels.

probability levels comparable to intermittent groups. This difference between rate and probability may be related to the temporal distribution of responding over the trial, discussed later.

Extinction. Following the 15 sessions of maintenance, all subjects were exposed to 10 sessions of extinction, with the same number of trials per session as during conditioning. Thus, the $P = 1.0$ groups experienced 10 sessions of 25 trials each, or 250 trials; while the $P = .10$ groups experienced 10 sessions of 250 trials each, or 2,500 trials. Extinction responding was analyzed by normalizing the response rate for each subject by dividing the rate in successive extinction sessions by the average response rate during the last 3 maintenance sessions. This normalization procedure was used because of the variability among the probability groups during maintenance, evident in Figure 2. Without normalizing the data, the strongest extinction effect is that subjects with lower rates at the end of maintenance training, appear to extinguish faster, thus potentially obscuring a graded effect of probability of reinforcement.

Extinction responding was first analyzed by comparing groups during their first 250 extinction trials. Group median normalized response rate is shown in Figure 3 as a function of probability of reinforcement.

ment with ITI as a parameter. When looked at from the perspective of the first 250 trials of extinction, there is a clear partial reinforcement effect. More responding in extinction is generated by lower reinforcement probabilities during conditioning. The effect is graded with reinforcement probability at the shorter ITIs, and tends to be restricted to the lowest probability value ($P = .10$) at the two longest ITI values. An analysis of variance revealed a large effect for probability of reinforcement [$F(4/92) = 11.18$, $p < .001$]. Somewhat greater resistance to extinction with longer ITIs was also found [$F(4/92) = 2.90$, $p < .05$], and the ITI by P interaction was not significant. In terms of an absolute number of extinction trials, then, intermittent reinforcement sustains responding longer than regular reinforcement.

An alternative view of extinction after partial reinforcement might look at extinction responding as a function of the "expected" number of reinforcers omitted. That is, intermittent reinforcement in conditioning may engender an expectancy of reinforcement in extinction that is very different from regular reinforcement training. It was with this possibility in mind that extinction sessions were conducted with the same number of trials as during conditioning. Thus, each session represented the omission of 25 "expected" reinforcers. One may ask, then, whether responding declines less rapidly for the low-probability groups after an equivalent number of omissions of "expected" reinforcers. One way to answer this question is shown in Figure 4. Median normalized

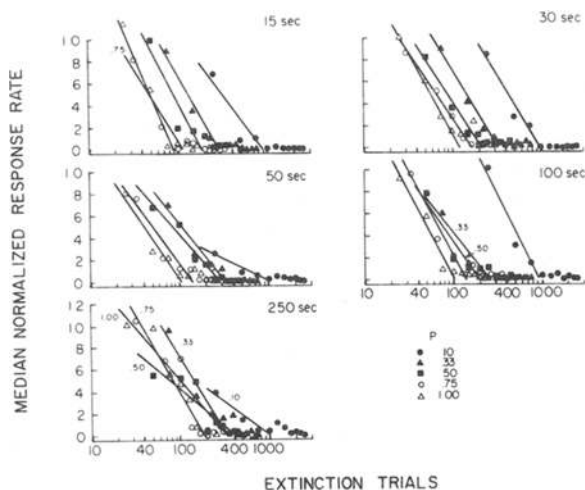


Figure 4. Median normalized response rate as a function of the log of extinction trials. Each point represents performance in blocks comprising 25 scheduled reinforcers omitted (one session). Successive panels, left to right and top to bottom, represent different increasing ITI groups, and within each panel, linear regressions are shown for the first four data points (sessions) for each probability of reinforcement group. The functions generally order themselves from left to right, with decreasing probability, and are labeled when this order is not clear. The steeper the slope of the regression function, the more rapid is extinction (the less resistance to extinction is evident).

rates over successive sessions are shown as a function of extinction trials on a logarithmic scale. A given distance at any point on this scale represents an equivalent number of "expected" reinforcers omitted for each group. For example, 10 sessions of extinction training for the $P = 1.0$ groups span the range from 25 trials (1st session) through 250 trials (10th session). This abscissa distance is the same as that spanned by the $P = .10$ groups, which ranges from 250 trials (1st session) to 2,500 trials (10th session). Each panel represents a different ITI, and within each panel five sets of data are shown corresponding to the five probability groups. The straight lines are least squares fits to the first four data points (sessions). These regression lines describe the *rate* at which responding declines in extinction when extinction trials are equated for the frequency of reinforcement delivered in training.⁴ The regression slopes characterize the decline of responding surprisingly well. The lowest correlation coefficient is .73, and all but one of the remaining correlations exceed .9. The slopes are quite comparable within panels, indicating that extinction progresses at about the same rate independently of probability of reinforcement. A small, but reliable, trend is discernible across panels, with the longer ITI groups showing shallower slopes and, thus, greater resistance to extinction.

The slopes in Figure 4 were subjected to an analysis of variance.⁵ The result was a reliable effect of ITI [$F(4/16) = 4.38$, $p < .025$] and no effect of probability of reinforcement.

Response probability data for extinction showed the same trends as seen in Figures 3 and 4. A partial reinforcement effect was obtained when response probability over the first 250 trials of extinction was analyzed, and there was no effect of probability of reinforcement when the data were examined in terms of sessions, as in Figure 4. The response probability data were more variable than the response rate data, however, and the ITI effect observed in Figure 4 was not reliable for response probability.

In summary, this extinction analysis shows that, while intermittent reinforcement generates more extinction responding *per trial*, it does not do so per "expected" reinforcer. When looked at from the perspective of an equivalent number of reinforcers omitted, the decline of responding over extinction is comparable across different probabilities of reinforcement within a given ITI. The ITI variable, in contrast, exerts a reliable effect on extinction responding. Long ITIs generate greater resistance to extinction just as they generate more rapid conditioning.

