

Changes in the distribution of response rates across the CS-US interval: Evidence that responding switches between two distinct states

Justin A. Harris

University of Sydney, Australia

Two experiments used the peak procedure to examine timing of conditioned responses in a magazine approach paradigm with rats. A conditioned stimulus (CS) was reinforced with food on 50% of trials. Food was delivered at a fixed time, either 20 s, 30 s or 40 s into the CS presentation. Response rates were recorded during non-reinforced CS presentations that extended well beyond the scheduled time of food delivery. The mean response rate (averaged over many trials) increased during the CS, peaking at the expected time of reinforcement, and decreased again. Detailed analyses of the frequency distribution of response rates showed that responding was described by two distinct distributions, consistent with the rat being in a low response state on some trials and in a high response state on other trials. Modeling of these frequency distributions showed that the systematic rise and fall in response rate across a trial was primarily explained by a change in the proportion of time that the rat spent in the low versus high response state. However, the change in responding was also explained in part by a continuous shift in the high response state, such that responding in that state increased and then decreased gradually across the trial. These results support accounts that describe response timing as an abrupt change from low to high responding during the CS, but also provide evidence for a continuous change in conditioning strength across the duration of the CS. The implications of these findings for timing and associative theories of conditioning are discussed.

Key words: Rat; fixed interval; variable interval; magazine approach; mixture model

Address correspondence to:

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

Author note: The work was supported by grants DP1092695 and FT100100091 from the Australian Research Council. The author thanks Peter Balsam for discussions on some of the analyses presented here.

Introduction

In conditioning preparations, when the interval between a conditioned stimulus (CS) and unconditioned stimulus (US) is fixed, the conditioned response (CR) shows evidence of timing, with the peak of responding coinciding with the time of arrival of the US. This timing of CRs was described by Pavlov (1927) and has been reported in many conditioning paradigms across many different species including rats, rabbits, pigeons and fish (e.g., Davis, Schlesinger, & Sorenson, 1989; Drew, Zupan, Cooke, Couvillon, & Balsam, 2005; Kehoe & Joscellyne, 2005; W. A. Roberts, Cheng, & Cohen, 1989; Smith, 1968; Williams, Lawson, Cook, Mather, & Johns, 2008). Timing of CRs is most clearly revealed using the peak procedure. This procedure intermixes reinforced trials, in which the CS is followed by the US after a fixed interval, with non-reinforced trials in which the CS is presented for an extended period, longer than the CS-US interval on reinforced trials. When response rates are tracked during the long non-reinforced trials, CRs initially increase in strength or frequency as elapsed time approaches the expected time of the US, and then decrease again as time extends beyond that point (e.g., Church, Meck, & Gibbon, 1994; S. Roberts, 1981; W. A. Roberts et al., 1989).

Pavlov (1927) proposed that CR timing involves the acquisition of inhibition to the onset of the CS which prevents the animal from responding early in the CS-US interval, but the CR gradually emerges as the inhibition diminishes during the trial. While this account does provide an explanation for the delay in responding prior to the usual time of US delivery, it does not explain why responding characteristically declines again as the CS extends beyond that time. More recent accounts have proposed that information about the CS-US interval itself is represented within the conditioning process. One class of theories treats the CS as a sequence of units or states that are distributed sequentially in time; the relationship between each unit and the US (or CR) is represented by a

continuous one-dimensional variable (e.g., Desmond & Moore, 1988; Machado, 1997; Sutton & Barto, 1981, 1990). In effect, the CS-US association is a vector of values describing the strength of an association at each moment in time. The CS units that coincide closely with the time of the US acquire the greatest strength, and other units acquire strength inversely proportional to their distance from the US. These more remote units acquire strength either because the temporal distribution of their activity overlaps with the time of the US or CR (Desmond & Moore, 1988; Machado, 1997), or by a propagation process that spreads associative strength between neighboring CS units (Sutton & Barto, 1981, 1990). In short, the changing pattern of responding within a trial is taken to reflect a continuous change in the recruitment of associative strength, which initially grows across time units during the CS, peaks at the expected time of the US, and then falls again.

