Pavlovian Contingencies and Temporal Information

Peter D. Balsam
Barnard College, Columbia University, and New York State
Psychiatric Institute

Stephen Fairhurst
New York State Psychiatric Institute

Charles R. Gallistel
Rutgers, The State University of New Jersey

The effects of altering the contingency between the conditioned stimulus (CS) and the unconditioned stimulus (US) on the acquisition of autoshaped responding was investigated by changing the frequency of unsignaled USs during the intertrial interval. The addition of the unsignaled USs had an effect on acquisition speed comparable with that of massing trials. The effects of these manipulations can be understood in terms of their effect on the amount of information (number of bits) that the average CS conveys to the subject about the timing of the next US. The number of reinforced CSs prior to acquisition is inversely related to the information content of the CS.

Keywords: Pavlovian conditioning, contingency, uncertainty, time, information

One of the important observations guiding learning theory for the past 35 years is that despite temporal contiguity between the conditioned stimulus (CS) and the unconditioned stimulus (US), if USs are also presented during the intertrial interval (ITI), conditioning is attenuated (Rescorla, 1968). If the rate of USs is the same in the presence and absence of the CS, as in a random control procedure, excitatory responding is diminished or prevented relative to conditioning without added USs (see Papini & Bitterman, 1990, for a review). Most modern theories attribute the deficit to strong context conditioning interfering with the development of the conditioned responses (CRs). For example, in the Rescorla and Wagner (1972) model, the presentation of the unsignaled USs during the ITI causes contextual cues to gain associative value. These contextual cues then compete with the CS for associative value on CS trials in which the US is presented. One would expect that the higher the rate of reinforcement in the absence of the CS, the more conditioning to the context and the less conditioning to the cue. Performance based comparator models make a similar prediction (Denniston, Savastano, & Miller, 2001; Gallistel & Gibbon, 2000; Gibbon & Balsam, 1981). In these models, subjects learn about both the context (background) and the cue. Performance is based on a comparison of the two values, and the strength of the CR is related to the extent to which the CS value is greater than the context value. For example, in scalar expectancy theory

(SET; Gibbon & Balsam, 1981), the asymptotic CS value was assumed to be determined by the delay to reward during a trial (T), and the background value was determined by the overall cycle time (C) between USs. Acquisition speed depended on the ratio of the context to trial value (C/T). Unsignaled ITI reinforcers should decrease C and thus slow acquisition. Support for a performance component comes from numerous studies that document excitatory associations between CS and US following exposure to the random control procedure (Benedict & Ayres, 1972; Kirkpatrick & Church, 2004; Kremer, 1971; Kremer & Kamin, 1971; Papini & Bitterman, 1990) even after excitatory CRs are no longer evident (Delamater, 1996; Rescorla, 2000). It is thus plausible that a good part of the detrimental effect of unsignaled ITI USs might be a performance effect produced by increasing the background value relative to the CS value.

In this article, we developed a variant of the idea that the strength of excitatory conditioning depends on the comparison between the expectation of reinforcement in the CS to the reinforcement expectation in the background (Gallistel, 2003; Gallistel & Gibbon, 2000; Gibbon & Balsam, 1981). The kernel of the approach is that acquisition of excitatory CRs is a function of information about *when* the US will be presented. It is assumed that subjects rapidly learn the times from one US to the next and the time from the onset of the CS until the US (Balsam, Drew, & Yang, 2002; Drew, Zupan, Cooke, Couvillon, & Balsam, 2005; Kirkpatrick & Church, 2000). We then modeled acquisition as a function of the difference between the temporal uncertainty after the onset of a CS and the background temporal uncertainty about the time of the next arrival of the US. The CS onset reduces the subject's uncertainty about when the next US will occur. The reduction in uncertainty is measurable by formulae that derive from Shannon's (1948) definition of information. This information conveyed by the CS about the expected time of the US determines the speed with which CRs emerge. Our approach quantifies intuitions about the key role that informativeness of the CS plays in conditioning (Egger & Miller, 1962; Gibbon, 1981; Kamin, 1969;

Rescorla, 1968). Consistent with these intuitions, a consideration of the information that CS onset provides about the time of the next US integrates the known effects of most parameters of simple conditioning protocols on the rate of acquisition of a CR.

First, we experimentally characterized the effects of changing contingencies by adding unsignaled USs to the ITI on the speed of acquisition in autoshaping and then developed the information-theoretic analysis.

Experiment 1

Though all of the models are similar in their qualitative predictions as to how unsignaled USs should affect acquisition, they differ in detail. In particular, rate expectancy theory (RET; Gallistel & Gibbon, 2000) makes the prediction not shared by the other theories that ITI reinforcers should have a catastrophic effect on acquisition.

In this view, acquisition depends on the ratio of the rate of reinforcement in the CS to the rate of reinforcement in the background. The model was developed to deal with a wide range of experiments, none of which actually presented ITI reinforcers. It was assumed that in the absence of background USs, the best estimate of the background rate was the inverse of cumulative time in the ITI, which is an estimate of the expected interval between reinforcements in the absence of a CS. Because the rate estimates are based on cumulative reinforcers divided by cumulative time, just two ITI reinforcers will double the background rate. Three reinforcers will triple it and so on. Thus, even a few ITI reinforcers ought to have a big impact on acquisition speed. The RET modeling indicated that acquisition occurred when the rate in the CS was about 300 times larger than the background estimate. If this is correct, then even a single unsignaled reinforcer during each ITI would mean that the average ITI duration would have to be 300 times longer than the CS duration for acquisition to occur.

This large effect of unsignaled USs seems unrealistic given that subjects do acquire CRs under partial positive contingencies (Hallam, Grahame, & Miller, 1992; Rescorla, 1968). Nevertheless, it seemed important to test this prediction in the same preparation that provided all of the data for the RET modeling (autoshaping). Additionally, the previous studies of partial positive contingencies were done in aversive conditioning paradigms. Thus, the current studies are, to our knowledge, the first ones to have manipulated the ITI reinforcement rate in appetitive conditioning. Finally, previous research has not quantified the effect of intertrial (background) reinforcements on the rate of conditioning, that is, on the number of CS–US pairings required for the appearance of a CR. These experiments are part of a long-standing effort to determine quantitatively the effects of key protocol parameters on the rate of conditioning (Gallistel & Gibbon, 2000; Gibbon & Balsam, 1981).

The first experiment assessed whether adding occasional unsignaled ITI reinforcers has the profoundly detrimental effect predicted by RET on the acquisition of conditioned pecking in the pigeon autoshaping paradigm. In both Experiments 1A and 1B, one group of pigeons was autoshaped in a standard procedure, and a second group of subjects was exposed to the identical procedure but with unsignaled reinforcers presented during the ITI.

Method

Subjects. Thirty-seven experimentally naive, white Carneaux pigeons (*Columba livia*) participated in the experiment. The birds were maintained at 85% of their free-feeding weight throughout the experiment.

Apparatus. The experimental chambers were six standard Lehigh Valley Electronics (Laurel, MD) pigeon conditioning chambers. Each chamber was 30 cm long, 34 cm wide, and 34 cm high. An aluminum wall of the chamber had three response keys, each 2.5 cm in diameter and mounted 25 cm above a mesh floor. Keys could be transilluminated by IEEE projectors, which came with the chambers. Only the center key was operative in this experiment. An aperture (5 cm × 5 cm), centered 10 cm above the floor, provided access to a solenoid-operated grain hopper. A 15-W houselight was mounted in the same wall, 32 cm above the floor. An externally mounted fan provided ventilation and masking noise. The chamber was housed in a light- and sound-attenuated enclosure.