Timing. Following acquisition, a reliable temporal pattern of responding over the trial was observed for all groups. Figure 5 shows the mean proportion of responses occurring in successive fifths of the trial during the 1st session with at least 25 responses (top row), during the 15th session of maintenance (middle

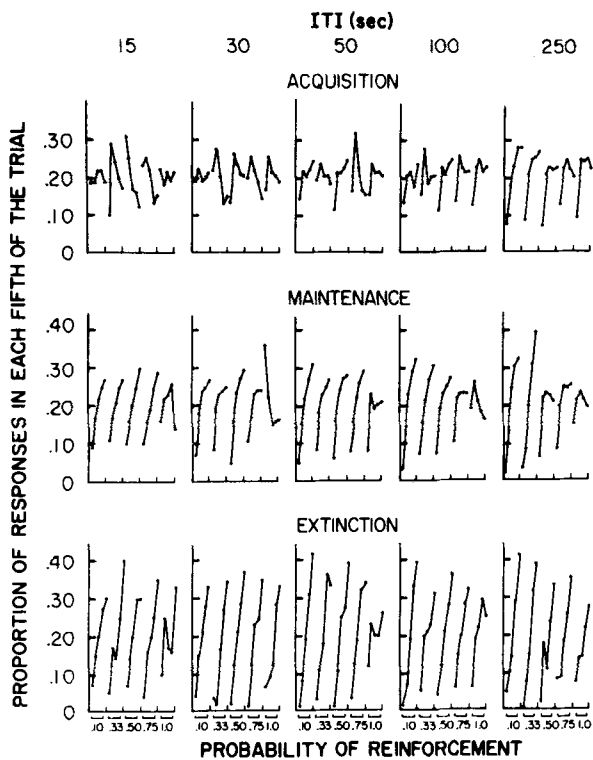


Figure 5. Proportion of total responses in successive fifths of the trial during the 1st day on which at least 25 responses were made (top row), the 15th session of maintenance training (middle row), and the last extinction session with at least 25 responses (bottom row). The five functions within each panel represent the five probability groups for a given ITI, and successive panels (columns) represent increasing ITI values.

row), and during the last extinction session with at least 25 responses (bottom row). Within each panel, the five functions represent the five probability groups for a given ITI. Points within each function represent successive fifths of the trial. A flat function would represent no temporal discrimination of trial duration. This form is approximated early in training. With the exception of the first fifth of the trial, particularly for the long ITI groups, responding is distributed approximately uniformly over the remainder of the trial duration.

By the end of training, acceleration of responding during the trial is evident for most of the groups (middle row). Responding is low at the outset of the trial and increases as the trial interval elapses. The temporal gradient is sharper for low probability of reinforcement and tends to flatten at high probabilities of reinforcement. Frequently, the $P = 1.0$ groups even show an inverted U shape, with peak rates in the middle of the trial. Compare, for example, the functions for continuous reinforcement in the 100- and 250-sec ITI groups with those for intermittent reinforcement. This finding of apparent lack of temporal control as reinforcement probability

increases is most pronounced with the long ITI groups and less clear at the shorter ITI values. For example, a very sharp gradient is obtained for $P = .10$ and $.33$ at the 250-sec ITI, but responding is low only for the first fifth of the trial for the higher probability groups at this ITI. For the 15-sec ITI, in contrast, the temporal gradient is quite sharp for all groups except the regular reinforcement group.

These data suggest that, after prolonged training, timing the trial duration is paradoxically less precise when reinforcement is more predictable at its end. However, data from the extinction sessions in the bottom row argue for an alternative view. When reinforcement was no longer forthcoming at the end of any trial, the depressed responding observed at the end of the trial for high-probability groups disappeared. The $P = 1.0$ groups during extinction responded at the end of the trial about as the intermittent groups responded at the end of the trial during training and extinction. This suggests that the flattening of the timing functions during training was a result of anticipation of food. Timing of the trial duration may have been just as precise for the continuous reinforcement groups, but the response distribution reflected this temporal discrimination differently. When food was always delivered at the end of the trial, the subjects may have engaged in more hopper-directed activity, thus reducing terminal response rates.

The depressed responding at the end of the trial for the continuous reinforcement groups is probably implicated in the finding of lower overall response rate for these groups at the end of maintenance training. If the lower response rate at the end of the trial under continuous reinforcement is ascribable to timing (Figure 5), the failure to find an effect of probability of reinforcement on response probability is not surprising. Probability of reinforcement evidently does not determine whether or not responding will occur in a trial, but rather whether, during the trial, responding will accelerate.

Response latency. The latency of the first response in a trial generally decreased with extended training and increased during extinction. Mean latency during the first three sessions of acquisition, the last three sessions of maintenance training, and the first three sessions of extinction is shown in Figure 6. Each row represents a given probability value, and each column a given ITI. Within each panel, the first three connected points are mean latencies for the first three sessions of acquisition, the next three points are latencies during the last three sessions of maintenance training, and the three points to the right of the dashed line are latencies during the first three sessions of extinction. Latency effects may be observed for both independent variables. As probability of reinforcement increases, latency is reduced. Increasing

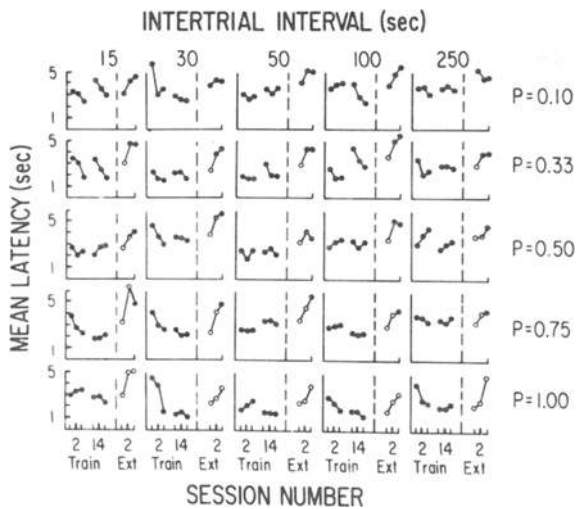


Figure 6. Mean latency of the first response in the trial during the first three sessions of acquisition, the last three sessions of maintenance training, and the first three sessions of extinction. Columns represent different ITI groups, and rows different probability groups. Latencies during extinction are shown to the right of the dashed line in each panel, and open points represent approximately the first 75 extinction trials.

ITI also is associated with somewhat shorter latencies, though the effect is not large and is not clear at all P values. The ITI effect appears to be restricted to maintenance training, while the probability effect is apparent early in acquisition as well. An analysis of variance of the data from the last maintenance session showed a significant P effect [$F(4/88) = 4.60, p < .01$] and ITI effect [$F(4/88) = 3.45, p < .01$] and no interaction.

During extinction, latencies increased rapidly. The increase did not vary systematically with either ITI or P when analyzed in terms of sessions. Note that, as with the response rate analysis, these data may be thought of as latency per 25 omitted reinforcers. Thus the failure to find an effect when the data are looked at over sessions suggests that there may be a partial reinforcement effect when the data are examined over trials. This is shown graphically in Figure 6 as the difference between open points in the extinction functions for successive rows. The three open points on the bottom row represent the first 75 extinction trials for the $P = 1.00$ groups; the two open points in the next row represent the first 66 trials for the $P = .75$ groups; the first 1.5 open points in the middle row represent the first 75 trials for the $P = .50$ groups; and the first point in the second row represents the first 75 trials for the $P = .33$ groups. Looking just at the open points, it is clear that there is a substantial increase in latency during extinction for the continuous groups in the bottom row. This increase is smaller in the $P = .75$ groups, and not evident at all for the first 75 trials

in the other groups. The increase in latency reflects more rapid extinction for the continuous reinforcement groups than for the intermittent reinforcement groups, when examined at a comparable number of trials. Again, therefore, a partial reinforcement extinction effect is evident on a per-trial, but not a per-session, basis.

