Published in *Journal of Experimental Psychology: Animal Behavior Processes, vol 37 (4), pp 375-384.*
http://www.apa.org/pubs/journals/xan/index.aspx

Response Rate and Reinforcement Rate in Pavlovian Conditioning

Justin A. Harris & Joanne S. Carpenter

The University of Sydney

Abstract

Four experiments used delay conditioning of magazine approach in rats to investigate the relationship between the rate of responding, R, to a conditioned stimulus (CS) and the rate, r, at which the CS is reinforced with the unconditioned stimulus (US). Rats were concurrently trained with four variable-duration CSs with different rs, either as a result of differences in the mean CS-US interval or in the proportion of CS presentations that ended with the US. In each case, R was systematically related to r, and the relationship was very accurately characterized by a hyperbolic function, $R = Ar/(r+c)$. Accordingly, the reciprocal of these two variables – response interval, I (=1/R), and CS-US interval, i (=1/r) – were related by a simple affine (straight line) transformation, $I = mi+b$. This latter relationship shows that each increment in the time that the rats had to wait for food produced a linear increment in the time they waited between magazine entries. We discuss the close agreement between our findings and the Matching Law (Herrnstein, 1970), and consider their implications for both associative theories (e.g., Rescorla & Wagner, 1972) and non-associative theories (Gallistel & Gibbon, 2000) of conditioning.

Address correspondence to:

Justin Harris
School of Psychology
University of Sydney
Sydney 2006
Australia
Tel: +61 2 9351 2864
Email: justin.harris@sydney.edu.au

Key words: Rat; delay conditioning; magazine approach; Matching Law; variable interval.

Author note: The authors thank Fred Westbrook, Evan Livesey, and Benjamin Andrew for comments on an earlier draft of this manuscript. This work was supported by grants DP0771154 and DP1092695 from the Australian Research Council.

The most significant theoretical advances in our understanding of Pavlovian conditioning have come from computational descriptions of the way that animals learn about the temporal correlation between a conditioned stimulus (CS) and unconditioned stimulus (US). A prevalent construct in this field is associative strength (V), a computationally-defined variable that represents the amount an animal has learned about the relationship between a CS and US (Rescorla & Wagner, 1972). One of the simplest yet most successful accounts describes this learning as the output of an error-correction algorithm that continuously updates the organism's knowledge of the CS-US association by bringing it closer to the most recent experience of those events (Bush & Mosteller, 1951; Rescorla & Wagner, 1972). On any trial when the CS and US are paired, V increases by a fixed proportion (k) of the difference between the asymptote (λ) to which the CS-US association converges and the associative strength that has already been learned (V_{t-1}). Thus, learning (ΔV) on trial t is defined as

$$\Delta V_t = k(\lambda - V_{t-1}) \qquad \text{Equation 1}$$

Moreover, when the CS is presented in the absence of the US, V decreases (the CS-US association extinguishes) by an amount proportional to V_{t-1}. This is equivalent to the rule stated in Equation 1, but where $\lambda = 0$ because the US is absent. Rescorla and Wagner (1972) greatly extended the explanatory scope of this simple description by proposing that learning on any trial is proportional to the difference between λ and the sum of what has been learned about all CSs present on the current trial (ΣV). This simple error-correction rule, or an operation that is functionally equivalent, has been incorporated into numerous more elaborate models of learning (Desmond & Moore, 1988; Grossberg & Schmajuk, 1989; Harris, 2006; McLaren & Mackintosh, 2000; Pearce, 1994; Sutton & Barto, 1981; Wagner, 1981; Wagner & Brandon, 2001).

This description of the processes underlying conditioning has received enormous empirical support, and continues to be one of the most influential theories in experimental psychology[1]. However, one limitation of the Rescorla-Wagner model, and models like it, stems from the fact that they fail to specify in detail how V is expressed in conditioned responding. In the absence of such detail, most researchers take the conservative position of assuming that there is a monotonically non-decreasing relationship between V and whatever variable is used to index conditioned responding, while remaining agnostic as to the precise form of that relationship. As a result, behavioral data are located on an ordinal scale of measurement, where a difference in responding is interpreted as a difference in V (Rescorla, 2000), but the reverse inference cannot always be made, nor can the size of a difference in responding be used to infer the size of a difference in V. In this regard, associative models lag behind so-called "timing" models that make richer theoretical contact with behavioral data, particularly in explaining patterns of responding within the trial (e.g., Church & Broadbent, 1990; Gallistel & Gibbon, 2000; Gibbon, 1977; Killeen & Fetterman, 1988; Kirkpatrick & Church, 2003).

The goal of the current research is to take an initial step to providing more detailed

---

[1] In a recent search, the original paper by Rescorla and Wagner (1972) had received over 3,000 citations, and its annual citation rate has risen consistently since its publication, reaching more than 150 citations in each of the last 4 years.

information about the relationship between response strength and conditioning strength, and in particular how responding might be related to V. A direct approach to uncovering this relationship would require independent measurement of both variables in order to map out changes in one as a function of changes in the other. This approach, however, is not available because V, as a purely theoretical construct with no physical instantiation, cannot be measured directly. Nonetheless, one can control how V changes by systematic manipulation of one of its input variables while concurrently measuring response strength. For example, from Equation 1, V will vary as a function of (1) the number of conditioning trials, (2) the learning rate parameter k, and (3) the presence versus absence of the US on each trial (the value taken by λ). This means that we can specify how V will change as a function of one of these variables. If, at the same time, we can characterize how response strength changes as a function of this variable, it should be possible to deduce the relationship between response strength and V.

To obtain reliable data on the relationship between responding and a given independent variable, it will be necessary to record responses at steady state while the variable is held constant over many trials. This requirement means we must select a variable that affects asymptotic response strength. The key variable that determines asymptotic strength in Equation 1 is λ. It is feasible to manipulate the US by varying, for example, its number or intensity, and this should vary λ. However, this will only achieve our goal if we can specify *a priori* exactly how this physical manipulation of the US will affect the value of λ. The only relationship we can specify precisely is that λ will take its nominal value (e.g., 1) when the US is present and equals 0 when the US is absent. Therefore, the most viable approach is to manipulate the rate, r, at which the CS is reinforced with the US, thereby manipulating the proportion of trials when λ takes a positive value versus a value of zero. As we will show shortly, the relationship between V and r can be specified precisely. Moreover, this between-trials description of r can be easily adapted to operate at the within-trial level, using Equation 1 to update V continuously in "real-time". In this case, r can be operationally defined as the proportion of time during the CS when the US is present versus absent.