Procedure. All subjects were first trained to eat from the feeder. The food hopper was activated on a variable-time 50-s schedule. When grain was presented, the hopper was held in the up position until there was a head entry and for 3.5 s thereafter. The hopper training session ended when a subject inserted its head in the feeder, with latencies less than 5 s on five consecutive trials. The session also ended if 50 reinforcements were presented without reaching criterion or if the subject went 1 hr without a head entry. When criterion was met in two consecutive sessions, subjects began the autoshaping experiment in a different experimental chamber. During autoshaping, the hopper was raised for a fixed 3.5 s during each reinforcer presentation.

Experiment 1A. This experiment consisted of two groups of subjects that were exposed to an autoshaping procedure in which each of twelve 21-s presentations of a keylight was followed by grain presentation. For half of the subjects, the keylight was red; for the other half, it was a white *X* on a dark field. The 8 subjects in Group 0 received no additional reinforcers. For this group, the average ITI timed from the offset of grain until the next CS presentation was 96 s, with individual ITIs drawn randomly from an exponential distribution. The 9 subjects in Group 4.6 received 12 additional unsignaled grain presentations during the ITI. For this group, the 96-s ITIs were generated by sampling from a list of 25 exponentially distributed durations, with a mean of 48. At the end of the first sampled interval, a US was delivered. The distribution was then resampled, and at the end of the second interval, a reinforced CS was presented. Thus, every ITI contained an unsignaled US. In this group, the rate of reinforcement during the CS was 4.6 times greater than the background reinforcement rate, as computed by RET (Gallistel & Gibbon, 2000). This ratio would not be expected to generate acquisition, according to the RET analysis. Subjects were run until they acquired the key-peck response, which was defined as a session in which they responded on three quarters of the trials.

Experiment 1B. This experiment combined partial reinforcement (for both groups) with background (ITI) reinforcement (for one group). Conceptually, sessions consisted of 120 12-s periods, during a random 10 of which the key was illuminated with a red light (the CS). In each session, four randomly chosen keylight periods terminated in grain presentation; thus, 4 out of 10 CS presentations were reinforced. For the 13 subjects in Group 0, there were no unsignaled reinforcers. For the 6 subjects in Group 7.3, there were 6 unsignaled reinforcers; that is, 6 of the 110 12-s intervals without a CS presentation were followed by unsignaled hopper presentations. For this group, rate of reinforcement in the CS was 7.3 times the background reinforcement rate in the RET computation, a ratio that also would not be expected to generate acquisition. Training continued until at least five sessions after key pecking had been acquired.

Data analyses. We examined acquisition with three independent methods to be sure that our conclusions were not specific to a single method of characterizing acquisition speed. First, we analyzed how many trials it took each subject to meet a criterion of pecks on three out of four consecutive trials, the criterion used in our previous studies of acquisition (Gallistel & Gibbon, 2000; Gibbon & Balsam, 1981).

Second, we used an algorithm that finds change points in cumulative records (Gallistel, Fairhurst, & Balsam, 2004). First, we computed the cumulative record of responses versus trials. If the average rate (responses per trial) remains constant, then the slope of the cumulative record will be constant. If the average rate changes at some point, then there will be a change in the slope of the cumulative record at that point. The algorithm proceeds datum by datum through the cumulative record testing for the occurrence of a change in slope. For the third and all subsequent data points in the record, the algorithm finds a putative change point (i.e., a would-be change point) prior to that datum. This putative change point is the previous datum that deviates maximally from a straight line drawn from the origin of the plot (or previous change point once one is detected) to the current cumulative record datum. It divides the prior data into those up to and including the putative change point and those after it up to and including the current datum. The algorithm then calculates the log of the odds against the null hypothesis that the observations on the two sides of the putative change point come from the same rate process.

For this computation, we used the binomial cumulative distribution function, $B(k, N, T_a/T)$, where N is the total number of pecks (up to the current datum), k is the number of pecks after the putative change point, and T_a/T is the probability of any one peck randomly occurring after the putative change point. This probability is the proportion of trials after the putative change point (T_a) to total trials (T); the smaller this proportion is, the less likely it is (on the null hypothesis) that any given peck in a sample of N pecks will happen to fall after the putative change point. The sensitivity of the algorithm to short-lasting changes depends on the choice of the decision criterion, which is the odds ratio, $(1 - p)/p$, at which the null hypothesis (no change) is rejected. The lower the decision criterion, the more sensitive the algorithm is to short-lasting changes. Long-lasting changes are detected more or less regardless of the decision criterion, because the longer a change lasts, the stronger the evidence for it becomes. That is why changing the criterion from 100:1 to 10,000:1 had very little impact on the results we report. Below we report the change points obtained with a criterion of 10,000:1.

The third method of analyzing acquisition is to assume a continuous change from a stable, low-rate state to a stable, high-rate state (Gallistel et al., 2004). The popular linear operator models are an example of such a description, though they generally assume a zero response rate for the initial low rate. More generally, acquisition functions are often sigmoidal in form (Prokasy, 1972) and can be described by the parameters of a Weibull function, which is commonly used to summarize psychometric scatter plots. Its three parameters, a, L, and s, correspond to the three aspects of acquisition—asymptote (a), location (L) of the transition to asymptote, and the slope (s) of the function during the transition. One reason for choosing this function is that different values for its slope parameter (s) cause it to assume widely different forms, so it can approximate most monotonically increasing data sets. For values of s close to 1, it has the form of the commonly assumed negatively accelerated learning curve. As s goes to infinity, the Weibull function becomes a step function, so it can also represent a maximally abrupt transition. The protean character of the Weibull function and the fact that it independently parameterizes the apparent between-subjects differences (asymptote, location, and abruptness) make it another method well suited to a theory-neutral quantitative description of the acquisition process. Because this last method requires an asymptotic level of responding, it was not appropriate for use in the analysis of Experiment 1A, which was terminated shortly after subjects started responding.

Results

In Experiment 1A, the addition of ITI reinforcers tended to slow acquisition. However, the experiment was terminated before some subjects satisfied an acquisition criterion. The t tests on acquisition scores of those subjects that did acquire failed to show a significant difference between groups in any criterion

($ts < 1.96$). In order to deal with the indeterminate acquisition scores of subjects that did not key peck, we also used nonparametric methods to analyze these data. The left panel of Figure 1 shows the median and range (dashed vertical lines extending from the bars) of the number of CS–US pairings to meet the acquisition criterion of pecks on three out of four consecutive trials. This provides a familiar but somewhat incomplete representation of the result. A more informative and compact, if somewhat unfamiliar, presentation of the results is by means of the cumulative distributions in the right-hand panels of Figure 1. The x-axis of these distributions is the number of CS–US pairings to acquisition. The cumulative distribution steps up at each point along this axis at which a subject acquired. Thus, it shows every datum (every acquisition point). The level to which it steps is the percentage of the subjects that have acquired up to that point (within that many CS–US pairings). The percentage that failed to acquire is the difference between the asymptote of the plot and the top of the graph. For all dependent variables, we constructed the cumulative distribution of subjects that achieved a value less than or equal to a given value. We then compared experimental groups by asking whether these distributions are significantly different from one another with Kolmogorov–Smirnov tests. We set our rejection criterion at $p < .05$, unless specified otherwise.

On all measures, there was a tendency for Group 4.6 to acquire more slowly than its control, but this difference was not significant in any acquisition measure. Similarly, in Experiment 1B, Group 7.33 and Group 0 did not differ significantly in either the trials to criteria, change point, or Weibull analyses.

We also examined whether the added reinforcers had an impact on maintained response rates. This analysis was restricted to Experiment 1B because postacquisition data were not collected in the prior experiment. There was no significant difference in the mean response rates during the final training session, nor was there a difference in Weibull asymptotes.