Another, very different, theoretical approach has been developed specifically around the timing of CRs. According to these theories, conditioning does not rely on graded changes in a continuous variable such as associative strength. Rather, when an animal learns about the relationship between CS and US, it is encoding the temporal interval between the CS and US (e.g., Gallistel & Gibbon, 2000; Gibbon, 1977; Guilhardi, Yi, & Church, 2007). On subsequent trials, the remembered CS-US interval is compared against the estimate of currently elapsed time, and the animal begins to respond when the ratio of these values reaches a threshold. According to this account, responding emerges abruptly at a decision point during the trial. The gradual rise in the mean response rate, reported in conditioning experiments that use a fixed CS-US interval, is argued to be an artifact of averaging over many trials (Schneider, 1969). It is assumed that the decision to respond does not occur at exactly the same time on each trial due to noise in any of the various components that go into the decision (e.g., variability in the memory of the CS-US

interval, variability in measurement of currently elapsed time, variability in the decision threshold itself). Therefore, response rates on individual trials take the form of a step function, transitioning abruptly from low to high when the decision threshold is reached, but averaging over many such step functions produces a continuously graded function that rises smoothly over the range of decision times (Gibbon & Church, 1990).

Evidence supporting an abrupt transition between two response states has come from break-run-break analyses that search for transition points in the response record of individual trials. For example, in a peak procedure, each non-reinforced trial is analysed to identify two time points, such that between those times the response rate (during the run) is maximally different from the response rate during the periods before and after those two time points (during the breaks, Cheng & Westwood, 1993; Church et al., 1994). When all trials are aligned by either of these time points, the mean response rates show a sharp discontinuity either side of that point, transitioning abruptly between a low but relatively steady rate and a high steady rate. This pattern of response change is clearly consistent with the argument that response rates within a trial change abruptly at a decision point as the expected time of reinforcement approaches. However, algorithms that search for break points in response records can be sensitive to random fluctuations in response rates within a trial, with the consequence that they produce artifactual evidence for discontinuity even when the underlying process generating the response data is continuous (Harris, 2011). Indeed, as described below, this problem exists for algorithms used to reveal break-run-break patterns in response records from individual trials.

To test the specificity of the algorithm used to reveal a break-run-break pattern, two different sets of simulated data were generated (see Appendix for details). One set was created to mimic a true low-high-low discontinuous response pattern, with the probability of responses transitioning abruptly from low to high and back to low at variable time points within each trial. The other set was created with a continuous change in the probability of responding across each trial. Figure 1A shows response rates at consecutive time points within the trial, average over 1000 simulated trials, for both the continuous data set and the discontinuous (2-state) data set. It is clear that, for both data sets, the averaged response rates rise and fall smoothly, with a peak at the middle of the trial. This demonstrates how averaging across trials will make a discontinuous response record appear smooth.

Both simulated data sets were then analyzed using an algorithm designed to find two transition points, t_1 and t_2 , in each trial (Cheng & Westwood, 1993; Church et al., 1994, see Appendix for details). Figure 1 shows the mean response rates, averaged across all trials, after trials were aligned by t_1 (Figure 1B) or by t_2 (Figure 1C). For the 2-state data, there is a very abrupt transition from a low rate to a high rate at t_1 , and an equally abrupt transition from high to low at t_2 , with response rates being steady either side of the transition points. This confirms that the method is sensitive in identifying change points in data that contain abrupt transitions in response probability. However, a very similar pattern, albeit with a more modest transition, is also extracted from the continuous data set (see Figure 1). This shows that the algorithm used to find transitions between low and high response states is not necessarily specific, and can produce false positive evidence for discontinuity in data where the change in probability of responding is smooth and continuous.

