

Information, certainty, and learning

Justin A Harris , CR Gallistel

The University of Sydney, Sydney, Australia • Rutgers University, New Brunswick, United States

Reviewed Preprint

v1 • November 20, 2024

Not revised

 https://en.wikipedia.org/wiki/Open_access

 Copyright information

eLife Assessment

This **important** paper shows that the acquisition and expression of Pavlovian conditioned responding are lawfully related to temporal characteristics of an animal's conditioning experience. It showcases a rigorous experimental design, several different approaches to data analysis, careful consideration of prior literature, and a thorough introduction. The evidence supporting the conclusions is strong and **convincing**. The paper will have a general appeal to those interested in the behavioral and neural analysis of Pavlovian conditioning.

<https://doi.org/10.7554/eLife.102155.1.sa2>

Abstract

More than four decades ago, Gibbon and Balsam (1981) showed that the acquisition of Pavlovian conditioning in pigeons is directly related to the informativeness of the conditioning stimulus (CS) about the unconditioned stimulus (US), where informativeness is defined as the ratio of the US-US interval (C) to the CS-US interval (T). However, the evidence for this relationship in other species has been equivocal. Here, we describe an experiment that measured the acquisition of appetitive Pavlovian conditioning in 14 groups of rats trained with different C/T ratios (ranging from 1.5 to 300) to establish how learning is related to informativeness. We show that the number of trials required for rats to start responding to the CS is determined by the C/T ratio and, remarkably, the specific scalar relationship between the rate of learning and informativeness aligns very closely to that previously obtained with pigeons. We also found that the response rate after extended conditioning is strongly related to T , with the terminal CS response rate being a scalar function of the CS reinforcement rate ($1/T$). Moreover, this same scalar relationship extended to the rats' response rates during the (never-reinforced) inter-trial interval, which was directly proportional to the contextual rate of reinforcement ($1/C$). The findings establish that animals encode rates of reinforcement, and that conditioning is directly related to how much information the CS provides about the US. The consistency of the data across species, captured by a simple regression function, suggests a universal model of conditioning.

More than a century of laboratory-based research has been devoted to investigating how animals learn about simple relationships between events, such as learning to respond to a conditioned stimulus (CS) that is followed by an unconditioned stimulus (US) or learning to perform a specific action that is reinforced by a rewarding US. Much of that research has focussed on identifying what properties about the CS-US or response-US relationship are most important for learning. There is widespread consensus about the importance of three particular properties. One is the

temporal contiguity between the events: conditioning emerges sooner when the US follows the CS or response closely in time. Another is the *spacing* of the learning trials: conditioning takes fewer trials when there is a long time-interval between each CS-US or response-US pairing. The third property is the *contingency* between the events: conditioning is more successful when the US occurs reliably in the presence of the CS or response, and does not occur in their absence.

Acquisition of Conditioning, C-over-T, and Informativeness

A landmark study, conducted more than 40 years ago, demonstrated that the first two of these properties—contiguity and trial spacing—are interdependent and subserved by a single principle that encompasses all three properties. In two large-scale experiments, [Gibbon and colleagues \(1977\)](#)  measured the number of trials required for pigeons to start pecking at a key-light (the CS) as they learned it was followed by food (the US). Different groups of pigeons were trained with different trial durations (i.e., the interval between onset of the CS and delivery of the US; henceforth T). T ranged from 1 s to 64 s across 41 groups in two experiments. The inter-trial interval (ITI) also varied between groups, ranging from 6 s to 768 s. [Gibbon et al.](#) recorded the number of training trials required for each pigeon to start responding reliably during the CS (responding on 3 out of 4 consecutive trials). The birds required more trials as T increased, confirming the effect of CS-US contiguity. They also required fewer trials as the ITI increased, confirming the trial-spacing effect. More importantly, however, these two effects were completely complementary, such that an increase in T had no effect if it was accompanied by a proportional increase in the ITI. Thus, the rate at which the pigeons acquired the conditioned response systematically increased as the ratio of ITI to T increased, but there was no separate effect of varying ITI or T when their ratio was held constant.

The extent of the relationship between contiguity and trial-spacing was established by a meta-analysis that combined data from all 41 groups tested by [Gibbon et al. \(1977\)](#)  with data from 11 other experiments investigating the acquisition of key-peck responses in pigeons ([Gibbon & Balsam, 1981](#)) . This analysis compared the number of trials to criterion against the ratio of C/T (where C is the interval between USs, equal to $T + \text{ITI}$). Their data on pigeon autoshaping are well described by a regression model whose only parameter is the x-intercept, which is the C/T value that produces acquisition after a single trial in the median pigeon (k in **Figure 1**) . When $C/T > 4$, as it is in most Pavlovian protocols, the slope of the regression model on loglog coordinates is approximately -1 . Thus, over most of the C/T range, the number of reinforcements at acquisition is inversely proportional to the C/T ratio; doubling the ratio reduces required reinforcements by a factor of 2.

When reinforcements are delivered only during the CS, the C/T ratio is the ratio of the rate of reinforcement during CSs to the *contextual rate*, the rate subjects expect when in the experimental chamber, without regard to whether the CS is or is not present. [Balsam and Gallistel \(2009\)](#)  termed the ratio of the CS rate of reinforcement to the contextual rate the *informativeness*, because the log of informativeness is the *mutual information* that CS onsets transmit to a subject about the expected wait for the next reinforcement (lower x-axis in **Figure 1**) . The mutual information transmitted by a CS is the upper limit on the extent to which the CS can reduce the subject's uncertainty about the wait for the next reinforcement. When the CS rate equals the contextual rate, no information is transmitted because the informativeness ratio is 1, and $\log(1) = 0$.

Contingency is mutual information divided by available information ([Gallistel & Latham, 2023](#)) . *Available information* is the amount that reduces a subject's uncertainty to 0. Because temporal measurement error scales with duration measured (Weber's Law, [Gibbon, 1977b](#)) , contingency = 1 only when two events coincide in time ([Gallistel & Latham, 2023](#)) .




Figure 1.

Median reinforcements to acquisition plotted against the C/T ratio

Note. Asterisks show the data from Gibbon and Balsam's (1981) metaanalysis of acquisition in pigeon acquisition. The two open circles are from Jenkins, Barnes and Barrera (1981) [who used even bigger C/T ratios in some of their groups](#). The log of the informativeness is the mutual information, I , between the CS and the expected wait for reinforcement (λ_R ; lower x axis). The learning rate is the reciprocal of reinforcements to acquisition (right axis). The regression model was fit only to the asterisks, but it also predicts the Jenkins et al. data. The value k is the x intercept, the informativeness that produces one-trial learning. The regression model curves upward to infinity as informativeness goes to 1 and mutual information to 0 because, when the CS rate does not differ from the contextual rate of reinforcement, differential responding to the CS does not emerge (Rescorla, 1967 [who used even bigger C/T ratios in some of their groups](#), 1968 [who used even bigger C/T ratios in some of their groups](#)).

The analysis presented thus far identifies a fundamental principle of associative learning; what animals learn about events is defined in terms of measurable properties of their temporal distributions. To claim this as a general principle of learning, we must seek evidence for the importance of C/T in conditioning paradigms with other species. Several studies have sought evidence that conditioning is related to C/T using an appetitive Pavlovian conditioning paradigm in which rats or mice learn to anticipate the arrival of a food reward (indexed by monitoring their activity at the food cup or licking at a spout) during a CS (Bouton & Sunsay, 2003 [DOI](#); Burke, Jeong, Wu, Lee, & Namboodiri, 2023 [DOI](#); Holland, 2000 [DOI](#); Kirkpatrick & Church, 2000 [DOI](#); Lattal, 1999 [DOI](#); Thraillkill, Todd, & Bouton, 2020 [DOI](#); Ward et al., 2012 [DOI](#)). However, as summarised below, the evidence from these studies is mixed.

The first evidence for an effect of C/T on the acquisition of responding in rats was reported by Lattal (1999) [DOI](#) and Holland (2000), who observed effects that persisted even when all rats were tested with an identical ITI (different groups had been trained with different itis but were shifted to a common interval between trials when tested for responding to the CS). More recently, Ward et al. (2012) [DOI](#) reported that the log of number of trials for mice to acquire responding scaled with $\log(C/T)$, and this relationship was similar to that described for pigeons (Gibbon et al., 1977 [DOI](#); Gibbon & Balsam, 1981 [DOI](#)). Bouton and Sunsay (2003) [DOI](#) also provided evidence for the importance of C/T on Pavlovian conditioning in rats. They showed that conditioning was negatively affected when T was increased 3-fold by inserting two CS-alone presentations between each CS-US trial (a manipulation that affects T but does not affect C and, therefore, reduces C/T) but that conditioning was not affected by a 3-fold increase in T brought about by omitting the US from two out of three CS-US trials (a manipulation that increases both T and C equally and, therefore, does not change C/T). Most recently, Burke et al. (2023) [DOI](#) have shown that removing 9 out of 10 CS-US trials from a Pavlovian conditioning schedule, thus increasing C 10-fold without changing T , reduced the number of trials required for learning by a factor of 10.

In addition to finding an effect of C/T on conditioning, both Lattal (1999) [DOI](#) and Holland (2000) also found an effect of T that was independent of the C/T ratio. They observed that an increase in T resulted in a decrease in responding even when the ITI was also increased to keep the C/T ratio constant. Further evidence against an influence of C/T ratio comes from a study by Kirkpatrick and Church (2000) [DOI](#) in which differences in C/T , ranging from 1.5 to 12, did not produce differences in the acquisition of conditioned responding in rats. Most recently, Thraillkill et al. (2020) [DOI](#) observed clear effects of T , but not C/T , on conditioning. Their rats responded at much higher rates to a 10-s CS than to a 60-s CS, even though the groups had identical C/T ratios. When comparing groups on how quickly responding was acquired, they found no systematic effect of C/T on the number of 4-trial conditioning blocks required to reach a response criterion.

In sum, in contrast with the impressive evidence from experiments with pigeons, studies of appetitive conditioning in rats and mice have provided inconsistent evidence for the importance of C/T to conditioning. At the same time, these studies have shown that T has an effect on responding that is independent of C/T . This latter observation is in fact consistent with the evidence from experiments with pigeons. In the original study which established that C/T determines when pigeons start to respond to a CS, Gibbon et al. (1977) [DOI](#) also observed that the rate at which the pigeons responded to the CS after extended training was negatively related to T and not related to C/T . In other words, while C/T affected how quickly conditioned responding emerged, T , rather than C/T , determined the level of responding that was ultimately acquired. This distinction may go some way to explaining the inconsistency in the evidence for the effect of C/T in rats and mice. The inconsistencies may be due to differences in how the point of acquisition was identified in different studies, and, in particular, whether differences in the amount of responding affected the measure of when responding was first acquired. This concern is not relevant to the results from pigeon experiments because the appearance of their key-peck response is unambiguous thanks to the fact that the baseline rate of that response is effectively zero. In

contrast, rats and mice show activity at the food-cup in the ITI, and therefore researchers using this appetitive paradigm must decide how to take account of the baseline response rate when quantifying conditioned responding during the CS (Lattal, 1999 [↗](#)).

The current experiment

The experiment described here attempts to elucidate the role of C/T and T in an appetitive Pavlovian conditioning paradigm with rats by distinguishing their impact on the emergence of responding from their effect on the level of responding subsequently acquired after extended conditioning. Fourteen groups of rats were trained with for 42 sessions with a single CS that was followed on every trial by delivery of food (the US). Each session contained either 10 CS-US trials (Groups 1-11) or 3 CS-US trials (Groups 12-14). Both T and C varied between the groups in an uncorrelated fashion ($r = -0.19$, $p = .519$) so that effects of T and C could be assessed independently (summarised in **Table 1** [↗](#)).

When the location of the CS is spatially separated from the US, as in most appetitive conditioning experiments with pigeons or rodents, the two types of response are mutually incompatible which can impact on the measurement of conditioned responding. For example, any factor that increases food-cup activity, such as reducing the ITI, may reduce evidence for conditioning that is indexed by CS-directed responses like key-pecks. Conversely, evidence for conditioning indexed by food-cup activity may be reduced to the extent that animals also acquire CS-directed responses. These problems can be mitigated if the CS and US are co-located. In the present experiment, the CS was illumination of a small LED inside the magazine. Conditioned responses were measured using an infra-red photobeam across the opening of the magazine that should detect both food-cup activity and approach responses to the CS. CS-US intervals varied randomly from trial to trial (around a mean equal to T) so that response rates remained constant across the length of each trial (Harris & Carpenter, 2011 [↗](#); Harris, Gharaei, & Pincham, 2011 [↗](#)).

Results

Rats in all groups eventually responded at a higher rate to the CS than during the ITI. **Figure 2** [↗](#) shows the mean CS and ITI response rates for each group across the 42 conditioning sessions. The equivalent of this figure for each rat is included in the Supplementary material. The last panel of **Figure 2** [↗](#) (Plot 15) shows the mean response rate per second during the CS for each group, averaged from all trials of the final 5 sessions of the experiment. This shows how responding increased rapidly over the first few seconds of the trial and then remained constant across the trial, consistent with previous experiments using variable CS-US intervals (Harris & Carpenter, 2011 [↗](#); Harris et al., 2011 [↗](#)).

Trials to acquisition

We used the cumulative response rates to identify the trial at which conditioned responding first emerged. The earliest evidence for responding to the CS was the trial after which the cumulative response rate during the CS permanently exceeded the cumulative ITI response rate, computed for each rat. Starting from this trial, the difference between CS and ITI cumulative response rates was subjected to one-tail t -tests in each rat to identify the trial at which CS responding was significantly greater than pre-CS responding. The results obtained using $p < .05$ are shown in **Figure 3A** [↗](#). This shows the number of trials required for each rat to reach this statistical criterion, plotted against the informativeness of the CS for that rat (C/T). The filled orange circles in the same figure show the median number of trials for each group of rats trained with the same C/T ratio. These data are

Group	<i>T</i>	<i>T</i> range	<i>C</i>	ITI [range]	<i>C/T</i>
1	42	2-82	63	21 [15-27]	1.5
2	48	2-94	144	96 [15-177]	3
3	21	2-40	94.5	73.5 [15-132]	4.5
4	12	2-22	72	60 [15-105]	6
5	6	2-10	54	48 [15-81]	9
6	62	2-122	930	868 [15-1721]	15
7	30	2-58	600	570 [15-1125]	20
8	18	2-34	486	468 [15-921]	27
9	36	2-70	1296	1260 [15-2505]	36
10	26	2-50	1404	1378 [15-2741]	54
11	8	2-14	576	568 [15-1121]	72
12	16	2-30	1760	1744 [15-3473]	110
13	10	2-18	1800	1790 [15-3565]	180
14	14	2-26	4200	4186 [15-8357]	300

Note. All values for *T* (CS-US interval), *C* (US-US interval), and the inter-trial interval (ITI) are in seconds.

Table 1.

Summary of groups.




Figure 2.

Mean response rate for each group across Sessions or across time since CS onset.

Note. There were 10 (groups with $C/T \leq 72$) or 3 (groups with $C/T \geq 110$) trials per session. Plot 15 (bottom right) shows the mean response rate per second during the CS, averaged from the last 5 sessions. Each group is identified by the length of the mean CS-US interval (T).

superimposed on the data shown in **Figure 1** that plots the results of a meta-analysis of many comparable experiments with pigeons (Gibbon & Balsam, 1981), revealing a close agreement between the present results and those previous results.