Discussion

The effect on the rate of acquisition of a modest number of reinforcers delivered during the ITIs is small. It is nowhere near the catastrophic effect predicted by RET. Indeed, with the lower of the two rates of background reinforcement (Experiment 1B), there was no hint of a detrimental effect of ITI USs. When there was a somewhat higher rate of background reinforcement (Experiment 1A), there was a trend toward slower acquisition. This trend suggests that ITI reinforcers might produce significant retardation if they were administered at still higher rates. The purpose of the next experiment was to explore the quantitative effects of adding ITI reinforcers now that we know the range within which measurable effects may be expected.

Experiment 2

Experiment 1 demonstrated that adding a low rate of unsignaled reinforcers to the ITI did not have a devastating effect on acquisition. However, there was a trend toward slowed acquisition in Experiment 1A, and there are other experiments on partial positive contingencies that demonstrated a strong depressive effect of ITI reinforcers (Hallam et al., 1992; Rescorla, 1968). There are also many examples of random control procedures in which equal
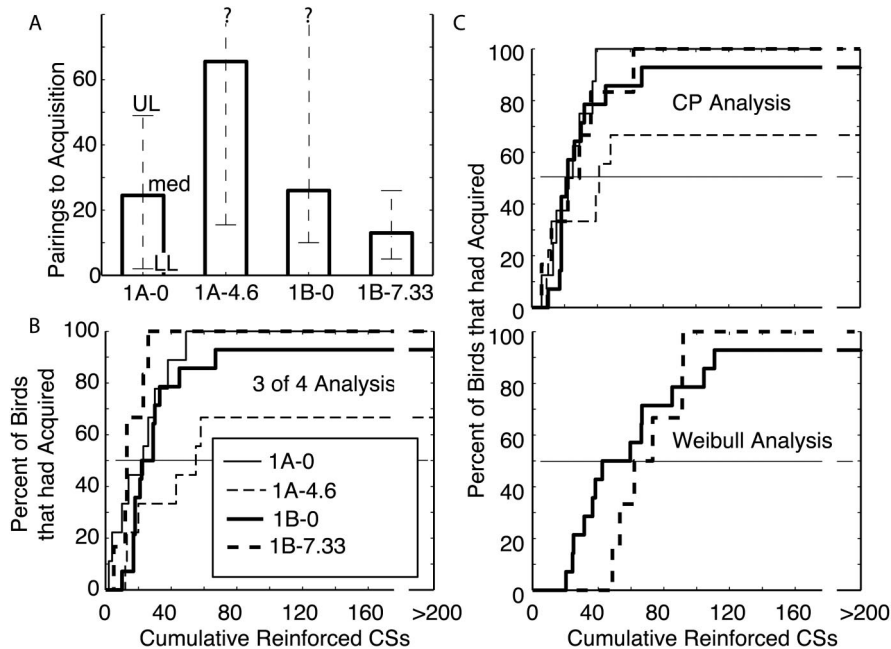
Figure 1. A: Median (med) trials to reach a criterion of pecks on three out of four consecutive trials are shown for all groups. The vertical lines represent the upper limit (UL) and lower limit (LL) of the range. B: The cumulative percentage of subjects meeting the three out of four criterion as a function of the number of CS–US pairings. C: The cumulative percentage of subjects meeting the other acquisition criteria as a function of the number of CS–US pairings. In the cumulative plots, the maximum separation is the Kolmogorov–Smirnov *D* statistic. For groups with 7–10 subjects, its critical values are near .6. Thus, if there is any point at which the vertical separation between two cumulative distributions clearly exceeds half the length of the *y*-axis, then they are probably significantly different distributions; otherwise, they are probably not. Unlike parametric statistics, this statistic is invariant under monotonic rescaling of the *x*-axis, which means that it is unaffected by assumptions about the form of the scale that relates the unobserved rate of learning to the observed trials to acquisition. It is also unaffected by assumptions about when subjects that did not acquire might have acquired. CS = conditioned stimulus; US = unconditioned stimulus; ? = undefined; 1A = Experiment 1A; 1B = Experiment 1B; CP = change point.

reinforcement rates in the CS and ITI attenuate or prevent acquisition of CRs (Rescorla, 1968; see Papini & Bitterman, 1990, for a review). The main purpose of Experiment 2 was to quantify this detrimental effect on acquisition. Specifically, we studied acquisition in procedures in which the rate of reinforcement in the CS was anywhere from twice to 16 times the background rate of reinforcement. We compared the acquisition speed in these groups with the acquisition function obtained when temporal variables are altered (C/T ratio) but no reinforcers are presented in the absence of the CS. This allowed us to directly compare the detrimental effects of adding the unsignaled ITI reinforcers with the detrimental effect decreasing the C/T ratio. Both manipulations were expected to degrade the temporal informativeness of the CS: how much its onset reduces the subject's uncertainty about when the next US will occur.

A second purpose of Experiment 2 was to evaluate the effects of feeder training on acquisition speed. Gallistel and Gibbon (2000) assumed that feeder training had little impact on acquisition. This idea was suggested by Gibbon and Balsam (1981), who analyzed whether the rate of reinforcement during feeder training had an impact on subsequent acquisition. They did not find such an effect and concluded that subjects adjusted their estimates of background reinforcement rates so rapidly that the feeder training did not have

an enduring impact on autoshaping. Though the spacing of food presentations during feeder training did not appreciably affect acquisition (Gibbon & Balsam, 1981), the number of unsignaled US presentations prior to acquisition does affect acquisition speed (Balsam & Schwartz, 1981; Balsam & Gibbon, 1988). Even a few unsignaled foods slowed acquisition relative to no pretraining. In these experiments, subjects were first trained to eat from a feeder in one context and then conditioned in a second context in which exposure to unsignaled USs was manipulated. These results raise the possibility that if feeder training and autoshaping are conducted in the same context, subjects might include the unsignaled USs during feeder training in their computation of the rate of reinforcement in the context. In all of the data modeled by Gallistel and Gibbon, feeder training occurred in the autoshaping context. Thus, Gallistel and Gibbon may have overestimated the threshold for responding. If the unsignaled USs from feeder training are included in the computation, then the ratio of CS to background reinforcement rate was much less than 300 at the time of acquisition.

Figure 2 shows data from 258 subjects run in our laboratory that had been exposed to feeder training prior to autoshaping in the same context. These include subjects from the studies summarized by Gallistel and Gibbon (2000) for which we had precise data
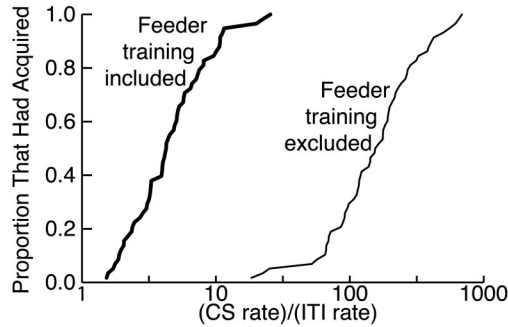
Figure 2. The cumulative percentage of subjects that met the three out of four criterion as a function of the ratio of conditioned stimulus (CS) to background rates as computed in rate expectancy theory. The lighter line excludes feeder training, and the darker line includes feeder training in the computation of the background rate. The distributions are based on 258 subjects that we had detailed records of for both feeder training and autoshaping. ITI = intertrial interval.

about their experiences both during feeder training and during autoshaping. For each subject, we computed the actual ratio of rates at the point at which each subject met the acquisition criteria. The rate of reinforcement in the background was calculated by dividing the cumulative time in the ITI by the cumulative number of unsignaled reinforcers, including those presented in feeder training. The rate of reinforcement in the CS was calculated by dividing the cumulative CS time by the number of reinforcers presented. These subjects were exposed to a very wide range of CS and ITI reinforcement rates. The figure shows the cumulative number of subjects that acquired by the time they reached a given ratio. The thin line is based on the computation excluding feeder training and shows that the median ratio for acquisition is 155. This value is lower than the original RET estimate of 300 but still quite substantial. The heavy line shows the same computation but includes the unsignaled USs presented during feeder training in the background rate. Note that the median ratio is 4.3, which is lower by a factor of 70 than the previous estimate offered by Gallistel and Gibbon. Also, note that if this were the computation used to determine acquisition, the results of Experiment 1 are consistent with RET as well as SET because the rate of reinforcement in the CS was at least 4.6 times the ITI rate. As Figure 2 shows, this ratio is sufficient for producing acquisition in most subjects.