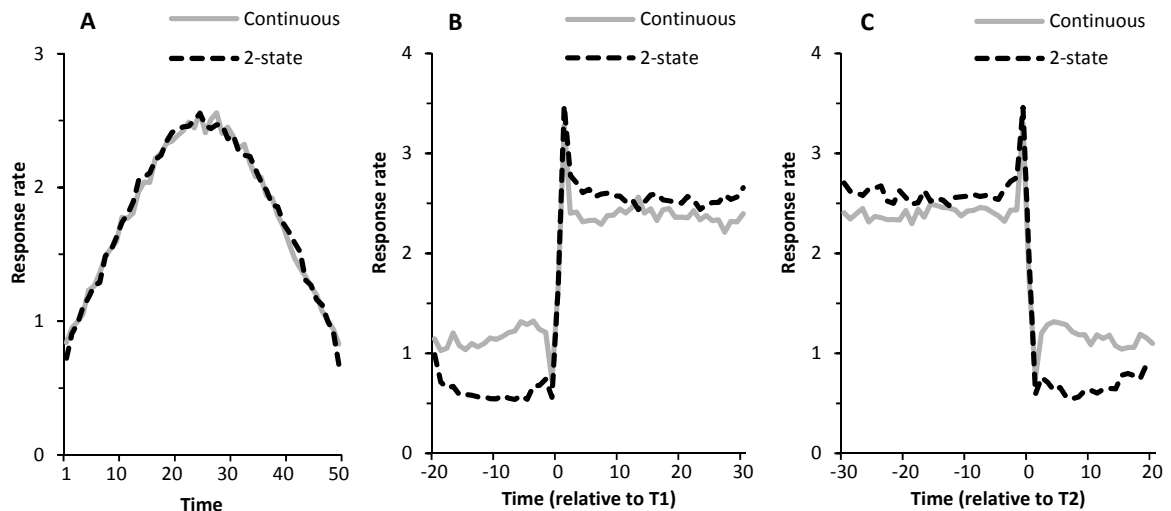


Figure 1. Response rates as a function of time in the trial for two simulated sets of data. In the Continuous data set, the probability of responding increased smoothly across each of the first 25 time bins, and decreased smoothly across the remaining 25 bins. In the 2-state data set, the response probability on each trial changed abruptly from zero to high at one randomly selected time during the first 25 time bins, and then decrease abruptly back to zero at a second randomly selected time in the last 25 bins. The plot in **A** shows that, despite the differences in the way the data were generated, there was a smooth increase and decrease in response rates for both data sets when responses were averaged over 1000 simulated trials. The same averaged data are plotted in **B** and **C** after all trials had been aligned to either the first (t_1) or second (t_2) break points in each trial. The break points were identified using a search algorithm that fits three horizontal line segments to response rates in three separate intervals, separated at t_1 and t_2 , within each trial.

As shown above, averaging response rates across many trials can produce artifactual evidence for both smooth and abrupt changes in response rates during the trial. Therefore, the mean response rate over trials cannot be used to distinguish between accounts that describe CRs as emerging in a graded or abrupt fashion. However, these descriptions of timed CRs do differ in what they predict about the distribution of response rates. If responding emerges in a graded fashion during each trial, as might be predicted by theories that identify conditioning with a continuously graded variable such as associative strength, then we should expect to see a single unimodal distribution of response rates that shifts as time elapses during the trial. By contrast, if responding steps abruptly between a low and high state during each trial, as suggested by timing theories, then the distribution of response rates should be a

mixture of two distinct distributions corresponding to the two underlying response states. In this case, the continuous change in response rate (averaged over trials) should be accompanied by a continuous change in the relative proportions of the two distributions, but the means of those distributions themselves should not change. That is, if one samples response rates from a time window near the beginning of a trial, responses on the majority of trials should fall within the lower distribution of response rates. If one samples responses from a later time window, when the likelihood is high that the subject has already commenced responding, the majority of trials should fall within the higher distribution of response rates.