Correlational analyses, with α set at .017 to correct for multiple tests, were used to assess the relationships between the log of trials to criterion (based on the $p < .05$ criterion shown in **Figure 3A**) and $\log(C/T)$, $\log(C)$, and $\log(T)$. $\log(C/T)$ was the strongest predictor of trials to criterion ($r = -0.90$, $p < .001$). $\log(C)$ was also significantly correlated with the log of trials to criterion, $r = -0.78$, $p < .001$, but $\log(T)$ was not correlated with trials to criterion, $r = 0.43$, $p = .13$. Given that $\log(C)$ and $\log(C/T)$ were strongly correlated with one another ($r = 0.90$), their partial correlations with the log of trials to criterion were calculated to test whether each made independent contributions to the rate of acquisition. After partialling out the effect of $\log(C)$, the correlation with $\log(C/T)$ remained significant, $r = -0.73$, $p = .005$. In contrast, after partialling out the effect of $\log(C/T)$, the correlation with $\log(C)$ was not significant, $r = 0.15$, $p = .629$. In sum, the number of trials required for conditioned responding to emerge was strongly affected by the C/T ratio, and neither C nor T alone had any independent effect.

Figure 3 also shows the number of reinforced trials to acquisition when the acquisition criterion was defined using the information-theoretic statistic, the nD_{KL} . **Figure 3B** plots the number of trials for the nD_{KL} to become permanently positive; from this trial on, there is consistently positive evidence that the distribution of CS response rates and Pre response rates diverge. The black regression line fitted to the pigeon data (see **Figure 1**) accounts for 67% of the variance in the median trials to criterion ($R^2 = 0.67$). **Figure 3C** plots the number of trials to reach an odds ratio of 10:1 in favour of a difference between CS and Pre response rates, based on the nD_{KL} .

In **Figure 3** it can be hard to appreciate the distribution of trials to acquisition data when the data for individual subjects (open grey diamonds) are superimposed. This is particularly the case for subjects learning after their first reinforcement, where individual data pile up on the x-axis, even when the informativeness is as low as 4. **Figure 4** plots trials to acquisition (shown with reversed scale on the right axis; for odds 10:1) as a function of the mutual information between CS and US (bottom axis; equal to $\log_2$ of the informativeness, shown on the top axis). In this plot, the individual data for trials to acquisition data have been grouped in bins of approximately equal width on a logarithmic scale. Thus, the tiles in the first 4 rows represent single numbers of trials (1, 2, 3, and 4 trials to acquisition) and tiles in subsequent rows represent bins with wider ranges (trials 5-6, 7-9, 10-15 etc). The darkness of each tile in **Figure 4** represents the number of rats within that bin. The left-hand axis shows the learning rate, computed as the reciprocal of number of trials to acquisition. **Figures 3** and **4** show the large individual differences in trials to acquisition; subjects within the same informativeness group may have trials-to-acquisition values that differ by two orders of magnitude. Nonetheless, overall, the learning rate increases as the informativeness increases. The figure also makes clear that the number of rats learning after just one reinforced trial (top row) increases as the informativeness increases.

As a measure of the terminal level of responding, the response rates during the CS were averaged over the last 5 sessions for each rat. This rate was calculated as total number of responses (summed across all trials over the 5 sessions) divided by the total CS duration (summed across all trials over the 5 sessions). Our initial analysis compared this terminal response rate with the log of T , C , and C/T ($\alpha = 0.017$ after correction for multiple comparisons). The rate of responding to the CS was marginally correlated with $\log(T)$ (see **Figure 5A**), $r = -0.53$, $p = .051$, but was not correlated with either $\log(C)$, $r = -0.3$, $p = .298$, or $\log(C/T)$, $r = -0.04$, $p = .894$.

More detailed analyses were conducted after segmenting the response rate into 3 separate components: (1) the latency to the first response in a trial; (2) the mean duration of each response (the time spent in the magazine); and (3) the inverse of the mean inter-response interval, $1/IRI$,




Figure 3.

Number of trials for acquisition of responding to the CS, plotted against CS informativeness (C/T ratio).

Note. The number of reinforced trials to reach an acquisition criterion for each rat (grey diamonds) and the median number for each group (orange circles) are plotted against informativeness of the CS, on double-log coordinates. These data are superimposed onto the data (black asterisks) from [Gibbon & Balsam \(1981\)](#) and the regression line shown in [Figure 1](#). The different plots show trials to acquisition using different acquisition criteria: **A.** responding during the CS is significantly greater than pre-CS responding by t-test ($p < .05$); **B.** when the nD_{KL} became permanently positive; **C.** when the odds reached 10:1 that the CS response rate was greater than the ITI rate (according to the nD_{KL}).




Figure 4.

The learning rate (left axis; as reciprocal of the number of trials to acquisition, right axis) plotted against the mutual information (log2 of informativeness) for the 14 groups.

Note. The criterion for acquisition was the trial on which the nD_{KL} reached odds of 10:1 that the CS response rate was greater than the Pre response rate. The number of trials to acquisition has been divided into 14 bins of approximately equal width on a logarithmic scale represented by the 14 rows of tiles. The darkness of each tile shows the number of rats that reached criterion within that range of trials (see greyscale bar on the right).

Terminal response strength.




Figure 5.

Response metrics, averaged over the final 5 conditioning sessions, plotted against T .

Note: Open grey circles show data for individual rats. Black filled circles show mean data for each group. The dashed orange lines are the best-fitting regression lines.

which equals the response rate after excluding the latency and response durations. None of these indices correlated significantly with $\log(C)$, largest $r = -0.35$, $p = .221$, or with $\log(C/T)$, largest $r = -0.31$, $p = .281$. On the other hand, $\log(T)$ correlated significantly with latency, $r = 0.77$, $p < .001$ (Figure 5B), and with $1/IRI$, $r = 0.89$, $p < .001$ (Figure 5C). There was a marginal correlation between duration of responding and $\log(T)$, $r = -0.55$, $p = .043$, that did not survive correction for multiple comparisons (Figure 5D). As is evident in Figure 5B, the positive correlation between latency and $\log(T)$ was largely confined to groups with long CS-US intervals ($T > 25$ s), whereas latency varied little among groups with shorter CS-US intervals. This invariance at short CS-US intervals may have been due to a floor effect because mean latencies to first response did not decrease below 2 s (horizontal dotted line in Figure 5B). This is also consistent with the plots of response by time-in-CS, shown in plot 15 of Figure 2, where responding to the CS was low for the first few seconds after CS onset. This apparent floor effect might reflect a constraint on how quickly the rats can commence responding after CS onset or it might have arisen because 2 s was the minimum CS-US interval used in this experiment. Regardless, these analyses indicate that neither duration of responding nor latency to first response are good markers of what the rats learn about the rate of reinforcement of the CS. By contrast, the response rate, $1/IRI$, varied systematically across the entire range of values of T (Figure 5C). This confirms earlier evidence that rats' response rates scale with the log of the reinforcement rate (Harris & Carpenter, 2011).

The response rates in Figure 5A are computed by dividing the response count by the duration of the interval over which the count was made. However, when the CS reinforcement rate was high, subjects made from 2, to more than 5, pokes per second and the pokes lasted substantial fractions of a second. A rat cannot make another poke when its head is in the magazine. The correct computation of a response rate should use the time over which it was possible for pokes to be counted, that is, the cumulative time minus the cumulative head-in-magazine time. Also, the latency of the first poke was rarely shorter than 2 s (Figure 5B); whereas the inter-poke intervals after the first poke were much shorter when T was short (Figure 5C). To make a proper estimate of the rate at which a rat poked once it had begun, we divided the number of pokes made during the CS (excluding the first poke) by the remaining time in which it was possible to initiate a poke (i.e., the CS duration minus the latency to 1st poke and the cumulative head-in-magazine time). Likewise, to estimate the contextual rate of poking, we divided the ITI response count by the time in which an ITI response could be initiated (i.e., the length of the pre-CS interval minus the cumulative time with head in magazine). These methods for computing the CS and contextual rates of responding correspond closely to $1/IRI$ (shown in Figure 5C) except that they can still be computed when there are fewer than 2 responses (unlike the IRI). Figure 6 plots the terminal CS response rates (black dots) and contextual response rates (red dots) as functions of the CS/contextual reinforcement rate on double logarithmic coordinates.

Figure 6 shows that the log response rate while in the experimental context and the log response rate during CSs are both described by the same linear regression when plotted against the relevant reinforcement rate. It accounts for 81% of the variance. Its slope is essentially 1, which means that the two poke rates are the same *scalar* function of the two reinforcement rates. The scalar that relates response rate to reinforcement rate is $10^{1.25} = 18$, where 1.25 is the constant in the logarithmic regression. Thus, the average rates of poking are 18 times faster than the reinforcement rates. This maximally simple relation between reinforcement rate and poke rate holds from reinforcement rates almost as low as 1 in 100 min and corresponding poke rates as slow as 5 or 6/min up to reinforcement rates as high as 10/min (1 every 6 s) and corresponding poke rates of 90-110/min (1 every 0.5 s).

Put another way, the average time the rat waits between magazine entries is $1/18 (= .06)$ times the average time it expects to wait for food (also Harris & Carpenter, 2011). The same scalar relation between the behavioural wait and the expected wait for reinforcement applies when the rat is in an intertrial interval—when reinforcement is never delivered—and when it is in a CS interval. As the informativeness of a protocol gets lower, that is, as the CS occupies a greater and greater




Figure 6.

Terminal response rates (final 5 sessions) of individual rats as a function of the reinforcement rate on double-logarithmic coordinates.

Note. Red circles show response rates during the CS (number of pokes, after the first poke, divided by the remaining CS time when the rat's head was out of the magazine) plotted against the CS reinforcement rate (1/T). Black circles show response rates during the inter-trial interval (ITI; calculated as number of pokes in the pre-CS ITI divided by the pre-CS time with head out of the magazine) plotted against the baseline reinforcement rate (1/C). Filled circles show data of rats given 420 reinforced trials (Groups with C/T ratios ≤ 72), and open circles show data of rats given 126 reinforced trials (Groups with C/T ratios ≥ 110). The inset gives the equation for the dashed black regression line ($R^2 = 0.81$, $\text{rmse} = 0.39$). The 95% confidence interval for the slope of this line is [0.93 1.03], and thus includes 1. The confidence interval of the intercept is [1.2 1.29].

fraction of the subject's time in the chamber, the difference between the CS poke rate and the contextual poke rate gets smaller because the difference between the CS reinforcement rate and the contextual reinforcement rate gets smaller.

As is also apparent in [Figure 6](#), the variability about the scalar relation between poke rate and reinforcement rate also scales with reinforcement rate. That is why the regression must be computed in the logarithmic domain. The rmse of the loglog base 10 regression is 0.39. Thus, 68% of the observed poke rates fall within a factor of between 7.2 and 45 times the reinforcement rate.

Trials from onset of responding to peak responding

Having established the relationship between response rate and reinforcement rate, we next analysed how response rate increased over trials towards its maximum value. Our first analysis assumes that there is a consistent (monotonic) increase in response rate starting from the initial point of acquisition. This analysis followed a method recently described ([Harris, 2022](#)) that uses the slope of the cumulative response rate over trials to identify the trial on which the response rate had reached each decile (from 10% to 90%) of the peak response rate. Based on our earlier analysis (see [Figure 5](#)), our measure of the CS response rate was calculated by dividing the response count by the total time out of the magazine during the CS. As shown in [Figure 7A](#), the response rate of an individual rat varies greatly from trial to trial. However, a clearer picture of the overall change in responding over trials can be obtained by plotting the cumulative response count against the cumulative opportunity to respond (cumulative time out of the magazine; [Figure 7B](#)).

The slope of the cumulative function can be used to estimate the rat's response rate across conditioning to find when responding had reached a given proportion of the peak response rate. To analyse how response rates changed across the course of conditioning, we extracted a segment of each rat's conditioning data starting from the trial, t_1 , on which the response rate during the CS became reliably greater than the ITI response rate and finishing at the trial, t_{end} , on which the response rate reached its peak (according to a moving average with a window width of 3 sessions). The total change in responding across conditioning was calculated by subtracting the response rate at the start of this segment of trials, R_1 , from the peak response rate, R_{max} (at the end of the segment): $\Delta R = R_{max} - R_1$. To identify when responding had increased by 10% of ΔR , we estimated what the cumulative response count, $cumR'$, would be at each trial, t , if the rat maintained a fixed level of responding equal to the starting rate plus 10% of ΔR . Thus, $cumR't = R_1 + 0.1 \cdot \Delta R \cdot cumTt$, where $cumTt$ is the cumulative CS-US interval from trial 1 to t . We then calculated the difference between $cumR't$ and the observed $cumRt$. The trial at which this difference was maximum was identified as the trial when the slope of $cumRt$ had increased by at least 10% of ΔR . This process was repeated for all deciles up to 90%. An example of the values obtained for one rat is shown in [Figure 7C](#). In this example, the rat's response rate had increased by 10% of ΔR on Trial 34, by 50% of ΔR by Trial 147, and by 90% of ΔR by Trial 233.

The analysis illustrated in [Figure 7](#) was conducted on the individual data of all rats (except those rats missing data from Session 1). We excluded rats with a ΔR less than 0.1 responses/s. The mean number of trials to reach each decile for every rat in the 14 groups is shown in [Figure 8](#). With some exceptions, for most rats, the relationship between the number of trials and response decile was roughly linear, meaning that their response rate increased uniformly over trials as it approached the peak response rate. This is clearest in the averaged functions (thick black lines in [Figure 8](#)). To investigate more precisely the relationship between trials to criterion, tc , and response decile, d , we compared 4 different functions for their fit to the data for each individual rat. Based on the apparent linear increase in trials across deciles, the first function tested was a straight line, $tc = m \cdot d + c$. The second was an exponential function, $tc = c \cdot e^{m \cdot d}$, which has a continuously increasing slope and thus predicts a systematic increase across deciles in the number




Figure 7.

Response rate and trials to criterion for an individual rat.

Note. The grey line in **A** shows the response rate on each trial for an individual rat (Rat 10 from Group C/T = 9). The response rate was calculated as the number of responses during the CS divided by the total time out of the magazine but excluding the latency to first response (i.e., the total time across which the animal could respond). In **B**, the cumulative response count across trials for the same rat is plotted against the cumulative opportunity to respond (cumulative time out of the magazine). The slope of this cumulative function was used to identify the trial on which the rat's response rate reached each decile (from 10% to 90%) of its peak response rate, as shown in **C**. These trials are also marked as circles on the response plots in **A** and **B**.

of trials between deciles. This function had most successfully captured the relationship between trials and response deciles in the data analysed by Harris (2022) [\[1\]](#). The third function was an inverse cumulative Gaussian: $tc = s \cdot 2^{1/2} \cdot \text{erf}^{-1}[2 \cdot (m \cdot d + 0.5) - 1] + c$. This was used to model a stepwise increase in responding, as would be observed if responding were governed by a decision process when evidence for the CS-US relationship exceeded some threshold (Gallistel, Fairhurst, & Balsam, 2004 [\[2\]](#)). The fourth function was a log function, $tc = -(\log_e[1-d])/k + c$, derived as the inverse of a cumulative exponential function. This function models the relationship between trials and response criterion predicted by an error-correction learning algorithm such as used by the Rescorla-Wagner model (Rescorla & Wagner, 1972 [\[3\]](#)). To compare between these four models of the data, the Bayesian Information Criterion (BIC) was calculated as

$$BIC = n \cdot \log_e \left(\frac{RSS}{n} \right) + p \cdot \log_e n$$

where RSS is the residual sum of squares for the difference between each observed tc and its corresponding point on the fitted function, n is the number of points being fitted ($= 9$) and p is the number of free parameters in the function.