In order to have better control over the effective number of unsignaled reinforcers, we used a procedure in Experiment 2 in which subjects never received a separate feeder training phase. All subjects were exposed to the autoshaping procedure from the start of training. This way, no unsignaled feedings were ever presented to the subjects prior to the start of the experiment proper. Different experimental groups were exposed to a range of CS and ITI values as well as several levels of unsignaled reinforcers. In this way, we could directly compare the effects of changing the ratio of CS to background rates by manipulating ITI reinforcers or by directly manipulating the C/T ratio. As noted above, RET posits that when there are no explicit ITI reinforcers, the background reinforcement rate is estimated from the cumulative nonreinforced time in the ITI. Figure 2 implies that without feeder training, as soon as the CS rate of reinforcement is a little over four times the cumulative background time, the majority of subjects should acquire the CR. Consequently, the current procedure should produce extremely

rapid acquisition in groups in which there are no unsignaled ITI reinforcers if the rate of reinforcement in the CS is compared with the cumulative ITI time. Alternatively, if the background rates are estimated from all reinforcers regardless of whether they are signaled (Gibbon & Balsam, 1981), then acquisition speed might not be very different in this procedure compared with the usual one. Furthermore, acquisition should be affected similarly regardless of whether one changes the C/T ratio by changing trial duration, ITI duration, or by adding unsignaled USs during the ITI.

Method

Subjects and apparatus. Sixty-two experimentally naive, white Carneaux pigeons (*Columba livia*) participated in the experiment. The birds were maintained at 85% of their free-feeding weight throughout the experiment. The apparatus was the same as that used in Experiment 1.

Procedure. The experiment was conducted in three replications. During the first two replications, all sessions lasted for 25 reinforcer presentations. During the third replication, all sessions lasted for 50 reinforcer presentations. We found no significant effect of replication, and we report the pooled data here.

All subjects were exposed to an autoshaping procedure from the start of training. When reinforcement was presented, the hopper was held in the up position until there was a head entry and for 3.5 s thereafter. When head entry occurred in less than 5 s for five successive reinforcements, the hold was removed and the hopper stayed up for only 3.5 s for the remainder of the session. When the bird met this hopper-conditioning criterion in two successive sessions, the hold was permanently removed.

For all subjects in the experiment, every CS was followed by grain presentation. For two groups of 12 subjects, the CS was a 10-s illumination of a red keylight. One of these groups (Group 10–0) received all of their reinforcers after a CS presentation—no unsignaled USs were presented during the ITI. The second group of subjects (Group 10–6.25) had on average a single reinforcer presented during each ITI. For the latter group, the rate of reinforcement during the CS was 6.25 times greater than the background reinforcement rate. The ITIs averaged 60 s. Exponential distributions of interevent intervals were generated by ending an interval during each second, with probability $1/M$, where M is the mean of the interval. Thus, the 60-s mean ITI in Group 10–0 was generated by ending the interval with $p = .0167$. For Group 10–6.25, an ITI period generated from an exponential distribution with a mean of 30 s was followed by an unsignaled US presentation. This was followed by a second ITI period averaging 30 s followed by a reinforced CS. In a session of 25 total reinforcers, there were 12 unsignaled USs preceded by an average of 30 s of ITI and 13 signaled USs preceded by an average of 30 s of ITI. This yields a rate of reinforcement in the CS that is 6.25 times greater than the rate in its absence.

Four other groups of subjects were exposed to a 6-s CS and an ITI that averaged 144 s. The groups differed in the rate of reinforcement during the ITI. For the 9 subjects in Group 6–0, no reinforcers were presented during the ITI, whereas for the 9 subjects in Group 6–2, the rate of CS reinforcement was only twice the background rate. Between these two extremes, two other groups of 10 subjects experienced different rates of unsignaled USs, such that the CS reinforcement rate was either 4 or 16 times greater than the background rate. Interevent intervals were generated as described for the groups with the 10-s CSs.

All subjects were trained for at least five sessions after they acquired key pecking or for a minimum of 100 CS–US pairings if they did not start key pecking. One subject in Group 10–0 did not learn to eat from the feeder and was dropped from the experiment.

Results

Figure 3 shows the number of trials for each subject to acquire the key peck response plotted as a function of the C/T ratio. The
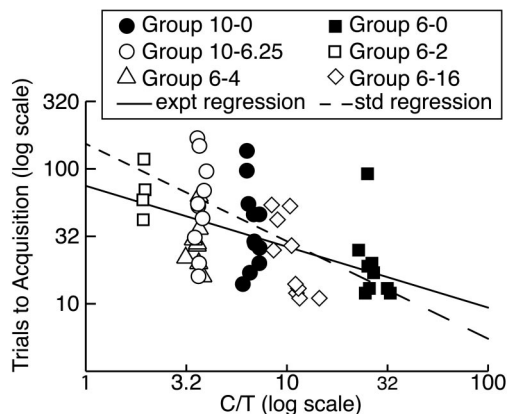
Figure 3. The number of conditioned stimulus–unconditioned stimulus pairings that each subject took before meeting the three out of four criterion is shown as a function of the ratio of cycle to trial duration (C/T) that each subject had experienced at the time of acquisition. All subjects in Experiment 2 that met criterion are shown in the figure. The solid regression line is based only on the data from Experiment 2 (expt). The dashed line is based on the subjects depicted in Figure 2 that had the standard feeder training (std) prior to autoshaping.

rate of reinforcement during the cycle was computed by dividing the cumulative time in the session by the cumulative number of reinforcers. The criterion shown in the figure is the number of CS–US pairings until each subject responded on three out of four consecutive trials. The Weibull start and the change point analyses show an identical pattern of results. Only those subjects that acquired key pecking are included in the plot. There was a significant effect of changing the C/T value, $F(5, 46) = 2.88, p < .05$.

The solid regression line shows the line of best fit for the subjects of Experiment 2. The dashed line is based on the regression from the 258 subjects (see Figure 2) exposed to the standard procedure of hopper training followed by conditioning without ITI reinforcers. A two-way analysis of variance, with C/T and presence or absence of hopper training as factors, confirmed a significant effect of C/T, $F(1, 308) = 28.66, p < .01$, and showed no significant effect of prior feeder training, $F(1, 308) < 1$. Thus, the elimination of hopper training in the current experiment did not significantly facilitate acquisition. Subjects did not acquire at the rapid rate predicted by the RET. Furthermore, though there is individual variability, there are no systematic deviations from the C/T analysis.

We also compared the effects of changing the C/T ratio by adding unsignaled ITI reinforcers with the effects of changing the ratio without ITI reinforcers by doing a two-way analysis of variance on the data from the current experiment, with C/T ratio and method of changing C/T (changing temporal parameters vs. adding ITI USs) as variables. Again, there was a significant effect of changing C/T, $F(4, 58) = 7.67, p < .01$, but no effect of the type of manipulation used to change it, $F(1, 58) < 1.0$.