The relationship between associative strength, V, and reinforcement rate, r

To uncover the relationship between V and response strength, we must first define the relationship between V and r as derived from Equation 1, and compare this to the empirically identified relationship between r and response strength. Accordingly, in this section we derive how V is related to r, and to its reciprocal, i, the CS-US interval. The full description of the derivation is presented in Appendix A. Its starting point is the assumption that the steady-state associative strength is reached after many conditioning trials once each increase in V during reinforcement is matched by the intervening decrease in V when the CS is present without reinforcement. This relation can be expressed as

$$k_0 \, i \, \overline{V} = k_1(\lambda \quad - V_i) \qquad \text{Equation 2}$$

The left side of the equation describes the total amount of extinction that accrues during the CS-US interval, i, as a function of the CS's mean associative strength across the interval, $\overline{V}$, and the learning rate parameter for non-reinforcement, k_0. The right side of the equation describes the increase in V when the US is presented at the i^{th} moment of the CS (i.e., the end of the CS for delay conditioning as modeled here), where k_1 is the learning rate parameter for reinforcement. If the rate of extinction during the CS is low (k_0 is small), the difference between V_i and $\overline{V}$ will be small. Therefore, if we treat V_i as equal to $\overline{V}$ (see Appendix A for further analyses without this assumption), Equation 2 can be re-arranged as

$$\frac{1}{\overline{V}} = \frac{k_0}{k_1\lambda} i + \frac{1}{\lambda} \qquad \text{Equation 3}$$

The above equation can be recognized as the formula for a straight line. That is, the inverse of the associative strength at steady state is related to the CS-US interval, i, by a straight line. Further, if we substitute r for $1/i$, we can obtain Equation 4:

$$\overline{V} = \frac{\lambda\, r}{r + \frac{k_0}{k_1}} \qquad \text{Equation 4}$$

In other words, the mean associative strength bears a hyperbolic relationship to r.

In light of the foregoing analysis, the present experiments were designed to map out the relationship between r and response rate, R, using delay conditioning of magazine approach in rats. We used this paradigm because, as we have recently confirmed, rats' conditioned magazine responses are affected by the rate at which CSs are reinforced with food (Harris, in press; Harris, Gharaei, & Pincham, in press). However, like many conditioned responses, the rate of magazine entry also tracks the timing of the US if there is a consistent temporal relationship between the CS and US. That is, if the US occurs at a fixed interval after CS onset, the response to the CS increases steadily up to the anticipated time of the US (Davis, Schlesigner, & Sorenson, 1989; Harris et al., in press; Kehoe & Joscelyne, 2005; Pavlov, 1927; Roberts, 1981; Smith, 1968; Williams, Lawson, Cook, Mather, & Johns, 2008). This change in responding across the CS would confound our efforts to establish the relationship between R and r. To avoid this, we have trained rats with variable CS-US intervals. We have recently shown that, when rats are trained with uniformly distributed random intervals, they maintain a stable (uniform) level of responding across the duration of the CS (Harris et al., in press). (By contrast, their response rate falls, approximating an exponential function, when they are trained with exponentially distributed random CS-US intervals.) The level at which rats respond to such variable CS-US intervals is affected by the mean interval (i.e., the average rate at which the CS is reinforced, Harris et al., in press). Thus, the present experiments trained rats with four variable CSs, each with a different mean CS-US interval, in order to determine how the strength of responding, R, varies as a function of r. The raw and summarized data from all experiments presented in this paper can be downloaded from the website:

http://sydney.edu.au/science/psychology/staff/justinh/downloads/

Experiments 1, 2, and 3

The first three experiments trained rats concurrently with four CSs that had different rs by virtue of differences in their mean CS-US intervals, as summarized in Table 1. The CS-US interval for a given CS varied randomly from trial-to-trial within a uniform distribution centered on the mean. Rats were trained until

Published in *Journal of Experimental Psychology: Animal Behavior Processes, vol 37 (4), pp 375-384.*
http://www.apa.org/pubs/journals/xan/index.aspx

Table 1. CS-US and inter-stimulus intervals used in Experiments 1 to 3.

	Mean CS-US interval (range)				**Mean inter-stimulus interval (range)**
	CS 1	CS 2	CS 3	CS 4	
Experiment 1	10 s (2 – 18 s)	24 s (2 –46 s)	60 s (2 – 118 s)	150 s (2 – 298 s)	120 s (30 – 210 s)
Experiment 2	10 s (2 – 18 s)	30 s (2 – 58 s)	50 s (2 – 98 s)	70 s (2 – 138 s)	90 s (30 – 150 s)
Experiment 3	10 s (2 – 18 s)	40 s (2 – 78 s)	70 s (2 – 138 s)	100 s (2 – 198 s)	65 s (30 – 100 s)

the mean response rate, R, to all CSs had reached a plateau. Training was then extended for a further 10 days, and all analyses were performed on steady-state Rs over the last 10 days of training.

Methods

Subjects

Experimentally naive male Hooded Wistar rats (Rattus norvegicus) were obtained from the Laboratory Animal Services breeding unit at The University of Adelaide, South Australia. The rats were 8 to 10 weeks of age at the start of Experiment 1 (n=16) and Experiment 2 (n=15), and were approximately 16 weeks of age at the start of Experiment 3 (n=16). During the experiment, they were housed in groups of eight in large white plastic tubs, measuring 26 x 59 x 37cm (height x length x depth), located in the animal colony maintained by the School of Psychology at the University of Sydney. They had unrestricted access to water in the home tubs. Three days prior to commencement of the experiment, their access to food was restricted to 2 hr per day (to commence half an hour after the end of the daily training sessions).

Apparatus

Rats were trained and tested in 16 Med Associates™ conditioning chambers measuring 28.5 x 30 x 25 cm (height x length x depth). The end walls of each chamber were made of aluminum; the sidewalls and ceiling were Plexiglas™. The floor of the chamber consisted of stainless-steel rods, 0.5 cm in diameter, spaced 1.5 cm apart. Each chamber had a recessed food magazine in the center of one end wall, with an infra-red LED and sensor located just inside the magazine to record entries by the rat. A small metal cup measuring 3.5 cm in diameter and 0.5 cm deep was fixed on the floor of each food magazine. Attached to the food magazine was a dispenser delivering 45 mg food pellets (Noyes Formula P; Research Diets Inc, New Brunswick, NJ). Each chamber was enclosed in a sound- and light-resistant wooden shell. Throughout all sessions, fans located in the rear wall provided ventilation;

the operation of these created a background noise level measuring 70dB. Experimental events were controlled and recorded automatically by computers and relays located in the same room.