The latency effects were small in absolute size. Latency ranged from about 1.5 sec at the long ITIs and high probabilities to about 3 sec at the short ITIs. Since the longest latencies averaged about 3 sec and since the trials were 10 sec long, latency effects would be largely restricted to the first third of the trial. This means that the timing distributions in Figure 5 do not reflect the development of long latencies during maintenance training. In any case, maintenance training tended to increase response speed. Just the reverse appears in the timing distributions, with response rates early in the trial becoming reduced with training rather than increasing, as shortened latencies would predict. Evidently the short-latency control over the first response in a trial is not matched by short interresponse times for subsequent responding.

Response duration. Response durations also changed over the course of training. Early in training, response durations were low, during maintenance they increased, and during extinction they again decreased. Analysis of variance revealed a very large effect of training [$F(2/150) = 109.3, p < .001$] and no effect for probability of reinforcement, ITI, or any interaction. The median response duration averaged over ITI and P was: 30.2 msec (SD = 18.7) for the 1st session on which the acquisition criterion was met, 55.4 msec (SD = 28.7) for the 15th session of training, and 35.1 msec (SD = 17.9) for the first 3 sessions of extinction. It is noteworthy that the timing of the trial, which only developed over extended training, was maintained in extinction, while the long peck durations, which also developed during maintenance, disappeared rapidly during extinction.

DISCUSSION

Acquisition

The intertrial interval in this experiment exerted a powerful and uniform effect on the speed with which autoshaped keypecking emerged. This finding is consonant with our previous results (Gibbon et al., 1977; Terrace et al., 1975) and with the spaced vs. massed training effects common in classical conditioning procedures (Gormezano & Moore, 1969; Kimble, 1961; Terrace, 1973). The intertrial interval effect is not surprising in the light of previous work, and it is evident whether one regards the appropriate measure of acquisition as the number of trials or the number of reinforcers before keypecking emerges.

Both measures decrease with increasing ITI in about the same ratio independently of probability of reinforcement (Figure 1).

The probability of reinforcement variable exerted a retarding effect on acquisition speed when that speed was assessed in terms of the number of trials required before reliable responding occurred. However, when acquisition speed was assessed in terms of the number of reinforced trials required before responding, little or no effect was evident. This finding is consonant with some data in the autoshaping literature in which both acquisition measures are reported (or inferable). Gonzalez (1973) reported autoshaping acquisition for two groups of birds exposed to 100% reinforcement and 25% reinforcement. His procedure was comparable to the present procedure, in that a fixed number of reinforcers (100) per session were programmed for each group. The 100% birds received 100 trials per session and the 25% group received 400 trials per session. The intertrial interval was the same for each group. The groups were small ($N = 3$ each) and statistical comparisons were not reported, but the median number of trials to the first peck was 172 for the 100% group and 263 for the 25% group, suggesting retarded acquisition per-trial under partial reinforcement. However, the median number of reinforcers for the 25% group was 66—*lower* than the median reinforcers to the first peck for the 100% group. It is likely from Gonzalez's report that a criterion measure would have equated the two groups in terms of reinforcers to acquisition, since the partial group appeared to achieve asymptote more slowly than the 100% group.

Wasserman, Hunter, Gutowski, and Bader (1975) reported acquisition scores for chicks responding to a key-light signal paired with heat reinforcement. They found that the median number of trials before the first response was 36, 23, and 13.5 for groups receiving probabilities of $P = .33, .67,$ and $1.0,$ respectively. The number of reinforcers to acquisition were therefore about 13 (11.5, 15.0 and 13.5, respectively, as reported in Wasserman, Deich, Hunter, and Nagamatsu, 1977, from a reanalysis of these data).

In both of these experiments and in the procedures we report here, ITI was held constant across probability groups. In the Gonzalez (1973) experiment this was effected by increasing session duration (as in the present report), while in the Wasserman et al. (1975) experiment this was accomplished by reducing the number of reinforcers delivered in each session, holding session duration constant. Evidently, the two procedures produce comparable results: a per-trial, but no per-reinforcer, effect for probability of reinforcement.

An alternative procedure for programming intermittent reinforcers has been studied in a provocative

series of experiments by Wasserman et al. (1977). In their procedure, the number of reinforcers and the session duration were held constant, but intermittent groups received *added* signals which were unreinforced. This results in a decrease in intertrial interval. In their Experiment 3, for example, two groups received 50% reinforcement procedures, one with 5 paired (and 5 unpaired) trials per day and another with 10 paired (and 10 unpaired) trials per day. The comparable continuous reinforcement groups received only the 5 or 10 paired trials per day. The results showed some retardation for the 50% group with 10 pairings per day and a slight *facilitation* for the 50% group with 5 pairings per day. Overall, there was no reliable partial reinforcement effect. This means that, even though both partial groups experienced a reduced intertrial interval, retardation of acquisition was not observed. A subsequent experiment, which added many more unpaired signals to groups receiving 5 paired trials per day, revealed a retardation at partial reinforcement values of $P \leq .20$. The procedure for these latter groups is not detailed, but presumably they also involved a decrease in intertrial interval. In light of the results reported here and elsewhere on the power of the ITI variable to modulate conditioning speed, the retardation observed at low probabilities might equally well be attributed to reduction in the intertrial interval. Two later experiments in this series confirmed the lack of an effect for a 50% partial schedule, again when the intertrial interval was also reduced (compare Groups 20-0 with Groups 0-0 in Experiments 4 and 5). The measure of conditioning speed in these experiments was percent CRs. This index is different from the one reported here, and may have different properties. However, the data seem quite clear that, at best, there is no discernible effect of intermittent reinforcement on the speed of acquisition when acquisition is assessed per reinforced trial and trial spacing is held constant.

The classical conditioning literature on acquisition in other preparations with pigeons is in accord with these findings. Two studies of activity in the pigeon using Pavlovian procedures with either food (Slivka & Bitterman, 1966) or shock (Longo, Milstein, & Bitterman, 1962) also found no per-reinforcer acquisition effect. In the Slivka and Bitterman (1966) study, there was as well no per-trial effect, while in the Longo et al. work with shock, acquisition of activity responding was retarded by intermittent reinforcement on a per-trial, but not on a per-reinforcer, measure.

In the literature on other infrahuman species, the effects of partial reinforcement on acquisition are mixed. Studies of rabbit eyelid (Thomas & Wagner, 1964) or nictitating membrane conditioning (Gormezano & Coleman, 1975); rat CER condition-

ing (Brimer & Dockrill, 1966); conditioned salivation (Fitzgerald, 1963; Sadler, 1968) and heart rate in the dog (Fitzgerald, Vardaris, & Teyler, 1966); and activity in the fish (Berger, Yarczower, & Bitterman, 1965; Gonzalez, Eskin, & Bitterman, 1963; Gonzalez, Longo, & Bitterman, 1961; Gonzalez, Milstein, & Bitterman, 1962), have found somewhat slower acquisition for $P < 1.0$ when acquisition was assessed on a per-trial basis. Other studies, as well as some of the experiments within the studies cited above, have reported no retardation on a per-trial basis (Fitzgerald, Vardaris, & Brown, 1966; Vardaris & Fitzgerald, 1969).