The overall performance of the functions can be compared by summing the BICs obtained from all rats. The function with the smallest ΣBIC is the function with most evidence. According to this analysis, the linear function had the most evidence, $\Sigma BIC = 9588$, followed by the exponential function, $\Sigma BIC = 9974$. The difference between these ($\Delta BIC = 386$) constitutes overwhelming evidence in favour of the linear model: $BF = e^{\Delta BIC/2} = 7.3 \times 10^{83}$ (Wagenmakers, 2007 [\[4\]](#)). The ΣBIC of the other two models were much higher again, 12128 and 10358, indicating that they had even less support from the data. In addition to this comparison of the aggregate BIC, a more specific comparison can be made by comparing the BIC for one model against the BIC for another for each rat. The scatter plot in the bottom right corner of **Figure 8** [\[5\]](#) plots, for each rat, the BIC for the exponential function against the BIC for the linear function (the two best-fitting functions). The orange line marks where these BICs are equal. The large majority of values (72%) sit above this line. These represent cases where the BIC for the linear function is lower than that for the exponential function, meaning that the evidence is stronger for the linear function. If we look at cases where the difference in BICs was greater than 4.6, corresponding to strong evidence in favour of one function over the other (a $BF \geq 10$), there is strong evidence favouring the linear function over the exponential in 39% of cases, whereas only 9% of cases provide strong evidence in favour of the exponential function. The evidence favouring the linear function over either the inverse cumulative Gaussian or the log function is even stronger: 69% of cases provide stronger evidence ($BF \geq 10$) in favour of the linear over the inverse cumulative Gaussian and 0% of cases favour the latter function; 59% of cases provide strong evidence in favour of the linear over the log function and only 8% of cases strongly favour the log function over the linear.

The above analyses reveal an overall tendency for the response rate to increase approximately linearly up to the point where the peak response rate is reached. To test how the rate of this increase varied across groups, we calculated the correlation coefficient between the slope of the line for each group, as shown in **Figure 8** [\[5\]](#), and the value of C , T , or C/T . These correlations did not include the 3 groups given just 126 trials because the smaller number of trials substantially reduced their slopes. For the other 11 groups, the slope was not significantly correlated with C/T , $r = 0.51$, $p = .110$, or with C , $r = 0.08$, $p = .938$, but was significantly negatively correlated with T , $r = -0.73$, $p = .011$ (see plot titled “Slope” in **Figure 8** [\[5\]](#)). This suggests that the response rate increased more quickly when the reinforcement rate of the CS ($1/T$) was higher, which is not surprising given that the peak rate of responding also increased systematically with reinforcement rate (**Figure 6** [\[5\]](#)).

The preceding analysis assumes that there is an overall monotonic increase in responding over trials. However, this trend is not always apparent, particularly when the CS informativeness is low. Such irregular and non-monotonic changes in responding are revealed by an analysis (see




Figure 8.

Number of trials to each decile of peak response rate.

Note. The coloured line plots show the mean number of trials to each decile for each rat in each of the 14 groups (identified by the C/T ratio). The thick black line in each plot is the average for that group. The slopes of those black lines for each group are plotted against each group's T in the rightmost plot of the third row. The final plot in the bottom right corner of the figure plots the Bayesian Information Criterion (BIC) for an exponential function against the BIC for a straight line when each was fitted to the trials-to-criterion data of individual rats (each open black circle shows the BICs for one rat). For each point in the top left of the plot (above the solid orange line), $BIC_{\text{line}} < BIC_{\text{exponential}}$, meaning that the data were better accounted for by the line than the exponential function (and vice versa for points in the bottom right half of the plot). The two dashed orange lines mark where the difference in BICs equals 4.6 which equates to odds of 10:1 in favour of the function with the lower BIC.

description of the Kullback-Leibler divergence, and the nD_{KL} , in Methods) that uses the nD_{KL} to parse the response rates into segments that have significantly different rates (either higher or lower). In such cases, the path to the peak rate can be bumpy and the peak can be higher than the terminal response rate, as illustrated in [Figure 9](#) for 6 rats from the group with $\iota = 1.5$. By contrast, when the informativeness is very high, the peak is usually reached almost immediately and there is little subsequent variation in the poke rate ([Figure 10](#)). The parsing of response rates shown in [Figures 9](#) and [10](#) used a stringent decision criterion; the algorithm found a step up or down only when the evidence for a divergence exceeded 6 nats. On the null hypothesis, a divergence that large would occur by chance about once in 1800 tests, well above the number of tests made when parsing the data over the 420 or 126 trials of the experiment. When the divergence between the contextual rate of reinforcement and the CS rate is high, as in [Figure 10](#), there are only between 0 and 2 steps in the CS poke rate over 126 trials. By contrast, when the divergence between the contextual rate and the CS rate of reinforcement is low, as in [Figure 9](#), there are many up and down steps, some lasting only a few trials. Thus, the linear increases to a peak revealed by the preceding analyses, should not be taken to indicate that a steady increase is routinely seen in the individual subjects regardless of the informativeness. When informativeness is low, the post-acquisition rate of responding during CSs and the difference between it and the rate during ITIs are unstable but trend upwards. When informativeness is high, a stable asymptotic rate generally appears after only one or two early steps, the first of which is often the step after the first trial.

Discussion

The present results have provided several important pieces evidence about the nature of the conditioned response acquired by rats in an appetitive Pavlovian paradigm. The first is that the initial acquisition of responding was determined by the C/T ratio. Neither T nor C had any independent effect after partialling out the effect of C/T . This is the same conclusion reached by [Gibbon and Balsam \(1981\)](#) in their meta-analysis of the acquisition of autoshaped key-pecking by pigeons. As shown in [Figure 3](#), responding to the CS emerged sooner as the C/T ratio increased, or, more precisely, the log of the number of trials to acquire a response scaled negatively with the log of C/T .

The dependence of trials to acquisition on C/T extends over the entire range from 1 to infinity. Contrary to what Gibbon and Balsam supposed, but in agreement with what [Jenkins et al. \(1981\)](#) showed, the regression continues to describe the data all the way to the abscissa, that is, to acquisition after 1 reinforcement. It is remarkable that the 1-parameter regression model that describes the pigeon acquisition data also describes the current rat magazine-poking data—and with essentially the same parameter value ([Figure 3](#)).

The second finding is that, when each CS is reinforced so that reinforcement rate scales inversely with T , the response rate after extended conditioning is strongly related to T (but not C or C/T). More specifically, the terminal CS response rate is a scalar function of the CS reinforcement rate (red symbols in [Figure 6](#)). Related to this is the third, and unexpected, finding that the response rates during the never-reinforced ITIs is proportional to the contextual rates of reinforcement (black symbols in [Figure 6](#)). It is particularly noteworthy that the constant of proportionality relating the ITI response rates to contextual reinforcement rates is the same as that relating the CS response rates to CS reinforcement rates. To the best of our knowledge, the dependence of the rate of responding during the ITIs on the *contextual* rate of reinforcement—rather than on the *probability* of reinforcement or the *rate* of reinforcement during the ITIs—has not been previously reported.

The fourth finding is that the post-acquisition response rate trends linearly upward to a peak. The higher the rate of reinforcement of the CS, the steeper that linear increase ([Figure 8](#)). However, the path to the peak rate has several ups and downs, particularly when informativeness is low,




Figure 9.

The parsed response rates during the CS (solid lines) and during the ITI (dashed lines) for the first 6 rats in the group with the lowest informative-ness ($\iota = 1.5$).




Figure 10.

The parsed response rates during the CS (solid lines) and during the ITI (dashed lines) for the first 6 rats in the group with the highest informativeness ($\iota = 299$).

and the terminal rate can be lower than the peak rate (**Figure 9** [↗](#)). When informativeness is high, the response rate generally rises to the peak in one or two steps. The first and often final step often occurs after the first trial (**Figure 10** [↗](#)).

The first finding confirms results reported previously in studies of Pavlovian conditioning in pigeons. Perhaps the most remarkable aspect of the observed relationship between trials to acquisition and C/T is that the data collected from conditioning experiments with pigeons and the data reported here are consistent with a simple regression function that has a slope of -1 over most of its range and an intercept of 255. As discussed later, this finding is strong evidence in favour of the notion that conditioning is directly related to how much information the CS provides about the US. That rats poking into a magazine would satisfy the same model with essentially the same value for its one parameter was not to be expected.

The previous published evidence concerning the relationship between C/T and acquisition of Pavlovian conditioning with rodents has been mixed. Lattal (1999) [↗](#) and Holland (2000) reported evidence that conditioning in rats was related to C/T , as did Ward *et al.* (2012) [↗](#) in a series of experiments with mice. However, Kirkpatrick and Church (2000) [↗](#), and, more recently, Thrailkill *et al.* (2020) [↗](#), found no evidence for an effect of C/T on trials to acquisition in rats. Moreover, Lattal, Holland, and Thrailkill *et al.* all found evidence for an effect of T when C/T was held constant. We suggest that these inconsistencies relate to differences in how response acquisition was measured across the studies. Our results show that the point at which evidence for conditioned responding first emerges is directly related to C/T , and not to C or T alone, but the level of responding subsequently acquired is directly related to T and not to C/T or C . Therefore, both C/T and T will affect any index of conditioning that is sensitive to both the time when responding emerges and how much responding is subsequently acquired. Lattal's evidence for an effect of T between groups matched on C/T was obtained in a test session conducted after 4 conditioning sessions totalling 48 trials. Holland observed more responding in the last 8 of 16 conditioning sessions and also found a steeper increase in responding over sessions for rats trained with short values of T despite matched C/T . In both cases, the observed effects of T could have been due to its effect on the level of responding acquired rather than how quickly responding emerged. A similar argument can be made for the evidence provided by Thrailkill *et al.*, based on the evidence that our rats took substantially longer to reach the Thrailkill *et al.* criterion for response acquisition than to reach the criterion we developed (see Supplementary Materials). This suggests that, to satisfy their criterion for acquisition, the rats acquired a higher level of responding which would have been affected by T . Finally, the absence of evidence for an effect of C/T in the study by Kirkpatrick and Church may also have been due to the particular method they used to assess trials to acquisition. Their method, when applied to the present data (see Supplementary materials), produced smaller differences between groups and therefore a weaker (albeit significant) correlation between trials to acquisition and $\log(C/T)$. Thus, their method may have been less sensitive to differences between groups in their rate of acquisition, which could explain why they failed to see an effect of C/T across the relatively limited range of C/T ratios they tested (from 1.5 to 12).

Had it been appreciated in the 1960s that the rate of responding during ITIs has the same scalar relation to the contextual rate of reinforcement as the rate of responding during CSs has to the CS rate of reinforcement, the results from Rescorla's (1967, 1968) truly random control and from Kamin's (1968) blocking experiments would have been considered entirely predictable. In both protocols, the reinforcement rate during the (target) CS is the same as the rate expected in the context in which it occurs. In the truly random protocol, the context is the test chamber; in a blocking protocol, it is the previously conditioned CS.

The contextual rate of reinforcement plays a fundamental role in Rate Estimation Theory (RET, Gallistel & Gibbon, 2000 [↗](#)). It is the first term in the vector of *uncorrected* rates. The uncorrected vector is the list of the *observed* rates of reinforcement for each possible predictor. The *corrected*

rate vector is the list of reinforcement rates subjects *ascribe* to the actions of the different predictors. The matrix equation that does the assignment is based on the assumption they act independently, in which case the ascribed rates must sum to the observed rates when the predictors co-occur. The contextual rate of reinforcement also plays a fundamental role in the information-theoretic model of acquisition (Balsam, Fairhurst, & Gallistel, 2006 [↗](#); Balsam & Gallistel, 2009 [↗](#); Ward et al., 2012 [↗](#)). In that model, the appearance of differential responding to the CS is determined by the informativeness of the protocol, which is the ratio between the CS reinforcement rate and the contextual reinforcement rate. The information transmitted to a subject by CS onset is the log of the informativeness.

The results in **Figure 6** [↗](#) confirm that the subjects compute the contextual rate of reinforcement. It determines how aroused they are when in that context. It has long been recognized that the contextual rate of reinforcement in appetitive operant conditioning has a scalar effect on foraging activities (Belke, 1992 [↗](#); Drew, Zupan, Cooke, Couvillon, & Balsam, 2005 [↗](#); Killeen, 2023 [↗](#); Killeen, Hall, & Bizo, 1999 [↗](#); Killeen, Hanson, & Osborne, 1984 [↗](#)). The effect of the operant contingencies is to channel that activity into the activity or activities on which reinforcement is contingent (Gallistel & Shahan, 2024 [↗](#)). It now appears that the same is true in Pavlovian conditioning. Reinforcement is contingent on poking into the magazine. During the ITIs, subject's poke rate is 18 times the reinforcement rate expected when in the experimental context; during the CSs, it is 18 times to the reinforcement rate expected when in the context and in the presence of the CS.

Information-Theoretic Contingency

Contingency in Pavlovian and operant conditioning has long resisted a mathematical definition that made it measurable in all circumstances, particularly when there is no time at which reinforcement may be anticipated hence no time at which failures of reinforcement to occur can be counted (Donahoe, 2006 [↗](#); Gallistel, 2021 [↗](#); Gibbon, Berryman, & Thompson, 1974 [↗](#); Granger & Schlimmer, 1986 [↗](#); Hallam, Grahame, & Miller, 1992 [↗](#); Hammond & Paynter, 1983 [↗](#)). Gallistel and Latham (2023) [↗](#) have developed a generally applicable measure based on the trivially computable prospective and retrospective mutual information between CSs and reinforcements or between responses and reinforcements. Given two distinguishable event streams X and Y—for example a stream of CS onsets and a stream of reinforcements at CS termination—there is prospective mutual information between the x events and the y events, $I(\bar{X}; Y)$, when the expected wait to the next y, conditional on an x, $(\mu_y|\bar{x})$, is reliably shorter than the expected wait between the y's (μ_y) , as defined in Equation (1) [↗](#). There is retrospective mutual information, $I(X; \bar{Y})$, when the expected wait looking back from a y to the most recent x, $(\mu_x|\bar{y})$, is shorter than the expected wait between the x's (μ_x) , as in Equation (2) [↗](#):

$$I(\bar{X}; Y) = \log \frac{\mu_y}{\mu_y|\bar{x}} = \log \frac{\lambda_y|\bar{x}}{\lambda_y} \quad (1)$$

$$I(X; \bar{Y}) = \log \frac{\mu_x}{\mu_x|\bar{y}} = \log \frac{\lambda_x|\bar{y}}{\lambda_x} \quad (2)$$

The arguments of the log function in Equations (1) [↗](#) and (2) [↗](#) are the ratio of the *unconditional wait* (in the numerator) to the *conditional wait* (in the denominator), or equivalently, the ratio of the inverse of the waits, *conditional rate* in numerator and *unconditional rate* in denominator. Because its logarithm is the mutual information, Balsam et al. (2006) [↗](#) have termed this ratio the *informativeness* of a Pavlovian protocol.

Gallistel and Latham (2023) [↗](#) define *contingency* as the ratio of the *mutual information* to the *available information*, which is the amount that reduces subjective uncertainty to 0. Because measurement error scales with latency (Weber's Law, see Gibbon, 1977a [↗](#)), contingency equals

one only when the x 's and y 's coincide. How to measure the available information is often unclear. Mutual information, however, is trivially computed, as in [Equations \(1\)](#) and [\(2\)](#).