Two auditory and two visual stimuli were presented from four spatially separated sources. White noise (78dB) was presented from a speaker mounted on the wall of each operant chamber above and to the right of the food magazine. A tone (78dB and 2.9 kHz) was produced from a piezo buzzer positioned on the floor of the sound-attenuating shell behind each operant chamber. A flashing light (2 Hz; 3.0cd/m^2) was emitted by a 3x3 array of white LEDs, located on the floor of the sound-attenuating shell in front of the operant chamber. A steady light (30cd/m^2) was produced by an incandescent bulb mounted high on the back wall of the sound-attenuating shell. The allocation of these four stimuli to the four CSs was counterbalanced across rats.

Procedure

On the day before training began, the rats received a single 20-min magazine training session during which 20 food pellets were presented on a variable-time (VT) 1 min schedule, with no stimulus presentations. The rats then received daily conditioning sessions 5 days per week for a total of 36 days (Experiment 1), 32 days (Experiment 1), or 35 days (Experiment 3). Each session consisted of 10 (Experiment 1) or 12 (Experiments 2 and 3) intermixed presentations of each of the four CSs, for a total of 40 (Experiment 1) or 48 (Experiments 2 and 3) presentations. The order of presentations was randomized within each half of the session. In each experiment, the inter-stimulus interval varied randomly according to a uniform distribution with a minimum of 30 s. The duration of each CS varied from trial-to-trial, but regardless of length, the termination of every CS coincided with delivery of a food pellet. The mean and range of CS-US intervals and inter-stimulus intervals are summarized in Table 1. Across all days, the timing of each photo-beam interruption by head entry into the magazine was recorded during each CS and 30-s pre-CS period.

The data from Days 28 and 29 were excluded for one rat (#3) in Experiment 3 due to a fault with the photo beam that measured nose pokes in one chamber. After the end of Experiment 1, we discovered that the white noise speaker was not working in the conditioning chamber of Rat 16. The response data suggested that the rat did not receive the noise stimulus during training, insofar as its response to the noise was equal to the response in the pre-CS interval, even though the noise served as the shortest (10-s) CS for this rat. Moreover, this rat’s pre-CS response rate was higher than that of any other rat, consistent with that rat receiving unsignaled food pellets on 25% of trials. Therefore, we excluded that rat’s data from all analyses, bringing n for this experiment to 15.

Results & Discussion

The mean Rs on each day of each experiment are plotted down the left side of Figure 1. Responding increased steadily for the first 20 to 25 days of training, and appeared to reach a plateau thereafter. The stable differences in R to the different CSs were confirmed by repeated-measures ANOVAs (with Greenhouse-Geisser correction) conducted on the data from the last 10 days of each experiment. There were significant main effects of CS in each experiment [$F(2.4,34.3) = 56.07$ in Experiment 1, $F(1.57,22.0) = 42.68$ in Experiment 2, $F(1.27,19.1) = 48.87$ in Experiment 3, all ps $< .001$], but not of Day [$F(3.6,50.4) = 1.70$, $F(3.33,46.6) = 2.57$, $F(2.08,31.2) = 1.60$, smallest $p = .06$], nor were there any significant linear trends across Days [largest $F(1,14) = 1.57$, smallest $p = .23$].

Rs during each CS were analyzed for the final 10 days of training. Each plot down the right side of Figure 1 shows the mean R during each second as time elapsed within the CS in each experiment. These plots show that the mean R was relatively stable across time during the CS but differed as a function of the mean CS duration (Harris et al., in press). ANOVAs conducted on these data confirmed that there were significant differences in responding among the four CSs, [$F(3,42) = 47.56$ in Experiment 1, $F(3,42) = 19.04$ in Experiment 2, $F(3,42) = 49.24$ in Experiment 3, ps $< .001$]. For the data from Experiment 1, paired t-tests identified significant differences in R between the 10-s and 24-s CSs, $t(14) = 2.61$, $p = .02$, between the 24-s and 60-s CSs, $t(14) = 6.99$, $p < .01$, and between the 60-s and 150-s CSs, $t(14) = 2.92$, $p = .01$. Similar comparisons identified differences between the 10-s and 30-s CSs, $t(14) = 2.91$, $p = .01$, the 30-s and 50-s CSs, $t(14) = 4.79$, $p < .001$, and the 50-s and 70-s CSs, $t(14) = 2.63$, $p = .02$ in Experiment 2; and between the 10-s and 40-s CSs, $t(15) = 7.90$, $p < .001$, between the 40-s and 70-s CSs, $t(15) = 6.71$, $p < .001$, and between the 70-s and 100-s CSs, $t(15) = 3.71$, $p = .002$ in Experiment 3.