We also examined possible effects of ITI reinforcers on maintained response rates. The top panel of Figure 4 shows the mean rate for each group during the last 10 trials of each subject's training. Though there was a trend toward lower rates with increased ITI USs, there were no significant differences between the groups trained with the 10-s cue or between the groups trained with the 6-s cue. The failure to find differences is the result of very

high intersubject variability. The cumulative distributions in the lower panel of Figure 4 make it evident that in each group, there were subjects that responded at very low rates (less than 0.1 responses per second) and others that responded at higher rates. The groups differed in the percentage of each type of subject. For example, in Group 10–0, 18% of the subjects had a low rate, but 33% of the subjects had a low rate in Group 10–5.8. Similarly, 11%, 40%, 20%, and 78% of the subjects in Groups 6–0, 6–16, 6–4, and 6–2, respectively, had low rates. Looked at in this way, the groups trained with the 6-s CS do differ significantly, $\chi^2(3, N = 38) = 10.43, p < .05$, in the percentage of low-rate subjects. Thus, it might be better to characterize the manipulation as affecting the likelihood that a subject will be in a low- or high-response state rather than thinking that varying contingency produces continuous changes in response strength.

Discussion

Many subjects in the groups exposed to added ITI reinforcers acquired key pecking. Hence, we did not observe the profound retardation predicted by the original formulation of RET. Nor did we find the extremely rapid acquisition predicted by our reformulation of the RET criterion based on the inclusion of feeder training. Thus, it appears that the computation proposed in RET does not accurately reflect the background reinforcement rate.

An alternative computation is proposed in SET (Gibbon & Balsam, 1981). Background reinforcement rate is computed by dividing the total number of reinforcers by the total session time. In this model, acquisition speed is a function of the rate ratio expressed as the ratio of C (cycle time: the overall delay to reinforcement in the background) to T (delay to reinforcement in the trial). The effect of adding reinforcers in the ITI is to decrease C. The current analyses indicate that there is a strong effect of changing the C/T ratio, but it does not matter whether the ratio is changed by adding ITI USs or by changing ITI and CS duration. The C/T ratio predicts acquisition speed equally well for all groups regardless of whether they are exposed to unsignaled USs. Note that we may have underestimated the effects of small C/T ratios because we excluded the subjects that did not acquire from the parametric analyses. The function relating acquisition speed to C/T may be steeper than the one we present here, but nothing in our data suggests different functions for procedures with and without ITI USs.

General Discussion

Acquisition and Rates of Reinforcement

The main finding in the current studies is that acquisition is controlled by the ratio of the rate of reinforcement during the CS to the overall reinforcement rate. The inverse of the rate ratio is the C/T ratio (Gibbon & Balsam, 1981), where C represents the expected interval between reinforcements (cycle time) and T represents the delay to reinforcement in the presence of the CS. The detrimental effect of adding reinforcers to the ITI is well described by assuming that the only effect of this manipulation is to increase the overall reinforcement rate. The implication of this view is that the added unsignaled reinforcers are no more detrimental than the effects of massing trials so long as the overall reinforcement rates are equivalent. Jenkins, Barnes, and Barrera (1981) actually re-
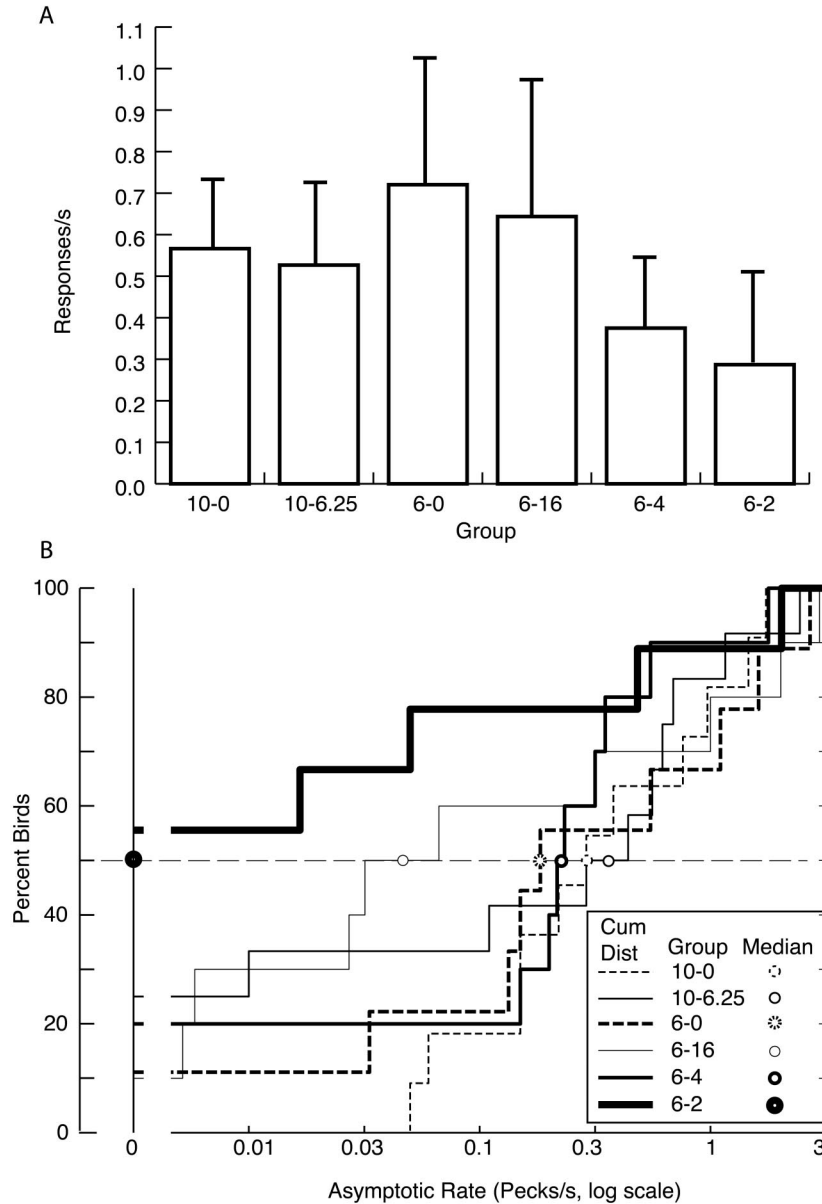
Figure 4. A: Mean rate of responding for each of the experimental groups in Experiment 2. The error bars represent the standard error. B: Cumulative distributions (Cum Dist) of subjects that responded with a rate less than or equal to a given rate (on the abscissa).

ported such a finding. In that experiment (Experiment 13), the percentage of ITI reinforcers preceded by the autoshaping cue was varied from 0% to 100%. All groups acquired key pecking at the same rate. This is consistent with the view that the detrimental effect of adding reinforcers to the ITI is mediated by changes in cycle time.

One apparent contradiction to this view comes from studies in which the ITI reinforcers are signaled by a cue that is different than the target cue (Durlach, 1983; Rescorla, 1972). If only the overall rate of reinforcement modulated acquisition, it would not matter whether the added reinforcers were unsignaled, signaled by the target cue (as in the Jenkins et al., 1981, experiment), or signaled by a different cue. However, signaling the ITI reinforcers with a

different cue does not decrement responding to the same extent as unsignaled USs (Durlach, 1983; Rescorla, 1972). To deal with this effect Cooper, Aronson, Balsam, and Gibbon (1990) posited that when the ITI cue and target cue differ, the two experiences are segregated. The differently signaled ITI reinforcers do not enter into the calculation of the overall rate for the target, and the target reinforcers do not enter into that calculation for the different ITI cue. There may indeed be different representations for both CS and background rate when there are distinct cues. Alternatively, generalization between cues may play a very important role in target acquisition. In most studies showing a facilitative effect of signaled ITI reinforcers, the ITI and target cues were similar and/or subjects were pretrained on a cue similar to the target cue (Cooper

et al., 1990; Durlach, 1983; Goddard & Jenkins, 1987). In studies using very distinct cues (different modalities), the facilitative effects of signaling ITI reinforcers are not so clear. Jenkins et al. (1981, Experiment 11) found no difference in autoshaping to a keylight between unsignaled ITI reinforcers and ITI reinforcers signaled by an auditory cue. In fact, just making the duration of the ITI cue different than the duration of the target cue reduces the facilitative effect of signaling the added reinforcers (Jakubow, Brown, & Hemmes, 2004; Williams, 1994). Future studies will have to decide the contribution of generalization from the ITI cue to the target cue in the two-cue procedure.