Another information-theoretic measure, nD_{KL} , measures the degree to which computed mutual information may be trusted. Its role, relative to mutual information, is analogous to the role of p values relative to correlations. The D_{KL} in the nD_{KL} is the Kullback- Leibler *divergence*. It is analogous to the effect size in conventional statistics. The effect size is the normalized *distance* between two distributions assumed to have the same variance. Distance is symmetric, i.e., $D(X,Y) = D(Y,X)$, but divergence is not: $D_{KL}(X|Y) \neq D_{KL}(Y|X)$. The asymmetric information-theoretic divergence is arguably the better measure, because the amount of data required to determine whether an X distribution differs from a Y distribution is not the same as the amount required to determine whether the opposite is true—and rats and mice are sensitive to this asymmetry (Kheifets, Freestone, & Gallistel, 2017; Kheifets & Gallistel, 2012).

The D_{KL} of one exponential distribution from another depends only on their rate parameters:

$$D_{KL}(X||Y)_{\text{exp}} = \ln \frac{\lambda_X}{\lambda_Y} + \frac{\lambda_Y}{X} - 1 \quad (3)$$

The uncertainty regarding the values of estimates for λ_X and λ_Y depends on the sample sizes, n_X and n_Y . Peter Latham (Gallistel & Latham, 2023, see their Appendix) has recently shown that when $\lambda_X = \lambda_Y$ (the null hypothesis),

$$n_e D_{KL}(X||Y) \sim \Gamma(n_p/2, 1), \quad (4)$$

where n_e is the *effective* sample size; $n_e = n_X(1 + n_X/n_Y)$; and n_p is the size of the parameter vector. For the exponential, $n_p = 1$. Therefore, as described earlier in [Equation \(2\)](#),

$$nD_{KL}(X||Y)_{\text{exp}} = \frac{n_X}{1+n_X/n_Y} \left(\ln \frac{\lambda_X}{\lambda_Y} + \frac{\lambda_Y}{\lambda_X} - 1 \right) \quad (5)$$

[Equation \(5\)](#) measures how “significant” an observed amount of mutual information is, how unlikely it is to have arisen by chance, and, by [Equation \(4\)](#), $nD_{KL}(X|Y)_{\text{exp}} \sim \Gamma(0.5,1)$, nD_{KL} 's, so it may be converted to the more familiar p values. Both are measures of the strength of statistical evidence. Conversion is motivated only by sociological considerations.

[Equations \(1\)](#), [\(2\)](#) and [\(5\)](#) and the RET equation $\hat{\lambda}_R = T^{-1}\lambda_R^I$ (Gallistel, 1990)—constitute a computationally simple, parameter-free model of associative learning. [Equation \(2\)](#) enables us to address the question, How sensitive are subjects to the strength of the evidence for observed mutual information? Put another way, Does their behavioural sensitivity to the accumulating mutual information suggest that they make a rational assessment of the extent to which they can trust the mutual information so far observed? **Figure 11** plots the cumulative distributions (CDFs) of the values for $nD_{KL}(\lambda_{R|CS} | \lambda_{R|C})$ —the nD_{KL} for the evidence that the CS rate of reinforcement, $\lambda_{R|CS}$, differs from the contextual rate of reinforcement, $\lambda_{R|C}$, when subjects' increased poke rates during CSs meets increasingly stringent evidentiary criteria for acquisition.

The medians of the blue CDFs in **Figure 11** generally fall at or below 3.2 nats worth of evidence, which corresponds to a p value of .01. This implies that they assess appropriately the strength of the accumulating evidence for reliable mutual information. The CDFs for the 10:1, 100:10, and 1000:1 evidentiary criteria cluster together well to the right, toward absurdly stringent criteria. (The odds against when the nD_{KL} is greater than 10 are more than 10,000:1.) The strong rightward shift in the strength of the stimulus evidence at acquisition occurs because the behavioural effect size is small and noisy (**Figure 11**). The low behavioural effect size requires larger n 's (more trials) to satisfy increasingly stringent evidentiary criteria (**Figure 3**). The Kullback-Leibler divergence of the CS reinforcement rate from the contextual rate is also small when the informativeness is low; when $\iota = 2$, it is only 0.2 nats. However, the effect of the small ratio of

the poke rates on the n 's required for increasingly stringent evidentiary criteria outweighs the small divergence in determining the amount of evidence for mutual information subjects have acquired after a given number of reinforcements.

The divergence is a decelerating function of informativeness, but its initial rise is steep: An informativeness of 10 yields a divergence of 1.5 nats. Thus, when a CS transmits 3.3 bits of mutual information per trial, 3 reinforced trials provide on average 4.5 nats of evidence, which corresponds to Odds of 370:1 in favour of the conclusion that the observed mutual information is reliable. This explains why CDFs for more stringent criteria in [Figure 11](#) migrate leftward as informativeness increases. When informativeness exceeds the value required for 1-trial acquisition in the median subject, the median subject satisfies the Odds 100:1 criterion when there is only 2.4 nats of evidence for the mutual information, which corresponds to Odds of 35:1 ($p < .03$). Again, this a rational evidentiary criterion.

Conclusion

Associative learning obeys simple equations that map from measurable properties of subjects' experience to measurable properties of their behaviour. The equations have at most one free parameter, and it is a scale factor with a data-anchored interpretation. The poke rate, for example, is 18 times the reinforcement rate in the median subject ([Figure 6](#)) over a range that covers 3 orders of magnitude. The variability about this regression is also multiplicative, with a scale factor of $10^{.39} = 2.5$. A second example is the simple relation between the learning rate (the reciprocal of reinforcements to acquisition) and the informativeness of a Pavlovian protocol. The log of the learning rate is -1 times the mutual information; the constant, k , in this regression is the informativeness that produces 1-trial acquisition ([Figures 1](#), [3](#), and [4](#)). A third example is the parameter free rate estimation equation. In a 1-CS protocol, that equation is:

$$\begin{bmatrix} \hat{\lambda}_{R|C\&CS} \\ \hat{\lambda}_{R|CS} \end{bmatrix} = \begin{bmatrix} 1 & 1/I \\ 1 & 1 \end{bmatrix}^{-1} \begin{bmatrix} \lambda_{R|C} \\ \lambda_{R|C\&CS} \end{bmatrix} \quad (6)$$

Equation (6) maps the *observed* contextual rate of reinforcement, $\lambda_{R|C}$, and the *observed* rate of reinforcement when the CS is also present, $\lambda_{R|C\&CS}$, to the rates of reinforcements that subjects *ascribe* to these predictors, $\hat{\lambda}_{R|C\&CS}$ and $\hat{\lambda}_{R|CS}$. It does so by way of the inverse of a simple matrix. The only element not equal to 1 in the matrix in **Equation (6)** is the inverse of the informativeness, which is also the protocol parameter in the learning-rate equation. **Equation (6)** explains the results in the cue-competition literature (a.k.a. the assignment-of-credit literature). (For review and comparison to the Rescorla-Wagner model, see [Gallistel, 1990](#), Chapter 13).

These equations bring out the formal structure of the *data* from associative learning experiments. They map an always defined and computable aspect of subjects' experience—*rate of reinforcement*—to measurable properties of subjects' behaviour. They contrast with equations that map from the often-undefinable *probability* of reinforcement to a *hypothetical construct* like associative strength or expected value ([Honey, Dwyer, & Iliescu, 2020](#); [Ludvig, Sutton, & Kehoe, 2012](#); [Pearce & Hall, 1980](#); [Rescorla & Wagner, 1972](#); [Sutton & Barto, 1990](#); [Vogel, Ponce, & Wagner, 2019](#)).

That the rates of reinforcement rather than the probabilities are the inputs to the equations that map perceived associations to behaviour calls into question the basic assumption in the common understanding of associative processes, the assumption that associative learning is mediated by the incrementing and decrementing of some scalar brain quantity (e.g., the conductivity of a plastic synapse) by reinforcement and non-reinforcement, respectively. This assumption is central to all models of Pavlovian and operant learning that map *probability* of reinforcement to associative strength and is central to all reinforcement learning models that map it to a weight




Figure 11.

Cumulative distributions of the nD_{KL} for the difference between the CS reinforcement rate and the contextual rate.

Note. The cumulative distributions are plotted as of the reinforcement after which the behavioural evidence for acquisition satisfied increasingly stringent criteria (Earliest evidence, and Odds of 10:1, 100:1 and 1,000:1 in favour of the trustworthiness of the observed mutual information). The informativeness (l) increases from left to right within a row of panels and top to bottom between panels.

state in a neural net (supervised and unsupervised learning) or to an expected long-term reward (reinforcement learning). They all fail to “leverage the statistical and computational structure of [the] problem” (Russo, Roy, Kazerouni, Osband, & Wen, 2018 [DOI](#), p. 6), because they must discretize time into unobserved trials or states with unspecified durations. They must do this because, unlike a rate that has a temporal unit, a probability is unitless.

A probability is the ratio between a count of reinforcements and the sum of the count of reinforcements and non-reinforcements. It is impossible to map from probability of reinforcement to behaviour in the general case because the non-reinforcements in the denominator are unobserved events with no physical properties, and, in the general case, often uncountable. They are countable only when reinforcement fails to occur at an expected time of reinforcement. In our protocol and many others the reinforcements occur at an unpredictable time or times following CS onset. The durations of our CS were drawn from a uniform distribution and the reinforcements coincided with its termination. In Rescorla’s (1967, 1968) experiments with truly random controls (and the many follow-ons), the reinforcements were programmed by Poisson processes. The defining feature of a Poisson process is its flat hazard function: there is no moment at which a reinforcement is any more likely than at any other moment.

Unlike models that take probability of reinforcement as the essential aspect of subjects’ experience, all of which have at least two free parameters, the model of associative learning provided by the simple equations we present here directly explains quantitative facts about associative learning that have gone unexplained for decades: One such fact is that reinforcements to acquisition are unaffected by partial reinforcement (Balsam & Gallistel, 2009 [DOI](#); Gallistel, 2003 [DOI](#); Gallistel, Craig, & Shahan, 2014 [DOI](#); Gibbon, Farrell, Locurto, Duncan, & Terrace, 1980 [DOI](#); Gottlieb, 2004 [DOI](#), 2005 [DOI](#)). This follows immediately from the equation for the learning rate as a function of informativeness (top of **Figure 1** [DOI](#)), because partial reinforcement does not alter informativeness. A second such fact is that deleting reinforced trials while retaining the spacing of the remaining reinforced trials does not alter the temporal progress of acquisition (Bouton & Sunsay, 2003 [DOI](#); Gottlieb, 2008 [DOI](#)). The number of reinforced trials in a given amount of training time is irrelevant, given only that there is at least one. This, too, follows from the equation that relates the learning rate to the informativeness or to its log, the mutual information (**Figures 1** [DOI](#) and **3** [DOI](#)). Because the slope on a loglog plot is essentially -1 over most of the useable range of values for informativeness, halving the number of trials doubles the informativeness and that doubles the learning rate.

We conclude that models of associative learning based on probability of reinforcement cannot be correct because the rate of responding, the learning rate and the assignment of credit are simple functions of rate of reinforcement, a quantity with temporal units. Only models that leverage the metric temporal structure of subject’s experience can be neurobiologically realisable.

Methods

Subjects

A total of 176 experimentally-naïve female albino Sprague Dawley rats (8 to 10 weeks of age) were obtained from Animal Resources Centre, Perth, Western Australia. They were housed in groups of 4 in split-level ventilated plastic tubs (TechniplastTM), measuring 40 x 46 x 40cm (length x width x height), located in an animal research facility at the University of Sydney. They had unrestricted access to water in their home tubs. Three days before commencing the experiment, they were placed on a restricted food schedule. Each day, half an hour after the end of the daily training session, each tub of rats received a ration of their regular dry chow (3.4 kcal/g) equal to 5% of the total weight of all rats in the tub. This amount is approximately equal to their required daily energy intake (Rogers, 1979 [DOI](#)), and took at least 2 h to be eaten (but was usually finished within 3

h). Rats on this schedule do not typically lose weight (and never more than 10%) but gain weight only very slowly. All experimental procedures were approved by the Animal Research Authority of the University of Sydney (protocol 2020/1840).

Apparatus

Rats were trained and tested in 32 Med Associates™ conditioning chambers distributed equally across four rooms. Twenty-four chambers (Set A) measured 28.5 x 30 x 25 cm (height x length x depth) and the other eight (Set B) were 21 x 30.5 x 24 cm (height x length x depth). Each chamber was individually enclosed in a sound- and light-resistant wooden shell (Set A) or PVC shell (Set B). The end walls of each chamber were made of aluminum; the sidewalls and ceiling were Plexiglas™. The floor 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 either in the center (Set A) or offset to the left of center (Set B). Attached to the food magazine was a dispenser delivering 45-mg food pellets (purified rodent pellets; Bioserve, Frenchtown, NJ). Illumination of an LED (Med Associates product ENV- 200RL-LED) mounted in the ceiling of the magazine served as the CS. Experimental events were controlled and recorded automatically by computers and relays located in the same room. Throughout all sessions, fans located in the rear wall of the outer shell provided ventilation and created background noise (between 61 and 66 dB, depending on the chamber).

Procedure

The rats were allocated to the 14 groups shown in [Table 1](#). The experiment was run with three separate cohorts of 64, 80, and 32 rats, so that by the end of the experiment each of the 14 groups had at least 12 rats (Groups 6 and 10 had 16 rats, as described next). In the first cohort, the data for 4 rats in each of Groups 6 and 10 were not saved in Session 1 due to human error. These 8 rats continued through the entire experiment but only their data for the final 5 sessions were used (to calculate their terminal response rate). An extra 4 rats were run in both groups in the second cohort.

The rats were not given magazine training. The experiment commenced with the first conditioning session and continued for 42 sessions over 42 consecutive days. Within each group, the CS-US interval varied from trial-to-trial according to a uniform distribution centred on the value of T for that group (the interval varied from a minimum of 2 s to a maximum of $2 \times T - 2$ s). The ITI in each group also varied from trial-to-trial as a uniform distribution with a minimum of 15 s. Between groups, T varied from 6 s to 62 s, and C varied from 63 s to 4,200 s (70 min) as summarised in [Table 1](#). The combinations of C and T gave rise to 14 distinct C/T ratios that were approximately evenly distributed on a log scale ranging from a ratio equal to 1.5 (63/42) up to 300 (4200/14). For the first 11 groups (those with ratios from 1.5 to 72), there were 10 trials per session; for the remaining three groups (with C/T ratios of 110, 180 and 300), each session contained only 3 trials so that the total session time remained within a manageable length (less than 4 h). All groups were trained with one session per day for a total of 42 sessions. Photo-beam interruptions by entry into the magazine were recorded during each CS and each ITI (recorded during the 20-s period immediately before CS onset).

Data analysis

Several different indices were used to identify when responding to the CS first appeared. The first index involved creating, for each rat, cumulative records of response counts during the CS and during the ITI, then converting these into cumulative records of response rates by dividing the cumulative count at each trial by the cumulative time (CS or ITI) at that trial. The primary index used here to establish when rats start responding to the CS is based on the difference between

these cumulative response rates (cumulative CS rate minus cumulative pre-CS ITI rate). The minimum value of this difference record identifies the trial after which the cumulative response rate during the CS becomes permanently greater than the cumulative ITI response rate. (This takes into account cases in which the CS response rate is initially lower than the ITI rate, wherein the difference in cumulative records decreases below zero but then reverses and starts to increase as soon as responding during the CS begins to outstrip ITI responding.) We used a one-sample t -test using to identify when this difference in the cumulative response rates became statistically reliable according to a significance threshold of $p < .05$. We have also used a new information-theoretic statistic, the nD_{KL} (see next section), to identify when the rate of responding during the CS began to exceed that during the ITI. Additional analyses were run adopting measures used in previous studies (Kirkpatrick & Church, 2000 [DOI](#); Thraill et al., 2020 [DOI](#)) to identify trials to a learning criterion. These analyses and results are described in the Supplementary Materials.