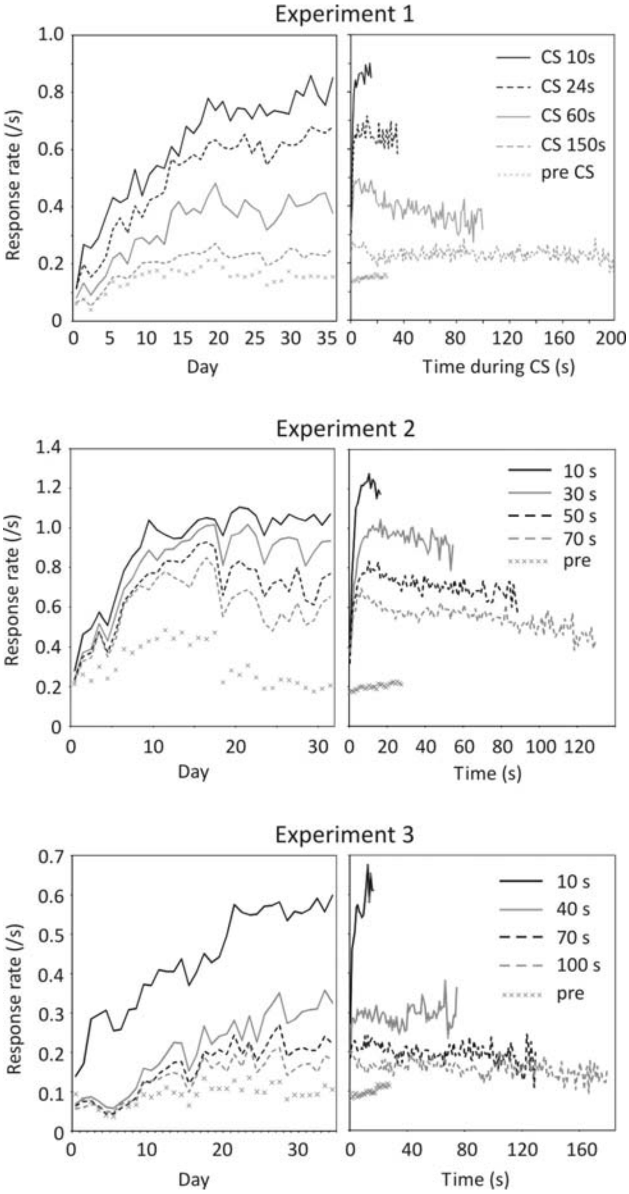

Figure 1. **Left**: Mean response rates during each CS, and pre-CS interval, across training days in Experiments 1 to 3. **Right**: Mean response rates as a function of elapsed time within the trial averaged over the last 10 days of each experiment.

The key analysis compares the mean R for each CS against its r. This relationship is shown in the plots down the left side of Figure 2. In each experiment, there was a very orderly monotonic relationship between R and r, that isvery well described by a hyperbolic function. This is equivalent to the relationship between V and r described in Equation 4. The observed relationship between reinforcement and responding is also described by a straight line when plotting the reciprocal of R (i.e., the response interval, I) against the reciprocal of r (the CS-US interval, i), as shown in each plot on the right side of Figure 2. Put informally, this relationship shows that the average time the rats wait between consecutive responses is proportional to the mean expected waiting time to reinforcement.

Comparing hyperbolic, log and power functions.

We conducted additional analyses across all three experiments that compared the hyperbolic function with two other two-parameter functions, both of which are monotonically non-decreasing but, unlike the hyperbola, are not asymptotic. The first of these is the simple power function: $R = Ar^p$, where p is < 1. The second is the log function: $R = Alog_e(cr+1)$. We have used this function, rather than $R = Alog_e(r)+c$, because the latter function is not resolvable when r = 0, and indeed becomes increasingly negative as $r \rightarrow 0$, whereas we would logically expect the relationship between R and r to pass through the origin (i.e., R = 0 when r = 0). In effect, requiring that all functions tested here pass through the origin adds a fifth data point to constrain each model.

Both log and power functions provide good fits to the data from each of the three experiments. The R^2s for these functions (and, for comparison, the hyperbola) are as follows: $R^2_{log} = 0.995$ and $R^2_{power} = 0.948$ ($R^2_h = 0.997$) in Experiment 1; $R^2_{log} = 0.898$ and $R^2_{power} = 0.853$ ($R^2_h = 0.954$) in Experiment 2; $R^2_{log} = 0.996$ and $R^2_{power} = 0.987$ ($R^2_h = 0.999$) in Experiment 3. To compare the functions statistically, we fitted all three functions to the response rate data of each individual rat across all three experiments (n = 46). Most fits were good, with the median

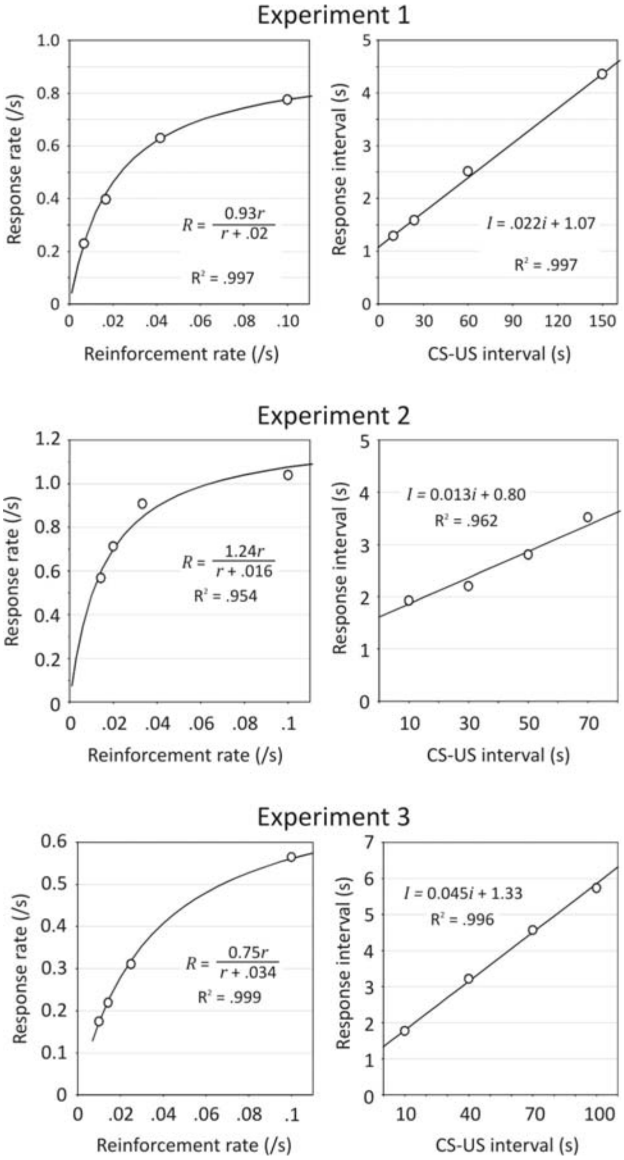

Figure 2. Data averaged over the last 10 days of Experiments 1 to 3. **Left**: The four circles mark the mean response rates (R) to each CS as a function of its reinforcement rate (r). The curve passing through these data points plots the best-fitting hyperbola described by the formula shown, along with the R^2 for its fit to the data. **Right**: the four circles mark the response interval, I (= $1/R$), for each CS as a function of its mean CS-US interval, i (= $1/r$). The best-fitting straight line through the data is shown, along with the formula of this line and R^2 for its fit.