The current studies are consistent with the idea that acquisition is based on a comparison of the CS rate with the overall rate of reinforcement—the SET formulation. One interesting aspect of the SET formulation is its requirement that subjects learn the delays to reinforcement even before responding emerges (Balsam et al., 2002; Drew et al., 2005). It seems reasonable to assume that the ease with which a time is learned does not depend on the C/T ratio, so why does acquisition speed change with changes in this ratio? In SET, it was assumed that contexts acquired associative value much more rapidly than cues. Thus, the speed of acquisition depends mainly on how long it takes a CS to acquire a value that is about twice the background value. Below we offer an alternate hypothesis as to what is being compared and why acquisition is not always immediate.

The Importance of Information in Acquisition

The discovery of blocking and overshadowing (Kamin, 1969), relative validity (Wagner, Logan, Haberlandt, & Price, 1968), and the importance of distinguishing contingency from temporal pairing (Rescorla, 1968) led students of animal learning to speak of the conditioning process in terms that invite an information-theoretic analysis. To many, the just-cited experiments suggested that conditioning does not occur unless the CS conveys information to the subject that is not already conveyed by other CSs (cf. Gallistel, 2003). We now show that this same perspective enables one to predict the effect on acquisition of manipulating the parameters of conditioning protocols.

Figure 5 portrays variations in simple conditioning protocols similar to the ones we used in the current experiments. From an information-processing perspective, the CS comes to elicit an excitatory CR only if its onset reduces the subject's uncertainty about the onset time of the next US (Gallistel, 2003). The greater the reduction in this uncertainty at CS onset, the sooner in training the CS elicits conditioned responding. The US onsets are represented by a dot on the time line. The CS is present during the intervals occupied by the rectangle.

The objective uncertainty about the time of occurrence of point events in an interval, T, is the entropy statistic, H (see the Appendix). The differences in uncertainty (i.e., entropy) between the intervals when the CS is not present and the intervals when it is quantify the objective information that the CS conveys about the timing of the US (Garner, 1962; Shannon, 1948). If reinforcers are distributed randomly in time, as in our feeder training procedure, the uncertainty about the time of the next US is $\bar{H} = k - \log_2 \lambda$, where λ is the parameter that specifies the reinforcement rate, and k is a constant that varies with the underlying time units.

What about the uncertainty after CS onset when the CS–US interval is fixed as in the current experiments (see Figures 5D and 5E)? The information conveyed by a CS can be separated into two
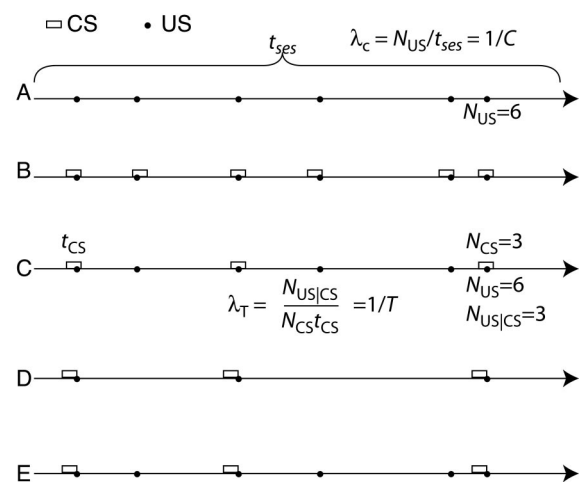

Figure 5. Basic conditioning protocols. A: The random time protocol. B: Continuous reinforcement; unconditioned stimuli (USs) occurring anywhere within the conditioned stimulus (CS), as in Rescorla (1968). C: Poor positive contingency; the CS–US contingency is degraded by USs that occur in the absence of the CS. D: Continuous reinforcement, with the CS–US latency fixed. Fixing the CS–US interval increases the information that the CS provides about the timing of the next US. E: Fixed CS–US interval and background reinforcers. The quantities used to calculate the entropies (uncertainties) are indicated as follows: C = average interval between USs; λ_T = rate of US occurrence when CS is on; λ_C = unconditional rate of US occurrence; N_{US} = number of USs; N_{CS} = number of CSs; $N_{US|CS}$ = number of USs during a CS; T = average delay to US during CS; t_{CS} = duration of a CS; t_{ses} = duration of a session.

parts (see the Appendix). First, there is information conveyed simply from the fact that the rate of US occurrence is higher in the presence of the CS. Second, there is information conveyed by virtue of the fixed interval between CS onset and US delivery. The first component depends on the rate ratios (C/T). The second component is a fixed 2 bits per reinforced CS (regardless of the CS duration). Though not tested here, one implication of this analysis is that there will be an advantage to fixed-over variable duration CSs that decreases as the C/T ratio increases.

Our hypothesis is that the number of CS–US pairings that a subject experiences before it begins to respond to the CS is inversely related to the informativeness of the CS: The greater the informativeness, the fewer the required number of reinforced CSs.

Specifically, our hypothesis is that bits to acquisition is constant for a given C/T ratio regardless of whether that manipulation is achieved through changing temporal parameters or contingencies. Bits to acquisition is CS informativeness multiplied by the number N_a, the number of reinforced CSs before the onset of responding. Thus, when the US comes at a fixed interval after CS onset, bits to acquisition is

$$N_a \left[\log_2 \left(\frac{\lambda_T}{\lambda_C} \right) + 2 \right] = N_a \log_2 \left(\frac{\lambda_T}{\lambda_C} \right) + 2N_a.$$

(Equation 3 in the Appendix)

When the CS–US interval is not fixed, one drops the $2N_a$ term.

We use Equation 3 to calculate for each bird the cumulative information conveyed by the CS up to the point at which conditioned

responding appeared. Figure 6 plots the cumulative distributions of this quantity (bits to acquisition). Overall, across a broad range of conditions, the median cumulative information at the time of acquisition is about 125 bits. Differences between groups appear to be more in the percentage of subjects that eventually acquire the CR rather than in when acquisition occurs for those that do start to respond. This is most evident in Group 6–2 in which 56% of the subjects never met the acquisition criterion. However, those in this group that did acquire were not particularly slow to do so. This is more clearly seen in Figure 7, in which all of the subjects from both experiments were segregated into three groups. One line depicts the acquisition scores for all of the subjects that received no intertrial USs. The second line depicts all subjects that received intertrial USs except for Group 6–2, and the third line replots the performance of this group alone. There is no significant difference between the first two distributions. The median of both distributions is about 125 bits. The distribution of scores in Group 6–2 is significantly different from both of the other distributions. Again, the main impact of the ITI reinforcers appears to be on the number of subjects that acquire the CR. The figure also shows the median and interquartile range for all subjects. Note that even in Group 6–2, 3 of the 4 subjects that acquired were within the interquartile range for the larger distribution. This is similar to the pattern of response rates shown in Experiment 2. Those subjects in all groups that did acquire key pecking responded with as much vigor as subjects in any of the other groups. The groups differed in the percentage of high- and low-rate subjects. It appears that the effect of varying the contingency was to modulate the number of subjects that acquired a reliable CR. Very poor positive contingencies may also slow acquisition speed, but we could not detect this within the limits of the current experiments.