The aim of the present experiments was to describe the distribution of response rates during conditioning with a fixed CS-US interval, and to analyze how this distribution changes as

time elapses within a trial. Both experiments trained rats in a magazine approach paradigm using the peak procedure. Experiment 1 trained two groups of rats, one with a fixed CS-US interval of 30 s, the other with a variable CS-US interval that had a mean of 30 s. For both groups, responding was measured continuously across non-reinforced presentations of the CS that lasted for 60 s. Experiment 2 also compared two groups, both with fixed CS-US intervals that were either 20 s or 40 s, and responding was measured across 120-s non-reinforced presentations of the CS. The goal of both experiments was to see whether the change in mean response rate during the CS in the groups with fixed CS-US intervals reflected the gradual shift in location of a single response distribution or a gradual shift in weight between two static response distributions.

Experiment 1

In Experiment 1, two groups of rats were trained with a fixed duration (60-s) CS that was reinforced on 50% of trials with a single food pellet (the US). The pellet was delivered either 30-s into the CS (Group FT30) or at a variable time between 1 and 59 s (Group VT30). On the remaining 50% of trials, the CS was presented for 60 s but no food pellet was delivered. Response rates were measured during the non-reinforced trials. (The relatively high proportion of non-reinforced trials served to increase the data yield, which was important when attempting to estimate the full distribution of response counts, rather than just a single measure of central tendency.) Rats in Group FT30 should show a systematic change in their response rate, rising across the first 30-s of the CS and then declining again beyond that time (Church et al., 1994). An interpretation of this pattern in terms of timing of the CR to the US would be confirmed by comparing Group FT30 to Group VT30, because rats in the latter group should maintain a relatively steady level of responding across the 60-s CS presentation (Harris & Carpenter, 2011; Harris, Gharaei, & Pincham, 2011). The rats in both groups were trained for 40 days, and their

response rates were analyzed from the last 20 days. There were 30 non-reinforced trials per day, giving a total of 600 trials for the data analysis. Before describing that experiment in more detail, I will first describe the analysis developed here. I will then use the simulated data described earlier to demonstrate how the analysis distinguishes between the contrasting descriptions of timed CRs as emerging in either a graded or an abrupt fashion.

Data Analysis

All analyses were performed on data from individual rats. Each non-reinforced trial was divided into 5-s time bins: a pre-CS bin covering the final 5-s of the inter-trial interval (of which the mean length was 90 s), and 12 bins covering the 60-s CS presentation. Response rates were calculated for each bin of each trial, and these were used to generate 13 frequency distributions of response rates.

In order to compare the two theoretical positions being investigated, the analysis combined elements of both positions. It modeled the frequency distribution of each rat's response rate using a Poisson probability density function (pdf) that shifted as time elapsed during the CS, reflecting a continuous change in the underlying conditioning strength. It then computed a weighted average of this positive response distribution and a no-response distribution with a mean and variance equal to zero. Thus, the analysis assumed that the full distribution was the composite of occasions when the rat was in a response state, the strength of which could change at different time points in the CS, and occasions when the rat was in a no-response state. The analysis also assumed that the two response states were exhaustive, such that the weights of the response distribution and no-response distribution summed to one. Accordingly, frequency distributions of response rates were analyzed by fitting the function, $F(x)$, shown in Equation 1.

$$F(x) = w \cdot P(x, \mu) + (1 - w) \cdot Z(x)$$

Equation 1

The equation describes the function for values of x (response counts, ≥ 0) as the weighted sum of a positive response distribution, modeled using the Poisson pdf, P , with mean μ , and a zero function, Z , where $Z(x) = 0$. The weight parameter, w (for $0 \leq w \leq 1$), describes the probability that the rat is in the response state, and thus $1-w$ is the probability that the rat is in the no-response state. When fitting Equation 1 to response distributions, parameters μ and w were free to vary.