A survey of this literature is summarized in Table 2. The studies are grouped by species and preparation, identified in the first four columns on the left. The right four columns indicate partial reinforcement findings for acquisition and extinction. A plus (+) represents a report of retardation due to partial reinforcement, a zero (0) represents no effect, and a minus (-) represents facilitation. An asterisk (*) indicates statistical reliability. For example, the top row shows that in the present experiment significantly retarded acquisition for $P < 1.0$ was found on a per-trial (+*), but not on a per-reinforcer (0), basis.⁶

Most frequently in this literature, retardation effects were reported on a per-trial basis. The effects indicated in the "per US" columns generally represent our reanalysis of the per-trial data. A minus in these columns means that *all* of the partial reinforcement data represent facilitation relative to continuous groups at the same number of reinforcement experiences (or omissions, for the extinction column). A zero means that there was some overlap between partial and continuous groups. If less than one-half of the partial reinforcement data was facilitated relative to the corresponding continuous reinforcement data, "(E)", for estimate, is placed next to the symbol. A plus means that *all* of the partial reinforcement groups remained retarded even when compared to continuous groups at comparable numbers of reinforcers.

When the partial reinforcement manipulation reduced the ITI, as in the Wasserman et al. (1977) procedure, "(I)" is placed next to the symbol to indicate this confound. In the per-reinforcer column for acquisition, this confound appears in four of the six cases of reported retardation. Thus, of the 24 experiments reviewed, 14 reported reliable retarded acquisition per trial, but only two (Fitzgerald, 1963, salivation in the dog; Gonzalez et al., 1963, activity in the mouth-breeder) showed unambiguous retardation per reinforcer. It is noteworthy that this frequency (2/24) is in the neighborhood of traditional Type I error levels. In fact, this correspondence may be closer than indicated, since no-effect findings

from several experiments within the same studies have been combined in Table 2.

Maintenance

Over the course of maintenance training, response rates in the trial gradually increased, with responding being somewhat more rapid for long ITIs than for short ITIs. The ITI effect was observable both at asymptote and as response rates approached their final levels (Figure 2). Probability of reinforcement exerted an effect on response rates primarily late in training, with lower reinforcement probabilities associated with faster responding.

A comparison of our maintenance findings with data from other infrahuman partial reinforcement classical conditioning studies is difficult on several grounds. First, the measure of response strength usually reported in this literature is percent CRs. In our data, and in the rest of the literature, percent CRs frequently do not show a difference between partially reinforced and continuously reinforced groups (Coleman & Gormezano, unpublished work as reported in Gormezano & Moore, 1969; Thomas & Wagner, 1964). We will argue below that the difference between the rate measure and the probability of response or percent-CR measure is important, and that the higher rates found in intermittent reinforcement groups here depend upon responses after the first response on any trial. A second difficulty derives from the fact that our higher rates developed after rather long maintenance training (375 trials). Generally, the focus in the partial-reinforcement classical conditioning literature is on PRE effects in acquisition or extinction, and it is rare to find long maintenance training. Those studies with extended training sometimes report no effect on asymptotic strength (e.g., Slivka & Bitterman, 1966) and sometimes a lower asymptote (e.g., Berger et al., 1965) under intermittent reinforcement (cf. review in Mackintosh, 1974). Thus, while the literature is mixed on whether a *reduction* in response strength is produced by partial reinforcement, our finding of an *increase* in response rate appears to be unique.

Temporal discrimination of the trial duration, which emerged gradually over training, is probably implicated in the partial reinforcement effect on overall rates. By the end of maintenance training, birds in the low-probability groups showed an accelerated response rate over the trial. High-probability subjects, however, tended to slow their rates of key-pecking as the end of the trial approached. This inverted U-shaped temporal gradient may be interpreted as showing control by anticipation of food at the end of the trial for high-probability groups. The temporal discrimination of reinforcement time in the high-probability groups may have "released" a summatory behavior pattern directed at the hopper,

Table 2
Survey of Partial Reinforcement Retardation Effects in the Acquisition and Extinction of Pavlovian Conditioned Responses

Study	Subject	Preparation		Acquisition		Extinction		Comments
		Response	US	per Trial	per US	per Trial	per US	
Present Experiment	Pigeon	Keypeck (Autoshaping)	Food	++	0	++	0	Experiments 3a, 4, 5 Experiment 3b Experiment 2
Gonzalez, 1973	Pigeon	Keypeck (Autoshaping)	Food	++	0			
Wasserman et al., 1975	Chick	Keypeck (Autoshaping)	Heat	++	0			
Wasserman et al., 1977	Chick	Keypeck (Autoshaping)	Heat		++(I)			
Slivka & Bitterman, 1966	Pigeon	Activity	Food	++	0(I)	0		
Longo et al., 1962	Pigeon	Activity	Shock	++	0(E)	0		
Thomas & Wagner, 1964	Rabbit	Eyelid	Air Puff	++	0	0		
Gormezano & Coleman, 1975	Rabbit	Nictitating Membrane	Shock	++	0(E)	0		
Gibbs et al., 1978	Rabbit	Nictitating Membrane	Shock			++	?(R)	U-shaped per trial function
Brimer & Dockhill, 1966	Rat	CER	Shock	++	0	++	0	Experiment 1, E vs. C-4; Experiment 2, E vs. C-4. Experiment 1, E vs. C-2; Experiment 2, E vs. C-1. Potentiated startle to CS
Wagner et al., 1967	Rat	Startle (Experiment 1) CER (Experiment 2)	Shock	+	0	0	0	
Rescorla, 1968	Rat	CER	Shock		++(I)	++	++(I)	
Fitzgerald et al., 1966a	Rat	Heart Rate	Shock	0	-	0	-	On-baseline Off-baseline Experiment 2, 4-0 vs. 2-0 and .1-0
Fitzgerald, 1963	Dog	Salivation	Acid	++	+(E)	++	0(E)	Experiment 1 Experiment 2 Experiment 1 Experiments 2 and 3 Experiment 1, 50 vs. 100S Experiment 1, 50 vs. 100L Experiments 2, 3, 4, 5
Wagner et al., 1964	Dog	Salivation	Food	++	0(E)	++	0	
Sadler, 1968	Dog	Salivation	Food	++	0(E)	0	-	
Fitzgerald et al., 1966b	Dog	Heart Rate	Shock	++	++(I)	++	0(E)	
Vardaris & Fitzgerald, 1969	Dog	Eye Blink	Air Puff	0	-	0	-	
Gonzalez et al., 1961	Goldfish	Activity	Shock	++	0(E)	0	-	
Gonzalez et al., 1962	Goldfish	Activity	Shock	0	-	0	-	
Berger et al., 1965	Goldfish	Activity	Shock	++	0(E)	0	-	
Gonzales et al., 1963	Mouthbreeder	Activity	Shock	++	++(I)	0	-	
				++	++	++	0(E)	

*Results statistically reliable.