Analyses were also conducted to assess how T , C , and C/T affect conditioned responding after the point when it first emerges. These analyses examined how the response rate increased over trials and what level of responding was reached after extended training. The level of responding ultimately acquired by each group was computed as the mean response rate over the last 5 conditioning sessions. Further analyses broke down the response rate into separate components: the latency to first response; the mean duration of each response (time in the magazine); and the interval between responses (time out of the magazine). The increase in responding over sessions was assessed using a method described by Harris (2022) [DOI](#) that uses the slope of the cumulative response record to identify the number of trials required for responding to reach successive response milestones corresponding to each decile of the rat's peak response rate. Based on findings relating response rates to reinforcement rates ($1/T$) using this paradigm (Harris & Carpenter, 2011 [DOI](#)), we hypothesised that responding at the end of conditioning would be related to T , and indeed would scale linearly with $\log(T)$, but would not be related to C or C/T . We had no clear hypotheses about whether T , C , or C/T would affect how quickly responding increases after it has emerged.

The Kullback-Leibler divergence, and the nD_{KL}

Comparing the CS poke rate to the ITI poke rate reinforcement-by-reinforcement is problematic in cases where one of both rates are undefined because the subject has not yet made a poke. When a subject has made 5 pokes during the first 2 CSs and no pokes during the much longer ITIs, one does not want to conclude that the subject has not yet acquired a conditioned response to the CS. A second problem when using conventional statistics like the t -test is that one is required to specify in advance the sample size.

We circumvented these difficulties by reformulating the null hypothesis as a comparison between the contextual rate of responding (during the ITI) and the rate during the CS and by using an information-theoretic statistic to measure the strength of the evidence that the CS rate of responding differs from the contextual rate (Gallistel & Latham, 2023 [DOI](#)). The information theoretic statistic measures the strength of the evidence against the hypothesis that the distribution of inter-poke intervals during CSs is the same as the distribution in the context in which the CS occurs. Put another way, the null hypothesis is that poking during CSs does not differ from the poking expected because the subject pokes into the magazine in response to the fact that pellets sometimes drop there without regard to the signal value of the CS.

The CS poke rate, $\lambda_{r|CS}$, at any point in training is the number of pokes made during the CSs divided by their cumulative duration: $\lambda_{r|CS} = n_{r|CS}/D_{CS}$. And likewise for the contextual poke rate: $\lambda_{r|C} = n_{r|C}/D_C$. In these calculations, D_{CS} is the cumulative duration of the CS and D_C is cumulative training time. Assume for example that cumulative CS time as of the 2nd reinforcement is 20 s, total training time is 1000 s, and the subject has made 5 pokes during the two CS intervals and no pokes during the ITIs. Then, $\lambda_{r|CS} = 5/20 = .25/s$ and $\lambda_{r|C} = 5/1000 = .005/s$. The ratio of the two

estimates is $.25/.005 = 50:1$, that is, the subject's observed poke rate at this point in training is 50 times faster during the CSs than would be expected given the estimate of how frequently it pokes in the training context. This discrepancy is unlikely to have arisen by chance.

An information theoretic statistic, the nD_{KL} , measures the unlikeliness (Gallistel & Latham, 2023 [↗](#)). The DKL is the Kullback-Leibler divergence, a measure of the distance between two distributions. The divergence of one exponential distribution from another and depends only on the rate parameters of the distributions. Therefore, as per Equation (3) [↗](#), the divergence of $\lambda_{R|CS}$ from $\lambda_{R|C}$ is

$$D_{KL}(\lambda_{R|CS} || \lambda_{R|C})_{exp} = \log_e \left(\frac{\lambda_{R|CS}}{\lambda_{R|C}} \right) + \frac{\lambda_{R|C}}{\lambda_{R|CS}} - 1$$

It is the information-theoretic measure of the extent to which an exponential distribution with rate parameter $\lambda_{R|CS}$ diverges from an exponential distribution with rate parameter $\lambda_{R|C}$. The divergence is, roughly speaking, the equivalent of the effect size in a conventional analysis. The equivalence is rough because the effect size—the normalized distance between the means—is symmetric, whereas the divergence is not: $D_{KL}(\lambda_{R|CS} || \lambda_{R|C}) \neq D_{KL}(\lambda_{R|C} || \lambda_{R|CS})$.

The nD_{KL} measures the additional cost of encoding n data drawn from the exponential distribution with parameter $\lambda_{R|CS}$ on the assumption that they come from an exponential distribution with rate parameter $\lambda_{R|C}$ (Cover & Thomas, 1991 [↗](#)). It multiplies the DKL by the effective sample size, $n = n_{R|CS} / (1 + n_{R|CS} / n_{R|C})$. Therefore, as per Equation (5) [↗](#),

$$nD_{KL}(\lambda_{R|CS} || \lambda_{R|C}) = \frac{n_{R|CS}}{1 + n_{R|CS} / n_{R|C}} \left(\log_e \left(\frac{\lambda_{R|CS}}{\lambda_{R|C}} \right) + \frac{\lambda_{R|C}}{\lambda_{R|CS}} - 1 \right).$$

When there is no divergence, the nD_{KL} is distributed gamma(.5,1) (for proof, see Appendix in Gallistel & Latham, 2023 [↗](#)). Thus, we can convert the information-theoretic measure of the strength of the evidence for divergence to the more familiar p value measure.

In the illustrative example, $D_{KL}(.25 || .005) = 2.93$ nats (The nat is a unit of information equal to the base of the natural logarithm. The Kullback-Leibler divergence must be computed using the natural logarithm, but the result in nats may be converted to bits by multiplying by $\log_2 e = 1.44$.), and $n = 5/(1+5/5) = 2.5$. Therefore, the $nD_{KL} = 2.5 \times 2.93 = 7.32$ nats and $1 - \text{gamcdf}(7.32, .5, 1) \cong .0001$. The evidence against the null hypothesis is very strong; the odds against the assumption that all 5 pokes have occurred only during the CSs only by chance are on the order of 10,000 to 1. This example illustrates one approach to the statistical comparison of poke rates following each of the first few variable duration CSs, each terminating in a reinforcement.

Another approach to estimating the onset of conditioned poking comes from parsing the CS and ITI inter-poke interval vectors into segments with significantly different rate parameters. Parsing of the inter-poke interval vectors was also motivated by a desire to capture differences between subjects in the often bumpy evolution of their post-acquisition rate of responding.

Our parsing algorithm recursively extends the length, n_e , of the vector of inter-poke intervals one interval at a time. After each extension, it compares the rate estimate for each successively longer subsequence to the rate estimate for the full sequence, using the nD_{KL} statistic. These comparisons generate the function $nD_{KL}(n_s)$ for $n_s \leq n_e$. Whenever $\max[nD_{KL}(n_s)] > c$, the parser truncates the inter-poke interval vector at the location of the maximum. The value for the decision criterion, c , is user supplied and generally falls between 2 and 6 nats. The rate estimate for the segment truncated is the number of pokes up to the truncation divided by the sum of the intervals up to the truncation. The algorithm operates recursively on the post-truncation portions of the vector until there is no significant subsequence. The Matlab™ code for the custom `expparser.m` function is provided in the Supplementary Materials. The `ParseTable.xlsx` file in the Supplementary Material

gives results for values of 2, 4 and 6 nats. The corresponding p values are .05, .005 and .0005. Because parsing inescapably involves multiple comparisons, highly conservative decision criteria are generally to be preferred. When reporting results, we give only the results for $c = 6$ nats.

In computing the parses, we did not include the first CS poke in any given CS in our estimate of the CS poke rate because the latencies to the first pokes clearly came from a different distribution than the distribution of inter-poke intervals. In protocols with a short mean CS, the mean inter-poke interval (the reciprocal of the poke rate) was as short as 0.1 s; whereas the mean latency to the first poke in a CS was rarely shorter than 2 s. The explanation for the slow first pokes is probably the fact that the duration of a CS—hence the minimum reinforcement latency—was never shorter than 2 s. In the denominator of our poke rate estimates, we excluded the interval to the first poke. We also excluded the intervals when the head was in the hopper, because a poke cannot be made when the head is in the hopper. CS durations with no pokes were also excluded because there was no way to estimate the first-poke latency. Thus, our estimates of the poke rates included only CSs where there were pokes and the n 's in the numerators of those estimates were only the counts after each first poke.

In the denominator was the cumulative duration of the CSs that had at least one poke minus the cumulative latencies to the first pokes minus the cumulative duration of the hopper entries. The estimates of the poke rates during the Pre periods were the cumulative number of pokes in those intervals divided by their cumulative duration. The estimates of the contextual poke rates were the Pre poke rate estimate and the CS poke rate estimate weighted respectively by the cumulative ITI duration divided by the cumulative training duration and by the cumulative CS duration divided by the cumulative training duration.

The initial rate estimate in a parse extends back to the beginning of observation, which makes it possible to plot, for example, an initial ITI poke rate that starts at the beginning of training even though the first poke during a Pre interval may not have occurred until the 5th ITI. The parsing gives an alternative way of comparing rates of poking as of each successive reinforcement for the first few reinforcements. This becomes important when informativeness is high because then conditioned poking appears very early in training, as early as the second CS presentation.

Given the novelty of the methods used to estimate the onset of conditioned poking, we thought it essential to provide plots of the results on which these estimates are based. The supplementary material gives one 4-plot figure for each subject. **Figures 12** and **13** are examples of these plots. The bottom row of plots in **Figures 12** and **13** show the functions plotted over the entire 420 reinforcements (hence CS trials). In the top row, these complete plots have been right-cropped to better show what happens early in training.

In the left column of **Figures 12** and **13**, three cumulative poke-rate functions are plotted in black against the left axes. The poke rate as of a given reinforcement is the cumulative number of pokes as of that reinforcement divided by the cumulative duration of the relevant interval. In other words, it is the average rate as of a given reinforcement. The CS average rate is a solid line, the Pre average rate a dashed line. The contextual average rate is a dotted line.

In the left column of **Figure 12**, all three averages rise rapidly over the first 40 trials and then decline. The Pre average (dashed line) is generally greater than the CS average (solid line) until about the 110th trial, after which it is consistently lower. Both averages steadily decline up to about Trial 300, implying a persistent drop in both rates following their peak at around Trial 40. At about Trial 300, the average CS poke rate begins to rise rapidly, implying that the CS poke rate increased at the start of the rise. The slight rise in the dashed curve indicates a slight increase in the Pre poke rate as well. The successive changes in the poke rates inferred from these plots of the cumulative rate estimates are confirmed by the parses plotted in the right column of **Figures 12**, where the solid black line plots the parse of the CS poke rate and the dashed black line the parse of




Figure 12.

Note. In panels **a** & **c**, the rate of poking during the CS, denoted $\lambda_{r|CS}$ (solid black line), during pre-CS ITIs, $\lambda_{r|Pre}$ (dashed black line), and the contextual rate, $\lambda_{r|C}$ (dotted black line), plotted as functions of the number of reinforcements (1 R/trial). The red curve is the signed nD_{KL} plotted against the right axis. The x-axis in **a** has been right-cropped to better reveal early changes. Panels **b** & **d** show the parsed estimates of $\lambda_{r|CS}$ and $\lambda_{r|Pre}$. In all plots, the vertical red lines mark different estimates of when CS-conditional appeared (see text).




Figure 13.

Note. Second example: from a subject in the group with the maximally informative protocol. The average pokes/min plots for the CS in Panels a and c (solid black curves) are cumulative poke count divided by cumulative CS duration; they include the single pokes made during CS 3 and CS 5. These first pokes were ignored by the parse algorithm.

the Pre rate. The Pre poke rate started lower than the CS rate but jumped after about 10 reinforcements to a higher value. At about Trial 40, both rates became the same, but at about Trial 85, the Pre poke rate dropped permanently below the CS rate. At about Trial 300, there is a sequence of increases in the CS poke rate, accompanied by a single smaller increase in the Pre poke rate. The decision criterion used in this and all the plotted parses was 6 nats, so one may have substantial confidence that the changes are statistically significant.

The signed nD_{KL} function is plotted in red against the right axis in the left columns of **Figures 12** and **13**. The nD_{KL} is always positive because the magnitude of a divergence cannot be negative. The direction of a divergence may, however, vary; the Pre rate may be greater than or less than the CS rate. To make the direction of divergence apparent, we give the nD_{KL} positive sign when the CS rate is greater than the Pre rate and negative sign when the reverse is true. In the left column of **Figure 12**, there is a short initial positive spike in the signed nD_{KL} followed by an interval of several tens of reinforcements when it is negative. After about 40 reinforcements, it hits its minimum and began a more or less steady climb. Its last upward crossing of the 0 line is at about Trial 110. The upper limit on the nD_{KL} axis (right axis) is set at 6 nats, because the evidence for acquisition is decisive when that value is exceeded. (When $nD_{KL} = 6$, the odds are greater than 2000:1 against the null hypothesis.)

The vertical red lines on the four plots are drawn at the values for various estimates of reinforcements to acquisition. The solid heavy red line is drawn at our earliest estimate of the onset of CS-conditional responding. In **Figure 12**, this is at the minimum of the signed nD_{KL} , which is more often than not what we take to be the best estimate of where conditioned poking to the CS first appeared. The thinner solid red line is for a generally more conservative estimate of first appearance, namely, the reinforcement after which the signed nD_{KL} becomes permanently positive (last upward 0 crossing). The five vertical dashed red lines are for increasingly stringent evidentiary criteria: odds against the null of 4:1, 10:1, 20:1 ($p = .05$), 200:1 ($p = .005$) and 2000:1 ($p = .0005$).

The subject whose results are plotted in **Figure 13** made no pokes during the first 2 CSs, one during the 3rd CS, none during the 4th, one during the 5th and 16 on the 6th. It made no pokes during the first five 30s Pre intervals. Consequently, the average ITI poke rates and the average contextual poke rates are undefined over the first 5 trials and so is the nD_{KL} . On Trial 6, all 3 statistics become defined—and the nD_{KL} is already off scale, which is why it does not appear in **Figure 13a&c**. The nD_{KL} on Trial 6 is 20.7 nats, which corresponds to odds of 8 billion to 1 against the null. Consequently, the minimum in the signed nD_{KL} , the point where it became permanently positive, and the 5 the increasingly stringent evidentiary criteria for acquisition all fell at Trial 6, as indicated by the superposed red verticals at the right edge of **Figure 13a**. One might conclude that CS-conditional responding in this subject did not appear until Trial 5.

However, as already remarked, the first segment of a rate parse always extends back to the onset of observation. In the parses of the CS poke rate and the ITI poke rate, which are plotted in the right column of **Figure 13** found there is no evidence for a changes in the poke rates at Trial 6. Indeed, the parse of the CS poke rate shows no change over all 126 CSs, while the parse of the ITI poke rate finds the first change to be at Trial 13.

The failure to find changes in the poke rates at Trial 6 is not a consequence of the high value for the decision criterion. Lowering it from 6 nats to 4 did not change the parse; lowering it to 2 nats produced a parse with a single change, a drop from 52 pokes/min to 38 pokes/min at Trial 40. Likewise for the parse of the ITI rate of poking: lowering the decision criterion from 6 nats to 2 nats did not alter it. Nor does this failure occur because the algorithm cannot find very short segments. Parses of the contextual poke rate at all three values for the decision criterion find the

same 1-trial long segment at Trial 13—see upward blip in dashed plot in **Figure 13d** [↗](#). In other subjects, upward or downward blips 1 or 2 trials long are sometimes found at the outset of training.