R^2s equal to 0.93 (range .449 to .999) for the hyperbola, 0.93 (.298 to .999) for the log function, and 0.91 (.451 to .999) for the power function. (In reporting these ranges of R^2s, we have excluded two rats from Experiment 2 that responded more to the least reinforced CS than to the most reinforced CS, and thus for which the R^2s were negative for all three monotonically increasing functions.) The mean R^2s were 0.85, 0.83, and 0.82 for the hyperbolic, log, and power functions, respectively. Student *t*-tests revealed that the R^2 values for the hyperbolic function were significantly higher than for the log function, $t(45) = 2.12$, $p = .036$, and the power function, $t(45) = 2.68$, $p = .01$. The log and power functions were not significantly different from each other, $t < 1$. Thus, while the relationship between R and r was captured by a log function and by a power function, each of these functions were slightly inferior to the hyperbolic function in accounting for the variance in the data. However, further analyses showed that the hyperbolic function is slightly more flexible than the log function in explaining data variance, inasmuch as the hyperbolic function performed as well as the log when fitting data generated from a log function with added noise, whereas the log function performed worse than the hyperbolic function when fitting data generated from a hyperbola with added noise (Pitt, Myung, & Zhang, 2002). This difference in flexibility puts the hyperbolic function at greater risk of over-fitting data, and thus we need to exercise caution when concluding that the hyperbolic function is superior to the log function in explaining the present data. (A similar test showed that the hyperbolic and power functions were equally flexible, in that each was better than the other in accounting for noisy data from its own function.)

Experiment 4

Experiments 1 to 3 revealed a very clear relationship between R and r when the latter was manipulated by varying the CS-US interval. Experiment 4 was designed to investigate this relationship by varying the proportion of trials on which each CS was reinforced. Rats were trained with four variable-duration CSs each of which had the same mean duration of 30 s. One CS was reinforced on 100% of its trials, a second was reinforced on 50% of trials, a third was reinforced on 25% of trials, and the fourth was reinforced on approximately 8% of trials (one trial in 12). This design extends the previous experiment by providing an alternative means to manipulate r. Its principal advantage is that Rs to all four CSs are recorded over the same time window, whereas this time window varied between the four CSs in the previous three experiments.

Methods

Subjects and apparatus

Sixteen experimentally naive male Hooded Wistar rats (8 to 10 weeks of age at the start of the experiment) were obtained from the same source and maintained in the same manner as in Experiments 1-3. They were trained and tested in the same conditioning chambers, and with the same counterbalanced stimuli, used in those previous experiments.

Procedure

On the day before training began, the rats received a single 20-min magazine training session as in Experiments 1-3. They then received daily conditioning sessions for a total of 40 days. Each session consisted of 12 intermixed presentations of each of the four CSs, for a total of 48 presentations. The order of presentations was randomized within each quarter of the session (i.e., each block of 12 trials contained three presentations of each of the four CSs). The average inter-stimulus-interval was 120 s as in Experiment 1. The duration of each CS varied from trial-to-trial between 2 s and 58 s. For one CS, the termination of every trial coincided with delivery of a food pellet. For the other three CSs, the end of the CS coincided with food on one in two trials (50%), one in four trials (25%), and one in 12 trials (8%). (This last CS was reinforced on every sixth trial for the first 12 days of training, but this was reduced to one in 12 from Day 13 in order to increase the difference in R to this CS compared to the 25% CS.) Across all 40 days, photo-beam interruptions by head entry into the magazine were recorded during each CS. Each session lasted approximately 2 hr. All analyses were conducted on Rs from the final 10 days of training.

Results and Discussion

The top left plot of Figure 3 shows the mean Rs for each CS, and the pre-CS period, on each day of the experiment. Responding to the four CSs increased steadily over the first 25 days of training, but changed little thereafter. A repeated measures ANOVA (with Greenhouse-Geisser correction) on the data from Day 26 onwards showed that there was a significant main effect of CS, $F(2.1,28.9) = 65.30$, $p < .001$, but no significant main effect of Day, $F(4.2,58.7) = 2.06$, $p = .094$, nor was there a significant linear trend across days, $F < 1$.

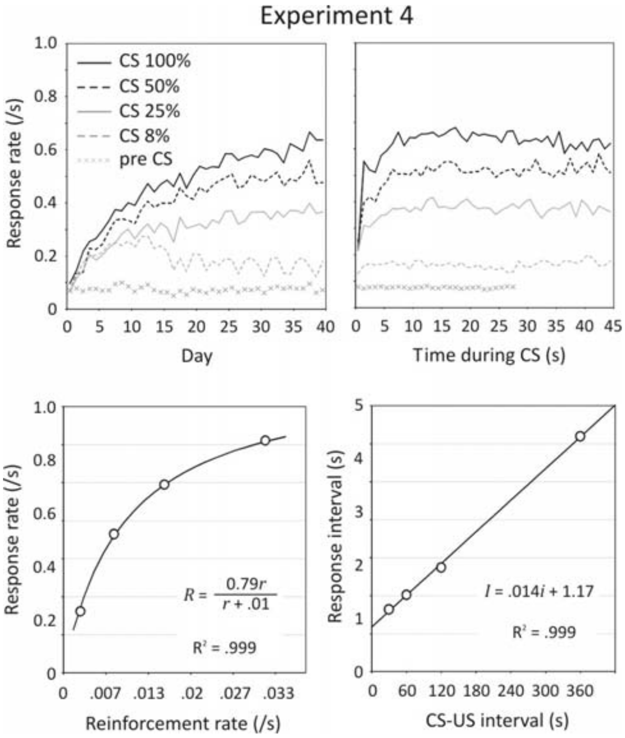

Figure 3. Results of Experiment 4. **Top**: Mean response rates during each CS, and pre-CS interval, across days (left) and across elapsed time within the trial averaged over the last 10 days of the experiment (right). **Bottom**: Both panels plot the same data from the last 10 days. On the left, the four circles mark the mean response rates (R) to each CS as a function of its reinforcement rate (r), and the continuous line plots the best-fitting hyperbolic function. The formula of this function is shown along with the proportion of explained variance (R^2) for its fit to the data. On the right, the four circles mark the response interval, I (= $1/R$), for each CS as a function of its mean CS-US interval, i (= $1/r$). The best-fitting straight line through the data is shown, along with the formula of this line and R^2 for its fit.