On-baseline

which was ethologically more appropriate to the imminent consumption of grain. Evidence in support of this interpretation is the uniform acceleration over the trial observed in *all* probability groups when food was omitted during extinction. Thus, subjects may have "known" when the trial would end, both in extinction and in maintenance training, but expressed that knowledge differently when food was forthcoming rarely or frequently at the end of the trial. This interpretation is consonant with the overall rate reduction in the $P = 1.0$ groups also, since inhibition of keypecking toward the end of the trial would lower average rates. It is not presently clear, however, whether the temporal gradient differences account for all of the rate differences or whether intermittency per se generates more vigorous responding.

A speculative, though intriguing, possibility with respect to the timing data parallels a well-established finding for partial reinforcement in instrumental alley-running behavior of rats. Rats run faster during the early portions of an alley under intermittent than under continuous reinforcement (Amsel, MacKinnon, Rashotte, & Surridge, 1964; Goodrich, 1959; Wagner, 1961; Weinstock, 1958). In the terminal section of the alley, running speeds tend to be lower for partial than for continuous groups. In the present experiment, if we equate early portions of an alley with early portions of the trial duration, then increased vigor in this section of the "temporal runway" is consonant with the partial reinforcement effect in the alley. Later portions of the trial duration were associated with decreasing response rates when reinforcement probability was high. This is the reverse of the increased vigor observed for high probability in the runway situation. This discrepancy is resolved, however, in terms of the postulated anticipatory, hopper-directed behavior. Continuous reinforcement in the autoshaping situation may increase the strength of anticipatory responses as opposed to, and perhaps incompatible with, keypecking near the end of the "temporal runway."

Overall running rate and temporal discrimination are both reflections of multiple responses in a trial. Two of the measures we have reported here reflect only the first response in each trial, namely its latency and its relative frequency or probability. These two measures covary, but show a different pattern of results than the timing and response rate indexes. The speed with which the first response in a trial is made ($1/\text{latency}$) was found to increase with long ITIs and increasing probability of reinforcement. The relative frequency of making at least one response per trial showed a similar pattern of increasing strength with long ITIs, and no reduction in strength with high reinforcement probability. Thus, the high response rates observed for the partial groups once responding had begun in a trial were not reflected in

the probability or latency with which the first response was made. Put another way, the strength of the first response in a trial increased with increasing probability of reinforcement, while the strength, or at least the frequency, of later responses in the trial decreased with increasing probability of reinforcement. However, if the decrement in rates late in the trial reflects an increment in anticipatory behavior, then this response also increased in strength with increasing P .

Response durations did not show reliable effects of either partial reinforcement or intertrial interval, though they confirmed Schwartz and William's (1972) and Schwartz's (Note 1) finding that duration of autoshaped pecks increases with extended training. These and related findings have prompted Schwartz to hypothesize that pecks of short duration are "reflexive" and controlled by the contingency between key light and food, while pecks of long duration are "operant" and controlled by instrumental relationships between responding and food that the autoshaping procedure permits after long training.

An alternative view of response duration effects has been proposed by Zirix and Silberberg (1978). They argue that the data adduced in favor of the two-process account of short- and long-duration pecks are equally susceptible to an interpretation in terms of variations in response strength from low to high. The present data do not bear directly on this issue. Durations were short in acquisition, longer during maintenance, and again short in extinction. In terms of the Schwartz interpretation, these findings might reflect signal contingency predominance in acquisition and operant relationships in maintenance. However, it is not at all clear why extinction should reflect primarily the signal-food contingency established in training. Certainly, if adventitious response-reinforcer relationships were responsible for the increased durations in maintenance training, these relationships must also extinguish when reinforcement is omitted. Our data would then require that the adventitious response-reinforcer relationships be weaker and extinguish faster than the signal-food relationship.

On the other hand, an interpretation of these data in terms of Zirix and Silberberg's proposal is straightforward. Responding is weak early in training, gains in strength in maintenance, and loses strength in extinction. Thus, the main effects are consonant with either view, though the response strength account seems more parsimonious.

Extinction

A persistent view of the presumed difference between classical and instrumental conditioning has been that partial reinforcement may distinguish the two. Partial reinforcement is claimed to have little

effect on resistance to extinction of classical Pavlovian CRs (Gray, 1975; Kimble, 1961, 1964; Terrace, 1973), in contrast to the greater resistance to extinction after partial reinforcement of instrumental responses (Amsel, 1958; Amsel, Rashotte, & MacKinnon, 1966; Felton & Lyon, 1966).

Our study is the first that we know of to report extinction of autoshaped keypecking after partial reinforcement training. In other infrahuman classical conditioning preparations, the literature is mixed on whether partial reinforcement retards extinction, assessed on a per-trial basis. Several studies have found little or no effect (Longo et al., 1962; Slivka & Bitterman, 1966, activity in the pigeon; Thomas & Wagner, 1964, rabbit eyelid; Gormezano & Coleman, 1975, rabbit nictitating membrane; Wagner et al., 1967, CER in the rat, Experiment 2; Fitzgerald, Vardaris, & Brown, 1966, heart rate in the rat; Sadler, 1968, salivation in the dog; Vardaris & Fitzgerald, 1969, eyeblink in the dog; Berger et al., 1965, Gonzalez et al., 1961, and Gonzalez et al., 1962, activity in the goldfish).

Several other studies have found a per-trial retardation (Gibbs, Latham, & Gormezano, 1978, rabbit nictitating membrane; Brimer & Dockrill, 1966; Wagner et al., 1967, CER in the rat; Fitzgerald, 1963; Wagner, Siegel, Thomas, & Ellison, 1964, salivation in the dog; Fitzgerald, Vardaris, & Teyler, 1966, heart rate in the dog; Gonzalez et al., 1963, activity in the mouth breeder).

Our survey of these effects is summarized in the two right-hand columns of Table 2. Nine of the 26 experiments reported retarded extinction (greater resistance to extinction) after intermittent reinforcement training. When examined at equivalent frequencies of reinforcers omitted, however, most of these studies showed *facilitated* extinction (17/26) or no effect (7/26). In only two cases did a partial reinforcement extinction effect remain when assessed per scheduled reinforcer. One of these was the study by Brimer and Dockrill (1966), in which a retarded extinction was found after CER training in the rat when continuous groups receiving an equivalent number of reinforced trials in training were compared in extinction to experimental groups on a 50% schedule. Retardation was found only when a 100% group was compared to a 50% group with *added* unreinforced trials. This is the same confound with shortened intertrial interval (I) as obtained in several of the acquisition reports. When a comparison was made on a per-reinforcer basis between groups for which the ITI remained constant (Experiment 1, E vs. C4; Experiment 2, E vs. C4), the partial reinforcement extinction effect disappeared.