There are significantly fewer pokes during the first 5 CSs than expected given the initial rate estimate for the CS poke rate. Much of this is attributable to the initially slow reaction to CS onset. The poke on Trial 3 came 20.3s into the 23s long CS; the poke on Trial 5 came 14.45s into the 16s long CS. The latency to make the first poke dropped rapidly over the first 9 trials from a mean greater than 11s for the first 5 trials to a mean of 3.6s for the Trials beyond 10. During the first 5 trials, the observation intervals during which it was possible to register a poke that would go into the parsing algorithm totaled only 4.18s. Given an initial poke rate estimate of 0.7/s, the expected number of pokes in that interval is 2.9 and there is a 5% chance of observing no pokes.

All considered, an argument can be made in this and several other cases where the informativeness was large that the parse results are a better indicator of the onset of CS-conditional poking, particularly when informativeness is high and the ITI poke rate very low. Blips in the ITI poke rate often put it momentarily greater than the CS poke rate (see dashed plot in **Figure 13d** [↗](#)) are common in post-acquisition protocols. Therefore, our parse-based estimate of the onset of conditioned is the trial after which the parsed CS poke rate is greater than the parsed ITI poke rate on 95% of the trials.

Columns 4, 5 and 6 of the Acquisition Table in the Supplementary Materials give the trial after which conditioned responding appeared as estimated in the above described three different ways — by the location of the minimum in the nD_{KL} , the last upward 0 crossing, and the CS parse consistently greater than the ITI parse, respectively. Column 3 in that table gives the minimum of the three estimates. Columns 7-11 of the Acquisition Table give the trial after which the evidence for acquisition exceeded increasingly stringent criteria (odds of 4, 10, 20, 100 and 1000:1). Four-panel figures, equivalent to **Figures 12** [↗](#) and **13** [↗](#), are included for each subject on pages 2 to 169 of the Supplementary Materials. The different estimates may coincide: In 81 subjects, the minimum of the nD_{KL} coincides with the earliest estimate; in 69 subjects the parse-based estimate does, and in 40 subjects the last upward 0 crossing does.

References

1. Balsam P. D., Fairhurst S., Gallistel C. R (2006) **Pavlovian contingencies and temporal information** *Journal of Experimental Psychology: Animal Behavior Processes* **32**:284–294
2. Balsam P. D., Gallistel C. R (2009) **Temporal maps and informativeness in associative learning** *Trends in Neurosciences* **32**:73–78
3. Belke T. W (1992) **Stimulus preference and the transitivity of preference** *Animal Learning and Behavior* **20**:401–406
4. Bouton M. E., Sunsay C (2003) **Importance of Trials Versus Accumulating Time Across Trials in Partially Reinforced Appetitive Conditioning** *Journal of Experimental Psychology: Animal Behaviour Processes* **29**:62–77
5. Burke D. A., Jeong H., Wu B., Lee S. A. F., Namboodiri V. M. K. (2023) **Few-shot learning: temporal scaling in behavioral and dopaminergic learning** *bioRxiv* <https://doi.org/10.1101/2023.03.31.535173>
6. Cover T. M., Thomas J. A (1991) **Information theory** New York: Wiley
7. Donahoe J. W (2006) **Contingency: Its meaning in the experimental analysis of behavior** *European Journal of Behavior Analysis* **7**:111–114
8. Drew M. R., Zupan B., Cooke A., Couvillon P. A., Balsam P (2005) **Temporal control of conditioned responding in goldfish** *Journal of Experimental Psychology: Animal Behavior Processes* **31**:31–39
9. Gallistel C. R (1990) **The organization of learning** Cambridge, MA: Bradford Books/MIT Press
10. Gallistel C. R (2003) **Conditioning from an information processing perspective** *Behavioural Processes* **62**:89–101
11. Gallistel C. R (2021) **Robert Rescorla: Time, Information and Contingency** *Revista de historia de la psicología* **42**:7–21
12. Gallistel C. R., Craig A. R., Shahan T. A (2014) **Temporal contingency** *Behavioural Processes* **101**:89–96 <https://doi.org/10.1016/j.beproc.2013.08.012>
13. Gallistel C. R., Fairhurst S., Balsam P (2004) **The learning curve: Implications of a quantitative analysis** *Proceedings of the National Academy of Sciences of the United States of America* **101**:13124–13131
14. Gallistel C. R., Gibbon J (2000) **Time, rate, and conditioning** *Psychological Review* **107**:289–344
15. Gallistel C. R., Latham P (2023) **Bringing Bayes and Shannon to the study of behavioral and neurobiological timing** *Timing & TIME Perception* **11**:29–89 <https://doi.org/10.1163/22134468-bja10069>

16. Gallistel C. R., Shahan T. A (2024) **Time-scale invariant contingency in reinforcement learning with extremely long delays to reinforcement**
17. Gibbon J (1977) **Scalar expectancy theory and Weber's Law in animal timing** *Psychological Review* **84**:279–335 <https://doi.org/10.1037/0033-295X.84.3.279>
18. Gibbon J (1977) **Scalar expectancy theory and Weber's law in animal timing** *Psychological Review* **84**:279–325
19. Gibbon J., Baldock M. D., Locurto C. M., Gold L., Terrace H. S (1977) **Trial and intertrial durations in autoshaping** *Journal of Experimental Psychology: Animal Behavior Processes* **3**:264–284
20. Gibbon J., Balsam P, Locurto C. M., Terrace H. S., Gibbon J. (1981) **Spreading association in time** *Autoshaping and conditioning theory* New York: Academic Press :219–253
21. Gibbon J., Berryman R., Thompson R. L (1974) **Contingency spaces and measures in classical and instrumental conditioning** *Journal of the Experimental Analysis of Behavior* **21**:585–605 <https://doi.org/10.1901/jeab.1974.21-585>
22. Gibbon J., Farrell L., Locurto C. M., Duncan H. J., Terrace H. S (1980) **Partial reinforcement in autoshaping with pigeons** *Animal Learning and Behavior* **8**:45–59
23. Gottlieb D. A (2004) **Acquisition with partial and continuous reinforcement in pigeon autoshaping** *Learning & Behavior* **32**:231–334
24. Gottlieb D. A (2005) **Acquisition with partial and continuous reinforcement in rat magazine approach** *Journal of Experimental Psychology: Animal Behavior Processes* **31**:319–333
25. Gottlieb D. A (2008) **Is the number of trials a primary determinant of conditioned responding?** *Journal of Experimental Psychology: Animal Behavior Processes* **34**:185–201
26. Granger R. H. J., Schlimmer J. C, Bower G. H. (1986) **The computation of contingency in classical conditioning** *The psychology of learning and motivation (Vol. 20* New York: Academic Press :137–192
27. 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
28. Hammond L. J., Paynter W. E (1983) **Probabilistic contingency theories of animal conditioning: A critical analysis** *Learning and Motivation* **14**:527–550 [https://doi.org/10.1016/0023-9690\(83\)90031-0](https://doi.org/10.1016/0023-9690(83)90031-0)
29. Harris J. A (2022) **The learning curve, revisited** *Journal of Experimental Psychology: Animal Learning and Cognition* **48**:265–280
30. Harris J. A., Carpenter J. S (2011) **Response rate and reinforcement rate in Pavlovian conditioning** *Journal of Experimental Psychology: Animal Behavior Processes* **37**:375–384
31. Harris J. A., Gharaei S., Pincham H. L (2011) **Response rates track the history of reinforcement times** *Journal of Experimental Psychology: Animal Behavior Processes* **37**:277–286

32. Holland P. C (2000) **Trial and intertrial durations in appetitive conditioning in rats** *Animal Learning and Behavior* **28**:121–135
33. Honey R. C., Dwyer D. M., Iliescu A. F (2020) **A model for Pavlovian learning and performance with reciprocal associations** *Psychological Review* **127**:829–852
34. Jenkins H. M., Barnes R. A., Barrera F. J., Locurto C. M., Terrace H. S., Gibbon J. (1981) **Why autoshaping depends on trial spacing** *Autoshaping and conditioning theory* New York: Academic Press :255–284
35. Kamin L. J., Jones M. R. (1968) **"Attention-like" processes in classical conditioning** *Miami symposium on the prediction of behavior: aversive stimulation* Miami: Miami University Press :9–31
36. Kheifets A., Freestone D., Gallistel C. R (2017) **Theoretical implications of quantitative properties of interval timing and probability estimation in mouse and rat** *Journal of the Experimental Analysis of Behavior* **108**:39–72 <https://doi.org/10.1002/jeab.261>
37. Kheifets A., Gallistel C. R (2012) **Mice take calculated risks** *Proceedings of the National Academy of Sciences of the United States of America* **109**:8776–8779 <https://doi.org/10.1073/pnas.1205131109>
38. Killeen P. R (2023) **Theory of reinforcement schedules** *Journal of the Experimental Analysis of Behavior* **120**:289–319 <https://doi.org/10.1002/jeab.880>
39. Killeen P. R., Hall S., Bizo L. A (1999) **A clock not wound runs down** *Behavioural Processes* **45**:129–139
40. Killeen P. R., Hanson S. J., Osborne S. R (1984) **Arousal: its genesis and manifestation as response rate** *Psychological Review* **85**:571–581 <https://doi.org/10.1037/0033-295X.85.6.571>
41. Kirkpatrick K., Church R. M (2000) **Independent effects of stimulus and cycle duration on conditioning: The role of timing processes** *Animal Learning and Behavior* **28**:373–388
42. Lattal K. M (1999) **Trial and intertrial durations in Pavlovian conditioning: Issues of learning and performance** *Journal of Experimental Psychology: Animal Behavior Processes* **25**
43. Ludvig E. A., Sutton R. S., Kehoe E. J (2012) **Evaluating the TD model of classical conditioning** *Learning & Behavior* **40**:305–319
44. Pearce J. M., Hall G (1980) **A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli** *Psychological Review* **87**:532–552
45. Rescorla R. A (1967) **Pavlovian conditioning and its proper control procedures** *Psychological Review* **74**:71–80
46. 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
47. Rescorla R. A., Wagner A. R, Black A. H., Prokasy W. F. (1972) **A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement** *Classical conditioning II: Current research and theory* New York: Appleton- Century-Crofts :64–99

- 48. Rogers A. E, Baker H. J., Lindsey J. R., Weisbroth S. H. (1979) **Nutrition** *The laboratory rat* New York: Academic Press :123–152
- 49. Russo D. J., Roy B. V., Kazerouni A., Osband I., Wen Z (2018) **A tutorial on Thompson Sampling** *Foundations and Trends in Machine Learning* **11**:1–96
- 50. Sutton R. S., Barto A. G, Gabriel M., Moore J. (1990) **Time-derivative models of Pavlovian reinforcement** *Learning and computational neuroscience: Foundations of adaptive networks* Cambridge, MA: Bradford Books/MIT Press :497–537
- 51. Thrailkill E. A., Todd T. P., Bouton M. E (2020) **Effects of conditioned stimulus (CS) duration, intertrial interval, and I/T ratio on appetitive Pavlovian conditioning** *Journal of Experimental Psychology: Animal Learning and Cognition* **46**:243–255
- 52. Vogel E. H., Ponce F. P., Wagner A. R (2019) **The development and present status of the SOP model of associative learning** *Quarterly Journal of Experimental Psychology* **72**:346–374 <https://doi.org/10.1177/1747021818777074>
- 53. Wagenmakers E.-J (2007) **A practical solution to the pervasive problems of p values** *Psychonomic Bulletin & Review* **14**:779–804
- 54. Ward R. D., Gallistel C. R., Jensen G., Richards V. L., Fairhurst S., Balsam P. D (2012) **Conditioned stimulus informativeness governs conditioned stimulus-unconditioned stimulus associability** *Journal of Experimental Psychology: Animal Behavior Processes* **38**:217–232

Author information

Justin A Harris

The University of Sydney, Sydney, Australia
ORCID iD: [0000-0003-3865-8097](https://orcid.org/0000-0003-3865-8097)

For correspondence: Justin.Harris@sydney.edu.au

CR Gallistel

Rutgers University, New Brunswick, United States
ORCID iD: [0000-0002-4860-5637](https://orcid.org/0000-0002-4860-5637)

Editors

Reviewing Editor

Nathan Holmes

UNSW Sydney, Sydney, Australia

Senior Editor

Kate Wassum

University of California, Los Angeles, Los Angeles, United States of America

Joint Public Review:

The subject area will have general appeal to those interested in the study of Pavlovian conditioning. The paper is important, showcasing a rigorous experimental design, several different approaches to data analysis, careful consideration of prior literature, and a thorough introduction. The results indicate that the rate of Pavlovian learning is determined by the ratio of reward rate during cue to the overall reward rate, and that the asymptotic response rate is determined by the reward rate during cue. These findings provide context to many conflicting recent results on this topic and are supported by strong/convincing evidence.

It is additionally claimed that the parameter that governs the acquisition and asymptote of responding in rats is exactly the same as that which governs the acquisition and asymptote of responding in the Gibbon and Balsam (1981) study that used pigeons as experimental subjects; and that the rates of responding during the inter-trial interval and the cue are proportional to the corresponding reward rates with the same proportionality constant. In both of these respects, there are several points that stand in need of clarification - at present, the strength of the evidence in support of these claims is solid. More generally, there are some points that could clarify aspects of rate estimation theory and, thereby, increase the rating of the paper from important to fundamental. These points range from analytical to conceptual and are presented below.

ANALYTICAL

(1) A key claim made here is that the same relationship (including the same parameter) describes data from pigeons by Gibbon and Balsam (1981; Figure 1) and the rats in this study (Figure 3). The evidence for this claim, as presented here, is not as strong as it could be. This is because the measure used for identifying trials to criterion in Figure 1 appears to differ from any of the criteria used in Figure 3, and the exact measure used for identifying trials to criterion influences the interpretation of Figure 3***. To make the claim that the quantitative relationship is one and the same in the Gibbon-Balsam and present datasets, one would need to use the same measure of learning on both datasets and show that the resultant plots are statistically indistinguishable, rather than simply plotting the dots from both data sets and spotlighting their visual similarity. In terms of their visual characteristics, it is worth noting that the plots are in log-log axis and, as such, slight visual changes can mean a big difference in actual numbers. For instance, between Figure 3B and 3C, the highest information group moves up only "slightly" on the y-axis but the difference is a factor of 5 in the real numbers. Thus, in order to support the strong claim that the quantitative relationships obtained in the Gibbon-Balsam and present datasets are identical, a more rigorous approach is needed for the comparisons.

***The measure of acquisition in Figure 3A is based on a previously established metric, whereas the measure in Figure 3B employs the relatively novel nDKL measure that is argued to be a better and theoretically based metric. Surprisingly, when r and r_2 values are converted to the same metric across analyses, it appears that this new metric (Figure 3B) does well but not as well as the approach in Figure 3A. This raises questions about why a theoretically derived measure might not be performing as well on this analysis, and whether the more effective measure is either more reliable or tapping into some aspect of the processes that underlie acquisition that is not accounted for by the nDKL metric.

(2) Another interesting claim here is that the rates of responding during ITI and the cue are proportional to the corresponding reward rates with the same proportionality constant. This too requires more quantification and conceptual explanation. For quantification, it would be more convincing to calculate the regression slope for the ITI data and the cue data separately and then show that the corresponding slopes are not statistically distinguishable from each

other. Conceptually, it is not clear why the data used to test the ITI proportionality came from the last 5 conditioning sessions. What were the decision criteria used to decide on averaging the final 5 sessions as terminal responses for the analyses in Figure 5? Was this based on consistency with previous work, or based on the greatest number of sessions where stable data for all animals could be extracted?