Because bits are related to acquisition speed across a wide range of parametric manipulations, one might be tempted to conclude that the accumulation of those bits is what leads to the appearance of the CR. This cannot be the whole story. This conclusion ignores the fact that our calculation of bits to acquisition assumes that the predictive

Figure 7. The cumulative percentage of subjects that met the three out of four criterion as a function of the cumulative bits of information. The thin solid line depicts all of the subjects exposed to procedures without any unsignaled unconditioned stimuli (USs) during the intertrial interval (ITI). The heavy solid line depicts all subjects exposed to poor positive contingencies (PP) created by adding unsignaled USs to the ITI except for (excpt) Group 6–2. The dashed line shows Group 6–2. The vertical lines show the median and interquartile range of all subjects.

relation between the CS and the US is known a priori, which, of course, it is not. We can compute the informativeness of the CS because we know the CS–US relation. At the start of the experiment, the subject does not know this relation; it must learn it in the course of the experiment. Thus, early on, the amount of information about US timing that the occurrence of a CS provides is less than its informativeness, because of the subject's uncertainty about the consistency and durability of the predictive relation between the CS and the US and between the one US and the next (background rate). In other words, exposure to the CS–US relation is required for the subject to learn about the informativeness of the CS. An information-theoretic model of the acquisition process will need to quantify this effect of initial ignorance.

In short, the accumulation of the bits that we get when we calculate CS informativeness does not wholly explain acquisition. Nonetheless, a consideration of CS informativeness integrates a large number of findings: When CS informativeness is zero, a vigorous CR to the CS does not develop (the truly random control). When a CS whose informativeness for the subject has not already been established is combined with a CS of equal and already established informativeness, a CR to the new CS does not develop (blocking). When two CSs of unequal informativeness are presented together, the CR develops primarily (or exclusively) to the CS with greater informativeness (the effect of relative CS validity). Finally, the current analysis establishes that the effects of trial spacing (ITI), CS duration, and the background rate of reinforcement can all be understood in terms of CS informativeness.

References

Balsam, P. D., Drew, M., & Yang, C. (2002). Timing at the start of associative learning. *Leaning and Motivation, 33,* 141–155.

Balsam, P. D., & Gibbon, J. G. (1988). Formation of tone–US associations does not interfere with the formation of context–US associations. *Journal of Experimental Psychology: Animal Behavior Processes, 14,* 401–412.

Balsam, P. D., & Schwartz, A. (1981). Rapid background learning in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes, 7,* 382–393.
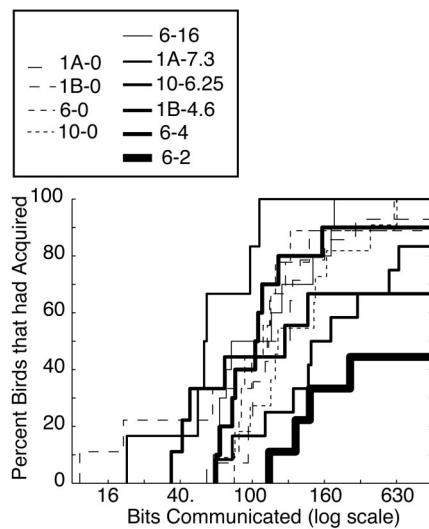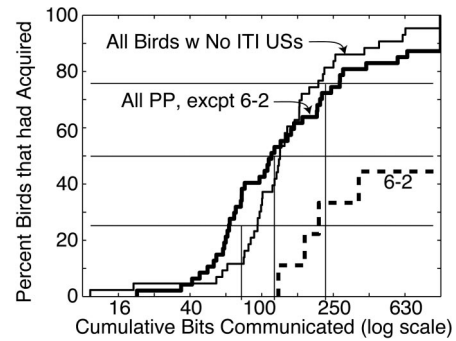
Figure 6. The cumulative percentage of subjects that met the three out of four criterion as a function of the cumulative bits of information conveyed by the conditioned stimulus about the time of occurrence of the unconditioned stimulus. All of the groups in both experiments are represented in the figure. 1A = Experiment 1A; 1B = Experiment 1B.

Benedict, J. O., & Ayres, J. J. B. (1972). Factors affecting conditioning in the truly random control procedure in the rat. *Journal of Comparative and Physiological Psychology, 78,* 323–330.

Bialek, W., Nemenman, I., & Tishby, N. (2001). Predictability, complexity, and learning. *Neural Computation, 13,* 2409–2463.

Church, R. M., & Gibbon, J. G. (1982). Temporal generalization. *Journal of Experimental Psychology: Animal Behavior Processes, 8,* 165–186.

Cooper, L., Aronson, L., Balsam, P., & Gibbon, J. (1990). Duration of nonreinforcement periods affect response acquisition in nondifferential training. *Journal of Experimental Psychology: Animal Behavior Processes, 16,* 14–26.

Delamater, A. R. (1996). Effects of several extinction treatments upon the integrity of Pavlovian stimulus–outcome associations. *Animal Learning & Behavior, 24,* 437–449.

Denniston, J. C., Savastano, H. I., & Miller, R. R. (2001). The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 65–117). Mahwah, NJ: Erlbaum.

Drew, M. R., Zupan, B., Cooke, A., Couvillon, P. A., & Balsam, P. D. (2005). Temporal control of conditioned responding in goldfish. *Journal of Experimental Psychology: Animal Behavior Processes, 31,* 31–39.

Durlach, P. J. (1983). Effect of signaling intertrial unconditioned stimuli in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes, 9,* 374–389.

Egger, M. D., & Miller, N. E. (1962). Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology, 64,* 97–104.

Gallistel, C. R. (2003). Conditioning from an information processing perspective. *Behavioural Processes, 62,* 89–101.

Gallistel, C. R., Fairhurst, S., & Balsam, P. D. (2004). The learning curve: Implications of a quantitative analysis. *Proceedings of the National Academy of Sciences, USA, 101,* 13124–13131.

Gallistel, C. R., & Gibbon, J. G. (2000). Time, rate, and conditioning. *Psychological Review, 107,* 289–344.

Gallistel, C. R., King, A., & McDonald, R. (2004). Sources of variability and systematic error in mouse timing behavior. *Journal of Experimental Psychology: Animal Behavior Processes, 30,* 3–16.

Garner, W. R. (1962). *Uncertainty and structure as psychological concepts.* New York: Wiley.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84,* 279–325.

Gibbon, J. (1981). The contingency problem in autoshaping. In C. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 285–308). New York: Academic Press.

Gibbon, J., & Balsam, P. D. (1981). The spread of association in time. In C. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 219–253). New York: Academic Press.

Goddard, M. J., & Jenkins, H. M. (1987). Effect of signaling extra unconditioned stimuli on autoshaping. *Animal Learning & Behavior, 15,* 40–46.

Hallam, S. C., Grahame, N. J., & Miller, R. R. (1992). Exploring the edges of Pavlovian contingency space: An assessment of contingency theory and its various metrics. *Learning and Motivation, 23,* 225–249.

Jakubow, J. J., Brown, B. L., & Hemmes, N. S. (2004). The cover-stimulus effect: Role of similarity in durations of the CS and cover cues. *Psychological Record, 54,* 293–317.

Jenkins, H. M., Barnes, R. A., & Barrera, F. J. (1981). Why autoshaping depends on trial spacing. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 255–284). New York: Academic Press.

Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 276–296). New York: Appleton-Century-Crofts.

Killeen, P. R., & Weiss, N. A. (1987). Optimal timing and the Weber function. *Psychological Review, 94,* 455–468.

Kirkpatrick, K., & Church, R. M. (2000). Independent effects of stimulus and cycle duration in conditioning: The role of timing processes. *Animal Learning & Behavior, 28,* 373–388.

Kirkpatrick, K., & Church, R. M. (2004). Temporal learning in random control procedures. *Journal of Experimental Psychology: Animal Behavior Processes, 30,* 213–228.

Kremer, E. F. (1971). Truly random and traditional control procedures in CER conditioning in the rat. *Journal of Comparative and Physiological Psychology, 76,* 441–448.

Kremer, E. F., & Kamin, L. J. (1971). The truly random control procedure: Associative or nonassociative effects in rats. *Journal of Comparative and Physiological Psychology, 74,* 203–210.