Using a Poisson pdf to model the positive response distribution assumes that, when the rat is in the high response state, the probability of making a response is uniform across time. Describing the full response distribution as a mixture of a Poisson pdf and a zero distribution is similar to other mixture-model approaches that have successfully described the distribution of response intervals in instrumental conditioning paradigms using variable-interval schedules of reinforcement (e.g., Brackney, Cheung, Neisewander, & Sanabria, 2011; Shull, Gaynor, & Grimes, 2001; Shull, Grimes, & Bennett, 2004). A notable difference between those approaches and the current one is that those approaches model behavior as a mixture of two Poisson processes, one describing high response rates (during response bouts) and the other capturing long response intervals (pauses between bouts). However, the evidence for two Poisson processes requires measurement of responding over long time windows in order to identify the long response intervals associated with the low frequency Poisson process. This requirement is incompatible with the present strategy of segmenting the trial into relatively short time bins in order to obtain independent response distributions at multiple time points within the trial. Moreover, the present approach of using one Poisson pdf and a zero pdf to model response distributions has the advantage that it involves only two parameters, whereas a model

that uses two Poisson pdfs requires three parameters (a mean for each Poisson pdf, and the weighting parameter).

Fitting Equation 1 to the response rates measured at different time points during the CS provides a means to test the strength of evidence for graded versus abrupt changes in responding within a trial. If responding emerges in a continuously graded fashion, the mean of the Poisson pdf should vary across time points within the CS, and this will account for the change in mean response rate across the trial. If, on the other hand, responding emerges abruptly, the relative weights of the Poisson and zero distributions will vary systematically over the trial, and this parameter will account for the change in mean response rate across the trial.

To assess the sensitivity of the analysis method described above, it was tested on the continuous and 2-state sets of simulated data described in the Appendix. Each simulated trial was divided into 10 bins, and frequency distributions of response counts in each bin were computed across trials. Equation 1 was fitted, using the method of least squares, to each of the 10 frequency distributions for both data sets. Figure 2 shows how the mean and weight of the Poisson change across the 10 bins. When modeling the continuous data set (Figure 2A), the mean of the Poisson distribution changed systematically across time bins, corresponding closely to the change in response rate in the same data set, but there was very little change in the relative weights of the two distributions (the weight of the Poisson pdf was consistently equal to 1). When modeling the 2-state data set (Figure 2B), the weight of the Poisson function changed systematically across time bins, closely matching the change in response rate in that data set. In contrast, the mean of the Poisson changed very little, and did not match the change in response rate.