The other exception is the study by Gibbs et al. (1978), which used the rabbit nictitating membrane preparation. They found an inverted U-shaped resistance to extinction function on a per-trial basis, with extinction fastest when probability of reinforce-

ment was either low or high. A per-reinforcer effect is questionable here also, however, since response levels prior to extinction differed greatly across probability of reinforcement groups. The response level confound (Hartman & Grant, 1962), indicated by "(R)" in Table 2, complicates an assessment of resistance to extinction.

Finally, at least one study has documented a reverse resistance to extinction effect per trial. Rescorla (1968) studied resistance to extinction after CER training in the rat off the baseline. He found that partial reinforcement schedules significantly *reduced*, rather than enhanced, resistance to extinction (Experiment 2, .4-0 vs. .2-0 and .1-0).

In summary, approximately one-third of the experiments we have reviewed find some retardation of extinction per trial after partial reinforcement training. With the possible exception of the U-shaped function obtained by Gibbs et al. (1978), retardation was not observed after partial reinforcement training on a per-scheduled-reinforcer basis. The present study fits nicely in this context, with a per-trial, but no per-reinforcer, effect.

The per-reinforcer perspective on extinction responding is conceptually similar to the "unit hypothesis" of extinction after partial reinforcement of instrumental behavior (Mowrer & Jones, 1945; Weissman & Crossman, 1966). The core of this view is that partial reinforcement establishes an expectation of a *rate* of reinforcement, or number of trials per reinforcement, which is different for different probability groups. When responding declines in extinction, it does so relative to this expected rate of reinforcement. We have not surveyed here the voluminous instrumental partial-reinforcement extinction literature (see Mackintosh, 1974, for an excellent review). In the present analysis, the decline in responding was taken relative to the preextinction response levels. The importance of this transform for obtaining monotone extinction effects rather than a U-shaped relation (cf. Grant & Schipper, 1952; Hartman & Grant, 1962) is presently not known. However, there is evidence in the instrumental literature that transforming to relative scores does not eliminate the partial-reinforcement extinction effect (e.g., Wagner, 1961). Also, interpolating continuous reinforcement prior to extinction eliminates the response level confound, yet the retardation of extinction for partial reinforcement subjects persists (Jenkins, 1962; Theios, 1962). Thus, while a "unit hypothesis" analysis will, perforce, attenuate the retardation effect in the instrumental case, it may not eliminate it entirely.

Role of Unreinforced Trials

Theoretical accounts of the conditioning process quite generally ascribe some special effects to unreinforced trial signals. In some accounts (e.g., Amsel,

1962; Capaldi, 1966), unreinforced trials generate effects via frustrative nonreward or via schedule-induced contextual cues (e.g., Neely & Wagner, 1974), which then become conditioned excitors in a manner similar to the nominal trial signal.

The influential theoretical account of conditioning offered by Rescorla and Wagner (1972) at one level regards unreinforced trial episodes as extinction occasions for the conditioned strength of both the signal and the background context cues. The result is a lower asymptotic strength for the signal, which depends directly on P .⁷ Similarly, contingency theory (Gibbon, Berryman, & Thompson, 1974; Rescorla, 1967) regards the role of unreinforced trial signals as equally important in that they reduce the correlation of the CS with the US.

In all of these conceptualizations, unreinforced trials play an important role in what is learned in a partial schedule, and effects may be expected to appear both in maintenance training and in extinction. The effects we have observed, assessed per trial, are consonant with these theoretical perspectives. The data show that partial reinforcement retards acquisition and retards extinction. However, the failure to find partial reinforcement effects when assessed per scheduled reinforcer poses a problem for all of these views. One would have to argue that the retarding effect of a given number of unreinforced trial signals is nicely balanced by the excitation accruing to endogenous, frustration-induced stimuli so that in extinction the two effects just counterbalance. It is less clear, however, why acquisition of conditioned strength is not retarded per reinforcer. Perhaps intermittent reinforcement generates schedule-induced cues only later in training, and early unreinforced trial episodes are in some sense "ignored" (cf. Amsel, 1958).

The present data suggest a potentially simpler view of the role of unreinforced trial signals in acquisition and extinction. This view stems from our prior finding of strong control by the *ratio* of the trial duration to the intertrial duration (Gibbon et al., 1977). Responding in the autoshaping situation was acquired after about the same number of reinforcements, when subjects were studied at widely different absolute trial and intertrial durations which maintained the same ratio. If one regards subjects as treating the entire period between reinforcers as a single conditioning episode, then partial reinforcement schedules result in the same ratio of trial-to-intertrial time as continuous reinforcement schedules. For example, 50% reinforcement at a 10-sec trial and 50-sec ITI results in 20 sec of trial time between reinforcers and 100 sec of ITI time between reinforcers—a 5 to 1 ratio. This ratio is readily seen to remain constant at any value of reinforcement probability, since both trial and intertrial time are multi-

plied by $1/P$. On this view, subjects assess the predictive value of the signal by the ratio of the total time that the signal appears between reinforcers to the total time that it is off between reinforcers—or, equivalently, total signal time to total time between reinforcers. Of course, in this context, it is important that at least one trial signal be contiguous with reinforcement, since trace procedures are known to produce severe retardation of autoshaping (Newlin & LoLordo, 1976; Wasserman et al., 1977, Experiment 2).

The failure to find extinction effects per reinforcer is similarly straightforward on this view. Intermittent reinforcement in training establishes an expectation of a *rate* of reinforcement, or total signal time per reinforcement. In extinction, the frequency with which the expectation of reinforcement is violated is just the frequency with which reinforcers are omitted relative to the training schedule. In the present procedures, scheduled reinforcement frequency was the same per session in training and extinction, perhaps maximizing the possibility that extinction rates per scheduled reinforcer would be equivalent.

While the "expected reinforcer" analysis is suggestive, a quantitative account is not presently available. If extinction is regarded as progressing at a constant rate per "violation of expectancy," then something like the surprise value of omitting a scheduled reinforcer appears to be involved. The core of the difficulty may be thought of as the spread of association in time. Evidently, subjects and experimenters are not always in agreement as to what constitutes a trial in the partial reinforcement experiment.

REFERENCE NOTE

1. Schwartz, B. *Studies of reflexive and operant key pecks in the pigeon*. Eastern Psychological Association, Symposium on Autoshaping and Conditioning Theory, 1976.