LaBarbera, J. D., & Church, R. M. (1974). Magnitude of fear as a function of the expected time to an aversive event. *Animal Learning & Behavior, 2,* 199–202.

Papini, M. R., & Bitterman, M. E. (1990). The role of contingency in classical conditioning. *Psychological Review, 97,* 396–403.

Prokasy, W. F. (1972). Developments with the two-phase model applied to human eyelid conditioning. In A. F. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current theory and research* (pp. 119–147). New York: Appleton-Century-Crofts.

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology, 66,* 1–5.

Rescorla, R. A. (1972). Informational variables in Pavlovian conditioning. In G. H. Bower & J. T. Spence (Eds.), *Psychology of learning and motivation* (pp. 1–46). New York: Academic Press.

Rescorla, R. A. (2000). Associative changes with a random CS–US relationship. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology, 53*(B)*,* 325–340.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Rieke, F., Warland, D., de Ruyter van Stevenick, R., & Bialek, W. (1997). *Spikes: Exploring the neural code.* Cambridge, MA: MIT Press.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27,* 379–423, 623–656.

Williams, B. A. (1994). Contingency theory and the effect of the duration of signals for noncontingent reinforcement. *Psychonomic Bulletin & Review, 1,* 111–114.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology, 76*(Pt. 1), 177–186.

(Appendix follows)

Appendix

Temporal Information

The objective uncertainty about the time of occurrence of point events in an interval, T, is the entropy statistic, H. The condition that maximizes the uncertainty is a random rate protocol (see Figure 5A), in which the onset of the next unconditioned stimulus (US) is equally likely at every moment. This is the protocol we used in hopper training. The temporal properties of the random rate process are entirely characterized by the rate parameter, λ, which is the average number of events (onsets) per unit of time. The expected interval between onsets is $1/\lambda$.

The differences in uncertainty (i.e., entropy) between the intervals when the conditioned stimulus (CS) is not present and the intervals when it is quantify the objective information that the CS conveys about the timing of the US (Garner, 1962; Shannon, 1948). In deriving formulae for calculating these differences, we divide time into a sequence of very short moments, each lasting $\Delta\tau$. The duration of a $\Delta\tau$ does not matter (provided it is at least a couple of orders of magnitude smaller than the $1/\lambda$s, the expected intervals between events), because the $\Delta\tau$s disappear from the final equations. Uncertainty, and hence, differences in uncertainty (information) are measured in bits.

The uncertainty rate (bits per unit of time) in a random rate process like the feeder training (see Figure 5A) is as follows (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997, p. 116, Equation 3.22; for a derivation, see their Appendix 10):

$$\dot{H} = \lambda \log_2\left(\frac{e}{\lambda\Delta\tau}\right). \tag{1}$$

The average interval between the events is $1/\lambda$, so the average uncertainty in one interevent interval is

$$\overline{H} = \frac{1}{\lambda} \lambda \log_2\left(\frac{e}{\lambda\Delta\tau}\right) = k - \log_2 \lambda,$$

where $k = \log_2\left(\frac{e}{\Delta\tau}\right)$

The basal or background uncertainty about the time of the next US is $\overline{H}_C = k - \log_2 \lambda_C$, where C denotes context. If we consider only the effect of the CS on the rate of reinforcement, ignoring the additional reduction in uncertainty that arises from the fixed interval between CS onset and US delivery, then the average uncertainty about the time of the next US during a CS is $\overline{H}_T = k - \log_2 \lambda_T$, where T denotes trial. The information conveyed by the average CS is the difference in these two uncertainties:

$$\overline{H}_C - \overline{H}_T = (k - \log_2 \lambda_C) - (k - \log_2 \lambda_T) = \log_2 \lambda_T - \log_2 \lambda_C.$$

Put another way, when the relation between the CS and the US has been learned, the information communicated to the subject per reinforced CS trial in the protocols shown in Figures 5B and 6C is $\log_2\left(\frac{\lambda_T}{\lambda_C}\right)$. We will call the amount of information that a CS provides about a US the *informativeness* of the CS. (In other contexts, what we here call *informativeness* is called the *mutual information*.) The calculation of CS informativeness assumes that the relation between the CS and the US is known a priori. It does not take into account the important fact that at the beginning of conditioning, the subject does not know the CS–US relation (cf. Bialek, Nemenman, & Tishby, 2001).

What about the uncertainty reduction (information) that arises when the US occurs at a fixed latency after CS onset as in the current experiments (see Figures 5D and 5E)? Given the empirically well-established scalar variability in the subjective representation of temporal intervals (Church & Gibbon, 1982; Gibbon, 1977; Killeen & Weiss, 1987; LaBarbera & Church, 1974), we assume that a subject's probability distribution for the time at which the US will occur after CS onset is a Gaussian distribution with $\sigma = wT$, where w is the Weber fraction for timing, and T is the fixed CS–US latency (the trial duration). The uncertainty (entropy) in a Gaussian distribution is as follows (Rieke et al., 1997, p. 109, Equation 3.11):

$$H = \frac{1}{2} \log_2\left[2\pi e \, \frac{\sigma^2}{(\Delta\tau)^2} \right].$$

By substitution:

$$\begin{aligned}
H &= \frac{1}{2} \log_2\left[2\pi e \, \frac{(wT)^2}{(\Delta\tau)^2} \right] \\
&= \frac{1}{2} \log_2\left(\frac{1}{\Delta\tau}\right)^2 + \frac{1}{2} \log_2 w^2 + \frac{1}{2} \log_2 T^2 + \frac{1}{2} \log_2 2\pi e \\
&= \log_2\left(\frac{1}{\Delta\tau}\right) + \log_2 T + \log_2 w + \frac{1}{2} \log_2 2\pi e \\
&= \frac{1}{2} \log_2 2\pi e + \log_2 w - \log_2 \Delta\tau - \log_2 \lambda.
\end{aligned}$$

The difference between this uncertainty and the uncertainty of a random rate process with the same expectation, which is

$$\log_2 e - \log_2(\Delta\tau) - \log_2 \lambda,$$

is

$$\begin{aligned}
&\log_2 e - \log_2(\Delta\tau) - \log_2 \lambda - \left(\frac{1}{2} \log_2 2\pi e + \log_2 w - \log_2 \Delta\tau - \log_2 \lambda\right) \\
&= \log_2 e - \log_2 w - \frac{1}{2} \log_2 2\pi - \frac{1}{2} \log_2 e = \frac{1}{2} \log_2\left(\frac{e}{2\pi}\right) - \log_2 w. \tag{2}
\end{aligned}$$

In Equation 2, we see that the difference in uncertainty between the case in which the US occurs at random within the CS and the case in which it occurs at a fixed latency after CS onset is a linear function of the log of the Weber fraction. We take as our estimate of w, the coefficient of variation in the distribution of the stop times in the peak procedure, which is about 0.16 in different species of subject, across a wide range of target intervals (Gallistel, King, & McDonald, 2004). For this value of w, the calculation in Equation 2 gives roughly 2 bits. Thus, fixing the CS–US interval increases by 2 bits the information about US timing conveyed by the onset of the CS.

In sum, we can separate the information conveyed simply from the fact that the rate of US occurrence is higher in the presence of the CS from the information conveyed by virtue of the fixed interval between CS onset and US delivery. The first component depends on the rate ratios. The second component is a fixed 2 bits per reinforced CS. Our hypothesis is that bits to acquisition is approximately a constant. Bits to acquisition is CS informativeness multiplied by the number N_a, the number of reinforced CSs before the onset of responding. Thus, when the US comes at a fixed interval after CS onset, bits to acquisition is

$$N_a\left[\log_2\left(\frac{\lambda_T}{\lambda_C}\right) + 2 \right] = N_a \log_2\left(\frac{\lambda_T}{\lambda_C}\right) + 2N_a. \tag{3}$$

When the CS–US interval is not fixed, one drops the $2N_a$ term.