The top right plot in Figure 3 shows the mean R to each CS as a function of time elapsed during the CS, from the last 10 days of the experiment. It is clear that R varied as a function of r, and that responding remained stable over time within the trial. Because each CS had the same duration in this experiment (unlike in Experiments 1 to 3), it was possible to conduct a repeated measures ANOVA that compared second-by-second Rs between CS. This analysis was confined to a time window starting from the 6th second of the trial, by which time responding had stabilized after CS onset, to the 45th second. (Sampling of responding beyond 45 s was too infrequent to provide reliable data per second.) The analysis identified a significant overall difference between the four CSs, $F(3,45) = 52.32$, $p < .001$, but found there to be no significant main effect of time during the trial, $F(39,585) = 1.08$, $p = .349$, nor was there an interaction between CS and time, $F < 1$. Therefore, after the initial rise in responding during the first 5-s of the CS, R was constant across the remainder of the trial, and was equally constant for all four CSs. The differences in level of responding between the CSs was confirmed by paired t-tests that identified significant differences in responding between the 100% and 50% CSs, $t(15) = 5.32$, $p < .001$, between the 50% and 25% CSs, $t(15) = 3.91$, $p = .001$, and between the 25% and 8% CSs, $t(15) = 5.03$, $p < .001$.

The key analysis involves plotting the observed R for each CS against its r, as shown in the lower left part of Figure 3. Because the length of the CS-US interval affects R, as demonstrated in Experiments 1-3, we have converted the proportional r for each CS into the absolute r per second by dividing the proportional rate by 30 s. Thus, for example, the 50% CS has a per-second r of 1/60. This allows us to plot r in this experiment on the same scale as in Experiments 1-3. It should be noted that this simple transformation has no substantive impact on the relationship between r and R. As shown in Figure 3, that relationship is almost perfectly described by a hyperbola. In the lower right part of the same figure, we have also plotted the relationship between the response interval I (the reciprocal of R) and the cumulative CS-US interval i (the reciprocal of r). This plot shows that the function relating i to I lies almost exactly along a straight line. Thus Experiment 4 has confirmed the key observations of Experiments 1 to 3.

General Discussion

Here we have presented four experiments investigating delay conditioning of magazine approach in rats. In each experiment, rats were trained with four variable CSs with different reinforcement rates, r, providing the opportunity to examine how the rate of responding, R, varied as a function of r. These variables were related by the simple hyperbolic function: $R = Ar/(r+c)$, where A is the asymptote of responding and c determines how quickly R approaches the asymptote. This corresponds to an affine (straight line) relationship, $I = mi + b$, between the CS-US interval, i (=$1/r$), and the response interval, I (=$1/R$), where the intercept, b, equals $1/A$, and the slope, m, equals c/A. In effect, this relationship means that each increase in the average time that the rat must wait for food produces a linearly proportional increase in the average time it waits between magazine entries. The relationship between R and r (and between I and i) held true when r was manipulated either by varying the mean CS-US interval (Experiments 1-3) or by varying the proportion of CS trials that were reinforced (Experiment 4). It is worth noting that the observed relationship between R and r could also be accounted for very successfully by a log function, $R = A\log_e(cr+1)$, and by a power function, $R = Ar^p$. However, these functions provided a slightly weaker fit to the mean data in most experiments, and were worse than the hyperbolic function when fitted to the data from each individual rat across Experiments 1 to 3. Nonetheless, with only four data points per experiment, the results presented here cannot be taken as definitive evidence that the relationship between R and r is hyperbolic rather than some similar curvilinear function.

In the Introduction we described how a hyperbolic relationship between V and r (Equation 4), and an affine (straight line) relation between the inverse of V and the CS-US interval (Equation 3), can be derived from the error-correction learning rule in Equation 1. For present purposes, the most important outcome of this analysis is that these relationships are isomorphic with the relationship between R and r (or between I and i) observed in each of the four experiments presented here. This greatly narrows the range of possible relationships that might exist between V and R. The most obvious relationship suggested by the present results is a direct linear transformation: $R = aV$, where a scales the units of V to units of R. If true, this very simple linear relationship would mean that R can be treated as a ratio scale of measurement for V. However, this conveniently simple and promising conclusion is not mandated by the current evidence. It is also possible that R and V are themselves related by another hyperbolic function because the composition of two hyperbolas is another hyperbola. Put differently, because any hyperbolic function can be decomposed into two other hyperbolic functions, the observed hyperbolic relationship between R and r can itself be composed of a hyperbolic relation between r and V and a second hyperbolic relation between V and R. Naturally, the same conclusions apply to the affine relationship between I and i. It is important to note, however, that these two possibilities are exhaustive, in that no other transformation between V and R will give the observed hyperbolic relationship between r and R. Therefore, on the basis of the present data, we can draw the general conclusion that $I = mV^{-1} + b$. If $b = 0$, the relation between R and V is linear, and we can treat responding as a ratio scale of measurement of V. If $b > 0$, the straight-line relating I and V^{-1} means that we can treat I as an interval scale of measurement of V^{-1}. Either way, the present findings reveal how we can use conditioned responding in this paradigm as a more informative index of associative strength than the ordinal scale normally advocated by experimenters in this field.

If the function relating what is learned (e.g., V) to responding is itself hyperbolic, what might be the basis for such a relation? One compelling account appeals to The Matching Law (Herrnstein, 1970). Herrnstein and others observed that, when animals are presented with a choice to make either of two trained instrumental responses to obtain reward on a variable-interval (VI) schedule, the relative frequency with which they performed each response was strictly proportional to the relative frequency with which each response produced reinforcement (Davison & McCarthy, 1988). From this, Herrnstein (1970) derived the general equation that, for any single instrumental response, B_1, trained on a VI schedule, $B_1 = kR_1 / (R_1 + R_e)$, where R_1 is the rate at which B_1 is reinforced, R_e is the summed rate of reinforcement of all other competing behaviors, and k is a scaling constant. Herrnstein's equation is clearly the same as the hyperbolic relationship we have observed between r and R.