REFERENCES

- AMSEL, A. The role of frustrative nonreward in noncontinuous reward situations. *Psychological Bulletin*, 1958, **55**, 102-119.
- AMSEL, A. Frustrative nonreward in partial reinforcement and discrimination learning: Some recent history and a theoretical extension. *Psychological Review*, 1962, **69**, 306-328.
- AMSEL, A., MACKINNON, J. R., RASHOTTE, M. E., & SURRIDGE, C. T. Partial reinforcement (acquisition) effects within subjects. *Journal of the Experimental Analysis of Behavior*, 1964, **7**, 135-138.
- AMSEL, A., RASHOTTE, M. E., & MACKINNON, J. R. Partial reinforcement effects within subject and between subjects. *Psychological Monographs: General and Applied*, 1966, **80**, 1-39.
- BALSAM, P. D., & PAYNE, D. Intertrial interval and unconditioned stimulus durations in autoshaping. *Animal Learning & Behavior*, 1979, **7**, 477-482.
- BERGER, B. D., YARCZOWER, M., & BITTERMAN, M. E. Effect of partial reinforcement on the extinction of a classically conditioned response in the goldfish. *Journal of Comparative and Physiological Psychology*, 1965, **59**, 399-405.

- BRIMER, C. J., & DOCKRILL, F. J. Partial reinforcement and the CER. *Psychonomic Science*, 1966, **5**, 185-186.
- BROWN, P. L., & JENKINS, H. M. Autoshaping of the pigeon's keypeck. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 1-8.
- CAPALDI, E. J. Partial reinforcement: A hypothesis of sequential effects. *Psychological Review*, 1966, **73**, 459-477.
- CATANIA, A. C., & REYNOLDS, G. S. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 327-383.
- FELTON, M., & LYON, D. O. The post-reinforcement pause. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 131-134.
- FERSTER, C. B., & SKINNER, B. F. *Schedules of reinforcement*. New York: Appleton-Century-Crofts, 1957.
- FITZBERALD, R. D. Effects of partial reinforcement with acid on the classically conditioned salivary response in dogs. *Journal of Comparative and Physiological Psychology*, 1963, **56**, 1056-1060.
- FITZGERALD, R. D., VARDARIS, R. M., & BROWN, J. S. Classical conditioning of heart-rate deceleration in the rat with continuous and partial reinforcement. *Psychonomic Science*, 1966, **6**, 437-438. (a)
- FITZGERALD, R. D., VARDARIS, R. M., & TEYLER, T. J. Effects of partial reinforcement followed by continuous reinforcement on classically conditioned heart rate in the dog. *Journal of Comparative and Physiological Psychology*, 1966, **62**, 483-486. (b)
- GIBBON, J., BALDOCK, M. D., LOCURTO, C., GOLD, L., & TERRACE, H. S. Trial and intertrial durations in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, 1977, **3**, 264-284.
- GIBBON, J., BERRYMAN, R., & THOMPSON, R. Contingency spaces and measures in classical and instrumental conditioning. *Journal of the Experimental Analysis of Behavior*, 1974, **21**, 585-605.
- GIBBS, C. M., LATHAM, S. B., & GORMEZANO, I. Classical conditioning of the rabbit nictitating membrane response: Effects of reinforcement schedule on response maintenance and resistance to extinction. *Animal Learning & Behavior*, 1978, **6**, 209-215.
- GONZALEZ, F. A. Effects of partial reinforcement (25%) in an autoshaping procedure. *Bulletin of the Psychonomic Society*, 1973, **2**, 299-301.
- GONZALEZ, F. A. Effects of varying the percentage of key illuminations paired with food in a positive automaintenance procedure. *Journal of the Experimental Analysis of Behavior*, 1974, **22**, 483-489.
- GONZALEZ, R. C., ESKIN, R. M., & BITTERMAN, M. E. Further experiments on partial reinforcement in the fish. *American Journal of Psychology*, 1963, **76**, 366-375.
- GONZALEZ, R. C., LONGO, N., & BITTERMAN, M. E. Classical conditioning in the fish: Exploratory studies of partial reinforcement. *Journal of Comparative and Physiological Psychology*, 1961, **54**, 452-456.
- GONZALEZ, R. C., MILSTEIN, S., & BITTERMAN, M. E. Classical conditioning in the fish: Further studies of partial reinforcement. *American Journal of Psychology*, 1962, **75**, 421-428.
- GOODRICH, K. P. Performance in different segments of an instrumental response chain as a function of reinforcement schedule. *Journal of Experimental Psychology*, 1959, **57**, 57-63.
- GORMEZANO, I., & COLEMAN, S. R. Effects of partial reinforcement on conditioning, conditional probabilities, asymptotic performance, and extinction of the rabbit's nictitating membrane response. *Pavlovian Journal of Biological Science*, 1975, **10**, 13-22.
- GORMEZANO, I., & MOORE, J. W. Classical conditioning. In M. H. Marx (Ed.), *Learning: Processes*. Toronto: Macmillan, 1969.
- GRANT, D. A., & SCHIPPER, L. M. The acquisition and extinction of conditioned eyelid responses as a function of the percentage of fixed-ratio random reinforcement. *Journal of Experimental Psychology*, 1952, **43**, 313-320.
- GRAY, J. A. *Elements of a two-process theory of learning*. London: Academic Press, 1975.
- HARTMAN, T. F., & GRANT, D. A. Effect of intermittent reinforcement on acquisition, extinction, and spontaneous recovery of the conditioned eye-lid response. *Journal of Experimental Psychology*, 1962, **63**, 217-226.
- HEARST, E., & JENKINS, H. M. *Sign-tracking: The stimulus reinforcer relation and directed action*. Austin, Tex: Psychonomic Society, 1974.
- JENKINS, H. M. Resistance to extinction when partial reinforcement is followed by regular reinforcement. *Journal of Experimental Psychology*, 1962, **64**, 441-450.
- KIMBLE, G. A. *Hilgard and Marquis: Conditioning and learning* (2nd ed.). New York: Appleton-Century-Crofts, 1961.
- KIMBLE, G. A. Categories of learning and the problem of definition. In A. W. Melton (Ed.), *Categories of human learning*. New York: Academic Press, 1964.
- LONGO, N., MILSTEIN, S., & BITTERMAN, M. E. Classical conditioning in the pigeon: Exploratory studies of partial reinforcement. *Journal of Comparative and Physiological Psychology*, 1962, **55**, 983-986.
- MACKINTOSH, N. J. *The psychology of animal learning*. New York: Academic Press, 1974.
- MOWRER, O. H., & JONES, H. Habit strength as a function of the pattern of reinforcement. *Journal of Experimental Psychology*, 1945, **35**, 293-311.
- NEELY, J. H., & WAGNER, A. R. Attenuation of blocking with shifts in reward: The involvement of schedule generated contextual cues. *Journal of Experimental Psychology*, 1974, **102**, 751-763.
- NEWLIN, R. J., & LoLORDO, V. M. A comparison of pecking generated by serial, delay, and trace autoshaping procedures. *Journal of the Experimental Analysis of Behavior*, 1976, **25**, 227-241.
- PERKINS, C. C., JR., BEAVERS, W. O., HANCOCK, R. A., JR., HEMMENDINGER, P. C., HEMMENDINGER, D., & RICCI, J. A. Some variables affecting rate of key pecking during response-independent procedures (autoshaping). *Journal of the Experimental Analysis of Behavior*, 1975, **24**, 59-72.
- RESCORLA, R. A. Pavlovian conditioning and its proper control procedures. *Psychological Review*, 1967, **74**, 71-80.
- RESCORLA, R. A. Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, 1968, **66**, 1-5.
- RESCORLA, R. A. Informational variables in Pavlovian conditioning. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 6). New York: Academic Press, 1972.
- RESCORLA, R. A., & WAGNER, A. R. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current research and theory*. New York: Appleton-Century-Crofts, 1972.
- SADLER, E. W. A within- and between-subjects comparison of partial reinforcement in classical salivary conditioning. *Journal of Comparative and Physiological Psychology*, 1968, **66**, 695-698.
- SCHWARTZ, B., & WILLIAMS, D. R. Two different kinds of key-peck in the pigeon: Some properties of responses maintained by negative and positive response-reinforcer contingencies. *Journal of the Experimental Analysis of Behavior*, 1972, **18**, 201-216.
- SLIVKA, R. M., & BITTERMAN, M. E. Classical appetitive conditioning in the pigeon: Partial reinforcement. *Psychonomic Science*, 1966, **4**, 181-182.
- TERRACE, H. S. Classical conditioning. In J. A. Nevin & G. S. Reynolds (Eds.), *The study of behavior*. New York: Scott, Foresman, 1973.
- TERRACE, H. S., GIBBON, J., FARRELL, L., & BALDOCK, M. D. Temporal factors influencing the acquisition and maintenance of an autoshaped keypeck. *Animal Learning & Behavior*, 1975, **3**, 53-62.