If the model is that animals produce response rates during the ITI (a period with no possible rewards) based on the overall rate of rewards in the context, wouldn't it be better to test this before the cue learning has occurred? Before cue learning, the animals would presumably only have attributed rewards in the context to the context and thus, produce overall response rates in proportion to the contextual reward rate. After cue learning, the animals could technically know that the rate of rewards during ITI is zero. Why wouldn't it be better to test the plotted relationship for ITI before cue learning has occurred? Further, based on Figure 1, it seems that the overall ITI response rate reduces considerably with cue learning. What is the expected ITI response rate prior to learning based on the authors' conceptual model? Why does this rate differ from pre and post-cue learning? Finally, if the authors' conceptual framework predicts that ITI response rate after cue learning should be proportional to contextual reward rate, why should the cue response rate be proportional to the cue reward rate instead of the cue reward rate plus the contextual reward rate?

(3) There is a disconnect between the gradual nature of learning shown in Figures 7 and 8 and the information-theoretic model proposed by the authors. To the extent that we understand the model, the animals should simply learn the association once the evidence crosses a threshold ($nDKL > \text{threshold}$) and then produce behavior in proportion to the expected reward rate. If so, why should there be a gradual component of learning as shown in these figures? In terms of the proportional response rule to the rate of rewards, why is it changing as animals go from 10% to 90% of peak response? The manuscript would be greatly strengthened if these results were explained within the authors' conceptual framework. If these results are not anticipated by the authors' conceptual framework, this should be explicitly stated in the manuscript.

(4) Page 27, Procedure, final sentence: The magazine responding during the ITI is defined as the 20 s period immediately before CS onset. The range of ITI values (Table 1) always starts as low as 15 s in all 14 groups. Even in the case of an ITI on a trial that was exactly 20 s, this would also mean that the start of this period overlaps with the termination of the CS from the previous trial and delivery (and presumably consumption) of a pellet. It should be indicated whether the definition of the ITI period was modified on trials where the preceding ITI was < 20 s, and if any other criteria were used to define the ITI. Were the rats exposed to the reinforcers/pellets in their home cage prior to acquisition?

(5) For all the analyses, the exact models that were fit and the software used should be provided. For example, it is not necessarily clear to the reader (particularly in the absence of degrees of freedom) that the model discussed in Figure 3 fits on the individual subject data points or the group medians. Similarly, in Figure 6 there is no indication of whether a single regression model was fit to all the plotted data or whether tests of different slopes for each of the conditions were compared. With regards to the statistics in Figure 6, depending on how this was run, it is also a potential problem that the analyses do not correct for the potentially highly correlated multiple measurements from the same subjects, i.e. each rat provides 4 data points which are very unlikely to be independent observations.

CONCEPTUAL

(1) We take the point that where traditional theories (e.g., Rescorla-Wagner) and rate estimation theory (RET) both explain some phenomenon, the explanation in terms of RET may be preferred as it will be grounded in aspects of an animal's experience rather than a hypothetical construct. However, like traditional theories, RET does not explain a range of

phenomena - notably, those that require some sort of expectancy/representation as part of their explanation. This being said, traditional theories have been incorporated within models that have the representational power to explain a broader array of phenomena, which makes me wonder: Can rate estimation be incorporated in models that have representational power; and, if so, what might this look like? Alternatively, do the authors intend to claim that expectancy and/or representation - which follow from probabilistic theories in the RW mould - are unnecessary for explanations of animal behaviour?***

***If the authors choose to reply to these points, they should consider taking advantage of an "Ideas and Speculation" subsection within the Discussion that is supported by eLife [<https://elifesciences.org/inside-elife/e3e52a93/elife-latest-including-ideas-and-speculation-in-elife-papers>].

(2) The discussion of Rescorla's (1967) and Kamin's (1968) findings needs some elaboration. These findings are already taken to mean that the target CS in each design is not informative about the occurrence of the US - hence, learning about this CS fails. In the case of blocking, we also know that changes in the rate of reinforcement across the shift from stage 1 to stage 2 of the protocol can produce unblocking. Perhaps more interesting from a rate estimation perspective, unblocking can also be achieved in a protocol that maintains the rate of reinforcement while varying the sensory properties of the US (Wagner). How does rate estimation theory account for these findings and/or the demonstrations of trans-reinforcer blocking (Pearce-Ganesan)? Are there other ways that the rate estimation account can be distinguished from traditional explanations of blocking and contingency effects? If so, these would be worth citing in the discussion. More generally, if one is going to highlight seminal findings (such as those by Rescorla and Kamin) that can be explained by rate estimation, it would be appropriate to acknowledge findings that challenge the theory - even if only to note that the theory, in its present form, is not all-encompassing. For example, it appears to me that the theory should not predict one-trial overshadowing or the overtraining reversal effect - both of which are amenable to discussion in terms of rates. I assume that the signature characteristics of latent inhibition and extinction would also pose a challenge to rate estimation theory, just as they pose a challenge to Rescorla-Wagner and other probability-based theories. Is this correct?

<https://doi.org/10.7554/eLife.102155.1.sa1>

Author response:

ANALYTICAL

(1) Figure 3 shows that the relationship between learning rate and informativeness for our rats was very similar to that shown with pigeons by Gibbon and Balsam (1981). We used multiple criteria to establish the number of trials to learn in our data, with the goal of demonstrating that the correspondence between the data sets was robust. To establish that they are effectively the same does require using an equivalent decision criterion for our data as was used for Gibbon and Balsam's data. However, the criterion they used—at least one peck at the response key on at least 3 out of 4 consecutive trials—cannot be sensibly applied to our magazine entry data because rats make magazine entries during the inter-trial interval (whereas pigeons do not peck at the response key in the inter-trial interval). Therefore, evidence for conditioning in our paradigm must involve comparison between the response rate during CS and the baseline response rate. There are two ways one could adapt the Gibbon and Balsam criterion to our data. One way is to use a non-parametric signed rank test for evidence that the CS response rate exceeds the pre-CS response rate, and adopting a statistical criterion equivalent to Gibbon and Balsam's 3-out-of-4 consecutive trials ($p < .3125$). The second method estimates the nDkl for the criterion used by Gibbon and Balsam. This could be done by assuming there are no responses in the inter-trial interval and a response

probability of at least 0.75 during the CS (their criterion). This would correspond to an nDkl of 2.2 (odds ratio 27:1). The obtained nDkl could then be applied to our data to identify when the distribution of CS response rates has diverged by an equivalent amount from the distribution of pre-CS response rates.

(2) A single regression line, as shown in Figure 6, is the simplest possible model of the relationship between response rate and reinforcement rate and it explains approximately 80% of the variance in response rate. Fixing the log-log slope at 1 yields the maximally simple model. (This regression is done in the logarithmic domain to satisfy the homoscedasticity assumption.) When transformed into the linear domain, this model assumes a truly scalar relation (linear, intercept at the origin) and assumes the same scale factor and the same scalar variability in response rates for both sets of data (ITI and CS). Our plot supports such a model. Its simplicity is its own motivation (Occam's razor).

If regression lines are fitted to the CS and ITI data separately, there is a small increase in explained variance ($R^2 = 0.82$). We leave it to further research to determine whether such a complex model, with 4 parameters, is required. However, we do not think the present data warrant comparing the simplest possible model, with one parameter, to any more complex model for the following reasons:

- When a brain—or any other machine—maps an observed (input) rate to a rate it produces (output rate), there is always an implicit scalar. In the special case where the produced rate equals the observed rate, the implicit scalar has value 1. Thus, there cannot be a simpler model than the one we propose, which is, in and of itself, interesting.
- The present case is an intuitively accessible example of why the MDL (Minimum Description Length) approach to model complexity (Barron, Rissanen, & Yu, 1998; Grünwald, Myung, & Pitt, 2005; Rissanen, 1999) can yield a very different conclusion from the conclusion reached using the Bayesian Information Criterion (BIC) approach. The MDL approach measures the complexity of a model when given N data specified with precision of B bits per datum by computing (or approximating) the sum of the maximum-likelihoods of the model's fits to all possible sets of N data with B precision per datum. The greater the sum over the maximum likelihoods, the more complex the model, that is, the greater its measured wiggle room, its capacity to fit data. Recall that von Neuman remarked to Fermi that with 4 parameters he could fit an elephant. His deeper point was that multi-parameter models bring neither insight nor predictive power; they explain only post-hoc, after one has adjusted their parameters in the light of the data. For realistic data sets like ours, the sums of maximum likelihoods are finite but astronomical. However, just as the Sterling approximation allows one to work with astronomical factorials, it has proved possible to develop readily computable approximations to these sums, which can be used to take model complexity into account when comparing models. Proponents of the MDL approach point out that the BIC is inadequate because models with the same number of parameters can have very different amounts of wiggle room. A standard illustration of this point is the contrast between logarithmic model and power-function model. Log regressions must be concave; whereas power function regressions can be concave, linear, or convex—yet they have the same number of parameters (one or two, depending on whether one counts the scale parameter that is always implicit). The MDL approach captures this difference in complexity because it measures wiggle room; the BIC approach does not, because it only counts parameters.
- In the present case, one is comparing a model with no pivot and no vertical displacement at the boundary between the black dots and the red dots (the 1-parameter unilinear model) to a bilinear model that allows both a change in slope and a vertical displacement for both lines. The 4-parameter model is superior if we use the BIC to take model complexity into account. However, 4-parameter has ludicrously more wiggle room. It will provide excellent fits—high maximum likelihood—to data sets in which the red points have slope > 1 , slope 0, or slope < 0 and in which it is also true that the intercept for the red points lies well below or well above

the black points (non-overlap in the marginal distribution of the red and black data). The 1-parameter model, on the other hand, will provide terrible fits to all such data (very low maximum likelihoods). Thus, we believe the BIC does not properly capture the immense actual difference in the complexity between the 1-parameter model (unilinear with slope 1) to the 4-parameter model (bilinear with neither the slope nor the intercept fixed in the linear domain).

· In any event, because the pivot (change in slope between black and red data sets), if any, is small and likewise for the displacement (vertical change), it suffices for now to know that the variance captured by the 1-parameter model is only marginally improved by adding three more parameters. Researchers using the properly corrected measured rate of head poking to measure the rate of reinforcement a subject expects can therefore assume that they have an approximately scalar measure of the subject's expectation. Given our data, they won't be far wrong even near the extremes of the values commonly used for rates of reinforcement. That is a major advance in current thinking, with strong implications for formal models of associative learning. It implies that the performance function that maps from the neurobiological realization of the subject's expectation is not an unknown function. On the contrary, it's the simplest possible function, the scalar function. That is a powerful constraint on brain-behavior linkage hypotheses, such as the many hypothesized relations between mesolimbic dopamine activity and the expectation that drives responding in Pavlovian conditioning (Berridge, 2012; Jeong et al., 2022; Y. Niv, Daw, Joel, & Dayan, 2007; Y. Niv & Schoenbaum, 2008).

The data in Figure 6 are taken from the last 5 sessions of training. The exact number of sessions was somewhat arbitrary but was chosen to meet two goals: (1) to capture asymptotic responding, which is why we restricted this to the end of the training, and (2) to obtain a sufficiently large sample of data to estimate reliably each rat's response rate. We have checked what the data look like using the last 10 sessions, and can confirm it makes very little difference to the results.

Finally, as noted by the reviews, the relationship between the contextual rate of reinforcement and ITI responding should also be evident if we had measured context responding prior to introducing the CS. However, there was no period in our experiment when rats were given unsignalled reinforcement (such as is done during "magazine training" in some experiments). Therefore, we could not measure responding based on contextual conditioning prior to the introduction of the CS. This is a question for future experiments that use an extended period of magazine training or "poor positive" protocols in which there are reinforcements during the ITIs as well as during the CSs. The learning rate equation has been shown to predict reinforcements to acquisition in the poor-positive case (Balsam, Fairhurst, & Gallistel, 2006).

(3) One of us (CRG) has earlier suggested that responding appears abruptly when the accumulated evidence that the CS reinforcement rate is greater than the contextual rate exceeds a decision threshold (C.R. Gallistel, Balsam, & Fairhurst, 2004). The new more extensive data require a more nuanced view. Evidence about the manner in which responding changes over the course of training is to some extent dependent on the analytic method used to track those changes. We presented two different approaches. The approach shown in Figures 7 and 8, extending on that developed by Harris (2022), assumes a monotonic increase in response rate and uses the slope of the cumulative response rate to identify when responding exceeds particular milestones (percentiles of the asymptotic response rate). This analysis suggests a steady rise in responding over trials. Within our theoretical model, this might reflect an increase in the animal's certainty about the CS reinforcement rate with accumulated evidence from each trial. While this method should be able to distinguish between a gradual change and a single abrupt change in responding (Harris, 2022) it may not distinguish between a gradual change and multiple step-like changes in responding and cannot account for decreases in response rate.

The other analytic method we used relies on the information theoretic measure of divergence, the nDkl (Gallistel & Latham, 2023), to identify each point of change (up or down) in the response record. With that method, we discern three trends. First, the onset tends to be abrupt in that the initial step up is often large (an increase in response rate by 50% or more of the difference between its initial value and its terminal value is common and there are instances where the initial step is to the terminal rate or higher). Second, there is marked within-subject variability in the response rate, characterised by large steps up and down in the parsed response rates following the initial step up, but this variability tends to decrease with further training (there tend to be fewer and smaller steps in both the ITI response rates and the CS response rate as training progresses). Third, the overall trend, seen most clearly when one averages across subjects within groups is to a moderately higher rate of responding later in training than after the initial rise. We think that the first tendency reflects an underlying decision process whose latency is controlled by diminishing uncertainty about the two reinforcement rates and hence about their ratio. We think that decreasing uncertainty about the true values of the estimated rates of reinforcement is also likely to be an important part of the explanation for the second tendency (decreasing within-subject variation in response rates). It is less clear whether diminishing uncertainty can explain the trend toward a somewhat greater difference in the two response rates as conditioning progresses. It is perhaps worth noting that the distribution of the estimates of the informativeness ratio is likely to be heavy tailed and have peculiar properties (as witness, for example, the distribution of the ratio of two gamma distributions with arbitrary shape and scale parameters) but we are unable at this time to propound an explanation of the third trend.

(4) There is an error in the description provided in the text. The pre-CS period used to measure the ITI responding was 10 s rather than 20 s. There was always at least a 5-s gap between the end of the previous trial and the start of the pre-CS period.

(5) Details about model fitting will be added in a revision. The question about fitting a single model or multiple models to the data in Figure 6 is addressed in response 2 above. In Figure 6, each rat provides 2 behavioural data points (ITI response rate and CS response rate) and 2 values for reinforcement rate ($1/C$ and $1/T$). There is a weak but significant correlation between the ITI and CS response rates ($r = 0.28$, $p < 0.01$; log transformed to correct for heteroscedasticity). By design, there is no correlation between the log reinforcement rates ($r = 0.06$, $p = .404$).