The correspondence between our data and Herrnstein's (1970) equations is unexpected because Herrnstein's laws were developed to explain an animal's choice to perform an instrumental response, whereas the present experiments were designed with Pavlovian contingencies between the CS and US. This raises the question whether the magazine approach response measured here is Pavlovian or instrumental. On the one hand, magazine activity could be considered an instrumental response to the extent that the collection of each food pellet is contingent on magazine entry. On the other hand, magazine entries made during the CS are usually considered to be Pavlovian responses because there is no contingency between any of these responses and the availability of the food pellet. Indeed, Holland (1979) reported that rats acquired magazine entry responses to a tone CS even though they were trained on an omission schedule in which the tone was followed by food except when the rat entered the magazine during the CS. This suggests that magazine entry can arise as a Pavlovian response rather than an instrumental one. Nonetheless, even if magazine entry were an instrumental response, one would expect the strength of that response to be systematically modulated by differences in V of each CS (Delamater, 1996; Holmes, Marchand, & Coutureau, 2010). Since V will be related to r by a hyperbolic function, as derived from Equation 1, the strength of the response would be expected to show the same hyperbolic relationship to r that we report here. Indeed, this very fact raises the question whether Pavlovian conditioning processes might contribute to the evidence of matching even in studies of instrumental conditioning.

Whether the behaviors in question are mediated by Pavlovian or instrumental processes, the hyperbolic relationship between response rate and reinforcement can be understood as a behavioral-normalization process – the impact of reinforcement on a specific behavior is not determined in isolation, but is computed relative to the total reinforcement of all behaviors. For example, in our experimental paradigm, the rat is continuously faced with competing motivations to either check the magazine for food or engage in any other behavior in its repertoire (e.g., exploration, grooming, sleeping). The frequency with which the rat checks the magazine would depend not only on how frequently the CS or response is reinforced, but also on the extent that other competing behaviors are reinforced. One implication of this particular description is that our discovery of the relationship between r and R has been crucially dependent on the use of a within-subject design, because this holds constant the amount that other competing behaviors are reinforced. That is, in our experiments, responding to each CS is normalized by the same level of reinforcement of other behaviors. This may not be true in a between-subjects design. On the other hand, in our within-subjects designs it is also possible that the response rate to each CS is affected by the conditioning strength of the other CSs, such as might arise from behavioral induction or contrast (Rachlin, 1973). However, while these relative influences may affect the precise details of the relationship between responding and reinforcement, by for example affecting the slope of the line between i and I, it is less likely that they will alter the general form of that relationship.

Thus far, we have considered how the present results can be accommodated by an error-correction learning rule (Equation 1), allowing that the expression of learning may be either linearly related to responding or subject to a behavioral normalization operation (the hyperbolic function). However, the present results also speak to a third possibility – that conditioned responding is subject to behavioral normalization, as just described, but learning itself is linearly related to *r*. That is, when rats are exposed to pairings of a CS and US, the content of their learning is the rate at which the US occurs during the CS. The idea that conditioning is based on the acquisition of knowledge about *r*, rather than the accumulation of associative strength, has been argued in detail by Gallistel and Gibbon (2000) as the basis of their Rate Estimation Theory (RET). Indeed, these theorists argue that animals use their knowledge about reinforcement rates to guide their ongoing choice between different behaviors. In conditioning schedules such as the present one, that choice would be determined by each CS's *r*, normalized by the reinforcement of all behaviors in the test environment, a computation that is captured by the hyperbolic function. Thus the present results are also fully consistent with this theoretical approach to conditioning. Further experimental work will be needed to distinguish between this particular approach and an associative account based on an error-correction learning rule.

References

Bush, R. R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review, 58*, 313-323.

Church, R. M., & Broadbent, H. A. (1990). Alternative representations of time, number, and rate. *Cognition, 37*, 55-81.

Davis, M., Schlesigner, L. S., & Sorenson, C. A. (1989). Temporal specificity of fear conditioning: Effects of different conditioned stimulus-unconditioned stimulus intervals on the fear-potentiated startle effect. *Journal of Experimental Psychology: Animal Behavior Processes, 15*, 295-310.

Davison, M., & McCarthy, D. (1988). *The matching law: A research review.* Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.

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

Desmond, J. E., & Moore, J. W. (1988). Adaptive timing in neural networks: The conditioned response. *Biological Cybernetics, 58*, 405-415.

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

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

Grossberg, S., & Schmajuk, N. A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks, 2*, 79-102.

Harris, J. A. (2006). Elemental representations of stimuli in associative learning. *Psychological Review, 113*, 584-605.

Harris, J. A. (in press). The acquisition of conditioned responding. *Journal of Experimental Psychology: Animal Behavior Processes*.

Harris, J. A., Gharaei, S., & Pincham, H. L. (in press). Response rates track the history of reinforcement times. *Journal of Experimental Psychology: Animal Behavior Processes*.

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13*, 243-266.

Holland, P. C. (1979). Differential effects of omission contingencies on various components of Pavlovian appetitive conditioned responding in rats. *Journal of Experimental Psychology: Animal Behavior Processes, 5*(2), 178-193.

Holmes, N. M., Marchand, A. R., & Coutureau, E. (2010). Pavlovian to instrumental transfer: a neurobehavioural perspective. *Neuroscience and Biobehavioral Reviews, 34*, 1277-1295.

Kehoe, E. J., & Joscelyne, A. (2005). Temporally specific extinction of conditioned responses in the rabbit (Oryctolagus cuniculus) nictitating membrane preparation. *Behavioral Neuroscience, 119*, 1011-1022.

Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review, 95*, 274-295.

Kirkpatrick, K., & Church, R. M. (2003). Tracking the expected time to reinforcement in temporal conditioning procedures. *Learning & Behavior, 31*, 3-21.

McLaren, I. P. L., & Mackintosh, N. J. (2000). An elemental model of associative learning: I. Latent Inhibition and perceptual learning. *Animal Learning & Behavior, 28*, 211-246.

Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex.* (G. V. Anrep, Trans.). New York: Dover.

Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review, 101*, 587-607.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109*, 472-491.

Rachlin, H. (1973). Contrast and matching. *Psychological Review, 80*, 217-234.

Rescorla, R. A. (2000). Associative changes in excitors and inhibitors differ when they are conditioned in compound. *Journal of Experimental Psychology: Animal Behavior Processes, 26*, 428-438.

Rescorla, R. A., & Wagner, A. R. (1972). 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.* (pp. 64-99). New York: Appleton-Century-Crofts.

Roberts, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes, 7*, 242-268.

Smith, M. C. (1968). CS-US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology, 66*, 679-687.

Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review, 88*, 135-171.

Wagner, A. R. (1981). SOP: a model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: memory mechanisms* (pp. 5-47). Hillsdale, NJ: Erlbaum.