- THEIOS, J. The partial reinforcement effect sustained through blocks of continuous reinforcement. *Journal of Experimental Psychology*, 1962, **64**, 1-6.
- THOMAS, E., & WAGNER, A. R. Partial reinforcement of the classically conditioned eyelid response in the rabbit. *Journal of Comparative and Physiological Psychology*, 1964, **58**, 157-158.
- VARDARIS, R. M., & FITZGERALD, R. D. Effects of partial reinforcement on a classically conditioned eyeblink response in dogs. *Journal of Comparative and Physiological Psychology*, 1969, **67**, 531-534.
- WAGNER, A. R. Effects of magnitude and percentage of reinforcement and number of acquisition trials on conditioning and extinction. *Journal of Experimental Psychology*, 1961, **62**, 234-242.
- WAGNER, A. R., SIEGEL, S., THOMAS, E., & ELLISON, G. D. Reinforcement history and the extinction of the classical reward response. *Journal of Comparative and Physiological Psychology*, 1964, **58**, 354-358.
- WAGNER, A. R., SIEGEL, L. S., & FEIN, G. G. Extinction of conditioned fear as a function of percentage of reinforcement. *Journal of Comparative and Physiological Psychology*, 1967, **63**, 160-164.
- WASSERMAN, E. A. Stimulus-reinforcer predictiveness and selective discrimination learning in pigeons. *Journal of Experimental Psychology*, 1974, **103**, 284-297.
- WASSERMAN, E. A., DEICH, J. D., HUNTER, N. B., & NAGAMATSU, L. S. Analyzing the random control procedure: Effects of paired and unpaired CSs and USs on autoshaping the chick's key peck with heat reinforcement. *Learning and Motivation*, 1977, **8**, 467-487.
- WASSERMAN, E. A., HUNTER, N. B., GUTOWSKI, K. A., & BADER, S. A. Autoshaping chicks with heat reinforcement: The role of stimulus-reinforcer and response-reinforcer relations. *Journal of Experimental Psychology: Animal Behavior Processes*, 1975, **104**, 158-169.
- WEINSTOCK, S. Acquisition and extinction of a partially reinforced running response at a 24-hour intertrial interval. *Journal of Experimental Psychology*, 1958, **46**, 151-158.
- WEISSMAN, N. W., & CROSSMAN, E. K. A comparison of two types of extinction following fixed-ratio training. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 41-46.
- ZIRIAX, J. A., & SILBERBERG, A. Discrimination and emission of different key-peck durations in the pigeon. *Journal of Experimental Psychology: Animal Behavior Processes*, 1978, **4**, 1-21.

NOTES

1. Balsam and Payne (1979) have shown that variation in rein-

forcement duration has little or no effect on acquisition of auto-shaped responding.

2. The relation between these two measures is that, on the average, number of trials to acquisition times P equals number of reinforcers to acquisition. In a sense, then, the trials-to-acquisition measure is confounded, since it contains the independent variable as a multiplier of the dependent variable.

3. These rates did not show a covariation of mean and standard deviation ($r = .04$), and so the log transform was not necessary here.

4. To accomplish equating for expected number of reinforcers, the abscissa scale might also have been simply extinction sessions. However, the regression fits were not as good with this variable.

5. Individual subjects' data were not used in this analysis, since regression fits did not characterize the decline in responding of individual subjects well. The analysis therefore did not contain repeated measures, and the interaction term was used as the estimate of error variance.

6. While no-effect findings from several experiments within the same study have been combined in some instances, main-effect findings (but not interactions) are, in every case, listed separately. These distinctions have not been uniformly observed in other literature surveys and result in some discrepancies between our Table 2 and the review by Gormezano and Moore (1969, Table 2) and Mackintosh (1974, p. 72). Some additional discrepancies we are unable to reconcile. Brimer and Dockrill (1966) and Thomas and Wagner (1964) both report reliable pertrial PREs, while Wagner et al. (1967, Experiment 2) and Wagner et al. (1964) do not report acquisition data. Despite the discrepancies, these surveys and our own present a similar overall picture.

7. The asymptotes of associative value for a partially reinforced cue, X, superimposed on a background, A, are

$$V_A(\infty) = 0, V_{AX} = V_X = \frac{P\beta_1}{P\beta_1 + (1-P)\beta_2},$$

where P is probability of reinforcement, and β_1, β_2 are the learning rates associated with reinforcement and extinction, respectively. The asymptotes for reinforcement and extinction are assumed to be 1.0 and 0. There is a typographical error in the statement of this result in Rescorla and Wagner, 1972 (p. 89) and in Rescorla, 1972 (p. 21). The denominator is a sum, not a difference.

(Received for publication October 21, 1977;
revision accepted April 18, 1979.)