CONCEPTUAL

(1) It is important for the field to realize that the RW model cannot be used to explain the results of Rescorla's (Rescorla, 1966; Rescorla, 1968, 1969) contingency-not-pairing experiments, despite what was claimed by Rescorla and Wagner (Rescorla & Wagner, 1972; Wagner & Rescorla, 1972) and has subsequently been claimed in many modelling papers and in most textbooks and reviews (Dayan & Niv, 2008; Y. Niv & Montague, 2008). Rescorla programmed reinforcements with a Poisson process. The defining property of a Poisson process is its flat hazard function; the reinforcements were equally likely at every moment in time when the process was running. This makes it impossible to say when non-reinforcements occurred and, a fortiori, to count them. The non-reinforcements are causal events in RW algorithm and subsequent versions of it. Their effects on associative strength are essential to the explanations proffered by these models. Non-reinforcements—failures to occur, updates when reinforcement is set to 0, hence also the lambda parameter—can have causal efficacy only when the successes may be predicted to occur at specified times (during “trials”). When reinforcements are programmed by a Poisson process, there are no such times. Attempts to apply the RW formula to reinforcement learning soon foundered on this problem (Gibbon, 1981; Gibbon, Berryman, & Thompson, 1974; Hallam, Grahame, & Miller, 1992; L.J. Hammond, 1980; L. J. Hammond & Paynter, 1983; Scott & Platt, 1985). The enduring popularity of the delta-rule updating equation in reinforcement learning depends on “big-

concept” papers that don’t fit models to real data and discretize time into states while claiming to be real-time models (Y. Niv, 2009; Y. Niv, Daw, & Dayan, 2005).

The information-theoretic approach to associative learning, which sometimes historically travels as RET (rate estimation theory), is unabashedly and inescapably representational. It assumes a temporal map and arithmetic machinery capable in principle of implementing any implementable computation. In short, it assumes a Turing-complete brain. It assumes that whatever the material basis of memory may be, it must make sense to ask of it how many bits can be stored in a given volume of material. This question is seldom posed in associative models of learning, nor by neurobiologists committed to the hypothesis that the Hebbian synapse is the material basis of memory. Many—including the new Nobelist, Geoffrey Hinton—would agree that the question makes no sense. When you assume that brains learn by rewiring themselves rather than by acquiring and storing information, it makes no sense.

When a subject learns a rate of reinforcement, it bases its behavior on that expectation, and it alters its behavior when that expectation is disappointed. Subjects also learn probabilities when they are defined. They base some aspects of their behavior on those expectations, making computationally sophisticated use of their representation of the uncertainties (Balci, Freestone, & Gallistel, 2009; Chan & Harris, 2019; J. A. Harris, 2019; J.A. Harris & Andrew, 2017; J. A. Harris & Bouton, 2020; J. A. Harris, Kwok, & Gottlieb, 2019; Kheifets, Freestone, & Gallistel, 2017; Kheifets & Gallistel, 2012; Mallea, Schulhof, Gallistel, & Balsam, 2024 in press).

(2) Rate estimation theory is oblivious to the temporal order in which experience with different predictors occurs. The matrix computation finds the additive solution, if it exists, to the data so far observed, on the assumption that predicted rates have remained the same. This is the stationarity assumption, which is implicit in a rate computation and was made explicit in the formulation of RET (C.R. Gallistel, 1990). When the additive solution does not exist, the RET algorithm treats the compound of two predictors as a third predictor, and computes the additive solution to the 3-predictor problem. Because it is oblivious to the order in which the data have been acquired, it predicts one-trial overshadowing and retroactive blocking and unblocking (C.R. Gallistel, 1990 pp 439 & 452-455).

The RET algorithm is but one component of the information-theoretic model of associative learning (aka, TATAL, The Analytic Theory of Associative Learning Wilkes & Gallistel, 2016)). It solves the assignment-of-credit problem, not the change-detection problem. Because rates of reinforcement do sometimes change, the stationarity assumption, which is essential to the RET algorithm, must be tested when each new reinforcement occurs and when the interval since the last reinforcement has become longer than would be expected or the number of reinforcements has become significantly fewer than would be expected given the current estimate of the probability of reinforcement (C. R. Gallistel, Krishan, Liu, Miller, & Latham, 2014). In the information-theoretic approach to associative learning, detecting non-stationarity is done by an information-theoretic change-detecting algorithm. The algorithm correctly predicts that omitted reinforcements to extinction will be a constant (C.R. Gallistel, 2024 under review; Gibbon, Farrell, Locurto, Duncan, & Terrace, 1980). To put the prediction another way, unreinforced trials to extinction will increase in proportional to the trials/reinforcement during training (C.R. Gallistel, 2012; Wilkes & Gallistel, 2016). In other words, it predicts the best and most systematic data on the partial reinforcement extinction effect (PREE) known to us. The profound challenge to neo-Hullian delta-rule updating models that is posed by the PREE has been recognized for the better part of a century. To the best of our knowledge, no other formalized model of associative learning has overcome this challenge (Dayan & Niv, 2008; Mellgren, 2012). Explaining extinction algorithmically is straightforward when one adopts an information-theoretic perspective, because computing reinforcement-by-reinforcement the Kullback-Leibler divergence in a sequence of earlier rate (or probability!) estimates from the most recent estimate and multiplying the vector of divergences by the vector of effective sample sizes (C. R. Gallistel & Latham, 2022) detects and

localized changes in rates and probabilities of reinforcement (C.R. Gallistel, 2024 under review). The computation presupposes the existence of a temporal map, a time-stamped record of past events. This supposition is strongly resisted by neuroscience-oriented reinforcement-learning modelers, who try to substitute the assumption of decaying eligibility traces.

The very interesting Pearce-Ganesan findings (Ganesan & Pearce, 1988) are not predicted by RET, but nor do they run counter its predictions. RET has nothing to say about how subjects categorize appetitive reinforcements; nor, at this time, does the information-theoretic approach to an understanding of associative have anything to say about that.

The same is not true for the Betts, Brandon & Wagner results (Betts, Brandon, & Wagner, 1996). They pretrained a blocking cue that predicted a painful paraorbital shock to one eye of a rabbit. This cue elicited an anticipatory blink in the threatened eye. It also potentiated the startle reflex made to a loud noise in one ear. A new cue that was then introduced, which always occurred in compound with the pretrained blocking cue. In one group, the painful shock continued to be delivered to the same eye as before; in another group, it was delivered to the skin around the other eye. In the group that continued to receive the shock to the same eye, the old cue effectively blocked conditioning of the new cue for both the eyeblink and the potentiated startle response. However, in the group for which the location of the shock changed to the other eye, the old cue did not block conditioning of the eyeblink response to the new cue but did block conditioning of the startle response to the new cue. The information-theoretic analysis of associative learning focusses on the encoding of measurable predictive temporal relationships, rather than on general and, to our mind, vague notions like CS processing and US processing. A painful shock elicits fear in a rabbit no matter where on the body surface it is experienced, because fear is a reaction to a very broad category of dangers, and fear potentiates the startle reflex regardless of the threat that causes fear. Once that prediction of such a threat is encoded; redundant cues will not be encoded that same way because the RET algorithm blocks the encoding of redundant predictions. A painful shock near an eye elicits a blink of the threatened eye as well as the fear that potentiates the startle. An appropriate encoding for the eye blink must specify the location of the threat. RET will attribute prediction of the threat to the new eye to the new cue—and not to the old cue, the pretrained blocker—while continuing to attribute to the old cue the prediction of a fear-causing threat, because the change in location does not alter that prediction. Therefore, the new cue will be encoded as predicting the new location of the threat to the eye, but not as predicting the large category non-specific threats that elicit fear and the potentiation of the startle, because that prediction remains valid. Changing that prediction would violate the stationarity assumption; predictive relations do not change unless the data imply that they must have changed. Unless we have made a slip in our logic, this would seem to explain Betts et al's (1996) results. It does so with no free parameters, unlike AESOP, which has a notoriously large number of free parameters.

Balci, F., Freestone, D., & Gallistel, C. R. (2009). Risk assessment in man and mouse. *Proceedings of the National Academy of Science U S A*, 106(7), 2459-2463. doi:10.1073/pnas.0812709106

Balsam, P. D., Fairhurst, S., & Gallistel, C. R. (2006). Pavlovian contingencies and temporal information. *Journal of Experimental Psychology: Animal Behavior Processes*, 32, 284-294.

Barron, A., Rissanen, J., & Yu, B. (1998). The minimum description length principle in coding and modeling. *IEEE Transactions on Information Theory*, 44(6), 2743-2760.

Berridge, K. C. (2012). From prediction error to incentive salience: Mesolimbic computation of reward motivation. *European Journal of Neuroscience*.

- Betts, S. L., Brandon, S. E., & Wagner, A. R. (1996). Dissociation of the blocking of conditioned eyeblink and conditioned fear following a shift in US locus. *Animal Learning and Behavior*, 24(4), 459-470.
- Chan, C. K. J., & Harris, J. A. (2019). The partial reinforcement extinction effect: The proportion of trials reinforced during conditioning predicts the number of trials to extinction. *Journal of Experimental Psychology: Animal Learning and Cognition*, 45(1). doi:<http://dx.doi.org/10.1037/xan0000190>
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, 18(2), 185-196.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: Bradford Books/MIT Press.
- Gallistel, C. R. (2012). Extinction from a rationalist perspective. *Behav Processes*, 90, 66-88. doi:10.1016/j.beproc.2012.02.008
- Gallistel, C. R. (2024 under review). Reconceptualized associative learning. *Perspectives on Behavioral Science (Special Issue for SQAB 2024)*.
- Gallistel, C. R., Balsam, P. D., & Fairhurst, S. (2004). The learning curve: Implications of a quantitative analysis. *Proceedings of the National Academy of Sciences*, 101(36), 13124-13131.
- Gallistel, C. R., Krishan, M., Liu, Y., Miller, R. R., & Latham, P. E. (2014). The perception of probability. *Psychological Review*, 121, 96-123. doi:10.1037/a0035232
- Gallistel, C. R., & Latham, P. E. (2022). Bringing Bayes and Shannon to the Study of Behavioral and Neurobiological Timing. Timing & Time Perception. *timing & TIME Perception*, 1-61. doi:10.1163/22134468-bja10069
- Ganesan, R., & Pearce, J. M. (1988). Effect of changing the unconditioned stimulus on appetitive blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, 14, 280-291.
- Gibbon, J. (1981). The contingency problem in autoshaping. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 285-308). New York: Academic.
- Gibbon, J., & Balsam, P. (1981). Spreading association in time. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 219-253). New York: Academic Press.
- Gibbon, J., Berryman, R., & Thompson, R. L. (1974). Contingency spaces and measures in classical and instrumental conditioning. *Journal of the Experimental Analysis of Behavior*, 21(3), 585-605. doi: 10.1901/jeab.1974.21-585
- Gibbon, J., Farrell, L., Locurto, C. M., Duncan, H. J., & Terrace, H. S. (1980). Partial reinforcement in autoshaping with pigeons. *Animal Learning and Behavior*, 8, 45-59. doi:[doi:10.3758/BF03209729](http://dx.doi.org/10.3758/BF03209729)
- Grünwald, P. D., Myung, I. J., & Pitt, M. A. (2005). *Advances in minimum description length: theory and applications*. Cambridge, MA: MIT Press.
- 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.
- Hammond, L. J. (1980). The effect of contingency upon the appetitive conditioning of free operant behavior. *Journal of the Experimental Analysis of Behavior*, 34, 297-304.

doi:10.1901/jeab.1980.34-297

Hammond, L. J., & Paynter, W. E. (1983). Probabilistic contingency theories of animal conditioning: A critical analysis. *Learning and Motivation*, 14, 527-550. doi:10.1016/0023-9690(83)90031-0

Harris, J. A. (2019). The importance of trials. *Journal of Experimental Psychology: Animal Learning and Cognition*, 45(4).

Harris, J. A. (2022). The learning curve, revisited. *Journal of Experimental Psychology: Animal Learning and Cognition*, 48, 265-280.

Harris, J. A., & Andrew, B. J. (2017). Time, Trials and Extinction. *Journal of Experimental Psychology: Animal Learning and Cognition*, 43(1), 15-29.

Harris, J. A., & Bouton, M. E. (2020). Pavlovian conditioning under partial reinforcement: The effects of non-reinforced trials versus cumulative CS duration. *The Journal of Experimental Psychology: Animal Learning & Cognition*, 46, 256-272.

Harris, J. A., Kwok, D. W. S., & Gottlieb, D. A. (2019). The partial reinforcement extinction effect depends on learning about nonreinforced trials rather than reinforcement rate. *Journal of Experimental Psychology: Animal Behavior Learning and Cognition*, 45(4). doi:10.1037/xan0000220

Jeong, H., Taylor, A., Floeder, J. R., Lohmann, M., Mihalas, S., Wu, B., . . . Namboodiri, V. M. K. (2022). Mesolimbic dopamine release conveys causal associations. *Science*. doi:10.1126/science.abq6740

Kheifets, A., Freestone, D., & Gallistel, C. R. (2017). Theoretical Implications of Quantitative Properties of Interval Timing and Probability Estimation in Mouse and Rat. *Journal of the Experimental Analysis of Behavior*, 108(1), 39-72. doi:doi.org/10.1002/jeab.261

Kheifets, A., & Gallistel, C. R. (2012). Mice take calculated risks. *Proceedings of the National Academy of Science*, 109, 8776-8779. doi:doi.org/10.1073/pnas.1205131109

Mallea, J., Schulhof, A., Gallistel, C. R., & Balsam, P. D. (2024 in press). Both probability and rate of reinforcement can affect the acquisition and maintenance of conditioned responses. *Journal of Experimental Psychology: Animal Learning and Cognition*.

Mellgren, R. (2012). Partial reinforcement extinction effect. In N. M. Seel (Ed.), *Encyclopedia of the Sciences of Learning*. Boston, MA: Springer.

Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53, 139-154.

Niv, Y., Daw, N. D., & Dayan, P. (2005). How fast to work: response vigor, motivation and tonic dopamine. In Y. Weiss, B. Schölkopf, & J. R. Platt (Eds.), *NIPS 18* (pp. 1019-1026). Cambridge, MA: MIT Press.

Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology*, 191(3), 507-520.

Niv, Y., & Montague, P. R. (2008). Theoretical and empirical studies of learning. In (., eds), pp. , Academic Press. In P. W. e. a. Glimcher (Ed.), *Neuroeconomics: Decision-Making and the Brain* (pp. 329-349). New York: Academic Press.

Niv, Y., & Schoenbaum, G. (2008). Dialogues on prediction errors. *Trends in Cognitive Sciences*, 12(7), 265-272. doi:10.1016/j.tics.2008.03.006

Rescorla, R. A. (1966). Predictability and the number of pairings in Pavlovian fear conditioning. *Psychonomic Science*, 4, 383-384.

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), 1-5. doi:10.1037/h0025984

Rescorla, R. A. (1969). Conditioned inhibition of fear resulting from negative CS-US contingencies. *Journal of Comparative and Physiological Psychology*, 67, 504-509.

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* (pp. 64-99). New York: Appleton-Century-Crofts.

Rissanen, J. (1999). Hypothesis selection and testing by the MDL principle. *The Computer Journal*, 42, 260-269. doi:10.1093/comjnl/42.4.260

Scott, G. K., & Platt, J. R. (1985). Model of response-reinforcement contingency. *Journal of Experimental Psychology: Animal Behavior Processes*, 11(2), 152-171.

Wagner, A. R., & Rescorla, R. A. (1972). Inhibition in Pavlovian conditioning: Application of a theory. In R. A. Boakes & S. Halliday (Eds.), *Inhibition and learning*. New York: Academic.

Wilkes, J. T., & Gallistel, C. R. (2016). *Information Theory, Memory, Prediction, and Timing in Associative Learning (original long version)*.

<https://doi.org/10.7554/eLife.102155.1.sa0>