Wagner, A. R., & Brandon, S. E. (2001). A componential theory of Pavlovian conditioning. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories.* (pp. 23-64). Mahwah NJ, USA: Lawrence Erlbaum Associates, Inc.

Williams, D. A., Lawson, C., Cook, R., Mather, A. A., & Johns, K. W. (2008). Timed Excitatory Conditioning Under Zero and Negative Contingencies. *Journal of Experimental Psychology: Animal Behavior Processes, 34*(1), 94-105.

Appendix A

Here we describe how the relationship between V and r (and between V and i) can be derived from Equation 1. When calculating V, λ takes its nominal value (e.g., 1) when the US is presented, and takes a value of zero when the US is absent. Thus, V will increase when the US occurs during the CS, and will decrease at each moment that the CS is present without the US. The amount of extinction, V^{-}, that accumulates across non-reinforced exposure to the CS is equal to the sum of each of these moment-by-moment decrements in V. That is,

$$V^{-} = (-k_0 V_0 + -k_0 V_1 + \ldots + -k_0 V_{i-1})$$

Here, V_t, is the associative strength at time t, and thus V_0 is V at the onset of the CS, and $V_{i\text{-}1}$ is V at the final moment of non-reinforced exposure to the CS (before the US occurs at time i); k_0 is the learning rate parameter for non-reinforcement. We can rewrite the above equation as:

$$\begin{aligned} V^{-} &= -k_0(V_0 + V_1 + \ldots + V_{i-1}) \\ &= -k_0\, i\, \overline{V} \end{aligned}$$

where $\overline{V}$ is the mean of V as it changes across the length, i, of the CS. The steady-state associative strength will be reached after many trials once the extinction of V during non-reinforced exposure to the CS is matched by the increase in V when the US occurs. Using Equation 1 to calculate the increase in V at the moment of reinforcement, extended conditioning will reach a steady state when

$$k_0\, i\, \overline{V} = k_1(\lambda \;\; - V_i) \qquad \text{Equation 2}$$

where k_1 is the learning rate parameter for reinforcement. In order to derive the relationship between $\overline{V}$ and i from the above equation, it is necessary to remove the term V_i, replacing it with an expression in terms of $\overline{V}$. To specify V_i, we must first define the value of V at any time, t, during a trial. V_t can be calculated using the integral of the extinction function:

$$V_t = V_0\, \mathrm{e}^{-k_0 t}$$

Therefore, at CS onset,

$$V_0 = V_i\, \mathrm{e}^{k_0 i} \qquad \text{Equation 5}$$

$\overline{V}$ can be calculated as

$$\overline{V} = \frac{\sum_0^i V_t}{i}$$

$\sum_0^i V_t$ is calculated by integrating V_t from $t = 0$ to i, as

$$\begin{aligned}
\sum\nolimits_0^i V_t &= \int_0^i V_0\, e^{-k_0 t}\, dt \\
&= \left[\frac{V_0}{-k_0}\, e^{-k_0 i} + C\right] - \left[\frac{V_0}{-k_0}\, e^{-k_0 0} + C\right] \\
&= \left[\frac{V_0}{-k_0} \cdot e^{-k_0 i}\right] - \left[\frac{V_0}{-k_0}\right] \\
&= \left[\frac{V_0}{k_0}\right] - \left[\frac{V_i\, \mathrm{e}^{k_0 i}}{k_0} \cdot e^{-k_0 i}\right] \\
&= \left[\frac{V_0}{k_0}\right] - \left[\frac{V_i}{k_0}\right]
\end{aligned}$$

Therefore,

$$\overline{V} = \frac{(V_0 - V_i)}{k_0 \cdot i}$$

Substituting Equation 5 into the above equation gives

$$\overline{V} = \frac{V_i(e^{k_0 i} - 1)}{k_0\, i}$$

which can be rearranged as

$$V_i = \overline{V} \frac{k_0\, i}{(e^{k_0 i} - 1)} \qquad \text{Equation 6}$$

Substituting this into Equation 2 gives

$$k_0\, i\, \overline{V} \;=\; k_1\lambda \;\; - k_1 \frac{\overline{V}\, k_0\, i}{(e^{k_0 i} - 1)}$$

and thus,

$$\overline{V} \;=\; \frac{k_1\lambda}{k_0\, i \left(1 + \frac{k_1}{(e^{k_0 i} - 1)}\right)}$$

The reciprocal of this equation can be rearranged to give Equation 7:

$$\frac{1}{\overline{V}} = \frac{k_0}{k_1\lambda} i + \frac{1}{\lambda}\left[\frac{k_0\, i}{(e^{k_0 i} - 1)}\right] \qquad \text{Equation 7}$$

The term inside the square brackets on the far right is the only feature that distinguishes Equation 7 from Equation 3. This term is an exponential decay function that approaches 1 as $i \rightarrow 0$, and gradually approaches 0 as $i \rightarrow \infty$. Therefore, Equation 7 describes a function that diverges from a straight line with slope = $k_0/k_1\lambda$ and intercept = $1/\lambda$, and converges on a parallel line with intercept = 0. In other words, Equation 7 is initially approximated by the straight line described by Equation 3, and as i increases it becomes more closely approximated by the linear function $1/\overline{V} = k_0 i/k_1\lambda$. The similarity between Equations 7 and 3 depends on k_0 and λ. The distance that the two equations diverge depends on λ, because Equation 7 ultimately converges on a parallel line whose distance from Equation 3 equals $1/\lambda$. The rate at which the equations diverge is determined by k_0: the smaller k_0, the closer the bracketed term in Equation 7 remains to 1, and thus the more similar the two Equations are as i increases. Therefore, the closeness with which Equation 3 approximates Equation 7 is lower when k_0 is large and λ is small. Under these circumstances, steady state associative strength will be very low. On the other hand, if λ is large relative to k_0, Equation 7 is closely approximated by Equation 3. This, we believe, is true for the experiments presented here, since the rate of extinction observed across the course of a trial was very low relative to the level of responding. In this regard, we note that the bracketed term in Equation 7 is identical to the expression on the far right of Equation 6, and from that equation we can see that, when k_0 is small, V_i is approximately equal to $\overline{V}$. This confirms the premise on which we derived Equation 3 in the Introduction.

Published in *Journal of Experimental Psychology: Animal Behavior Processes, vol 37 (4), pp 375-384.*
http://www.apa.org/pubs/journals/xan/index.aspx