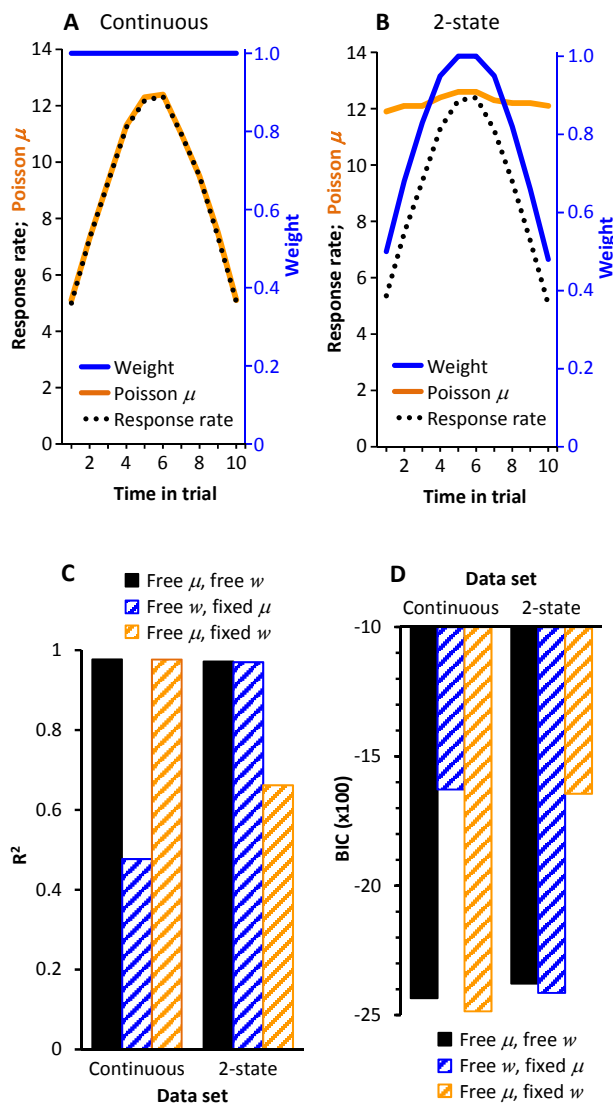


Figure 2. [LEFT] The plots in **A** and **B** show response rates (black dotted line) as a function of time, for the same two simulated data sets plotted in Figure 1, but with trial length compressed into 10 bins. The plots also show the values obtained for the mean (μ , orange line) and weight (blue line) of the Poisson pdf when Equation 1 was fitted to the frequency distributions of response rates in each time bin. For the Continuous data set (**A**) the mean of the Poisson tracked the change in response rate whereas the weight did not vary across time bins; for the 2-state data set (**B**), the weight of the Poisson changed in close step with the change in response rate whereas the mean of the Poisson changed very little. **C** shows mean goodness of fit (R^2) for different versions of Equation 1 fitted to the frequency distribution of response rates across time bins. In one function, both the Poisson mean and weight (w) varied as free parameters during the fitting operation; in the other functions, either the mean was fixed and the weight varied as a free parameter or the weight was fixed and the mean varied as a free parameter. **D** shows the Bayesian Information Criterion (BIC) values obtained from the same analysis fitting different versions of Equation 1 to the response distributions. Note that more negative values for BIC indicate better evidence for that equation given the data.

To quantify the differences in how Equation 1 described the two simulated data sets, two different versions of the function were fitted to the same data sets. One version used a single fixed value for μ , selected to optimize the fit across all time bins, but allowed w to vary between time bins. This model assumes that responding is described by two fixed distributions—the zero distribution and a Poisson pdf—and variations in responses across time can only be accounted for by changes in the relative weights of these two distributions. The second version of the function used a single fixed value for w , selected to optimize the fit across all time bins, but allowed μ to vary across time bins.

This model assumes that changes in responding across the CS reflect a continuous change in the mean of the Poisson pdf, and the relative weighting of this distribution does not account for systematic changes in response rate. Figure 2C shows the average fit (R^2) to the 10 time bins of the two data sets for each of the three versions of Equation 1. When modeling the continuous data set, there was no loss of fit when fixing w and allowing only the mean of the Poisson pdf to vary across time bins, compared with the fit obtained when both μ and w varied for each fit. There was, however, a very substantial loss of fit when μ was fixed and only w varied. The reverse was true when modeling the 2-state data set. There was almost no decrement in fit when μ was fixed and w was free, but there was a large decrement in fit when w was fixed and μ was free.

Comparing the different models for their fit (as R^2) does not take into account the difference between them in the number of their free parameters. This can be done by calculating the Bayesian Information Criterion (BIC) for each model (Wagenmakers, 2007), according to the formula

$$\text{BIC} = n \cdot \log_e(\text{SSE}/n) + k \cdot \log_e(n)$$

Equation 2

where n is the total number of data points being fitted, SSE is the sum of squares of the difference between the fitted and observed values of each data point, and k is the number of free parameters in the model. When comparing two models, a lower BIC value indicates a better fit to the data (a smaller SSE). This is also true for negative BIC values which arise when SSE/n is much less than 1 (since the log of a fraction less than 1 is negative). To calculate a single BIC value for all fits to each data set, we pooled the SSE across fits at each time bin (i.e., for 10 bins x 25 data points, $n = 250$), and pooled the number of free parameters across bins. Thus, $k = 20$ (2 parameters x 10 bins) for Equation 1 with two free parameters; $k = 11$ ($1 \times 10 + 1$) for both models with one fixed and one free parameter. The BIC values are shown in Figure 2D; smaller (in this case, more negative) values signify better evidence for the model given the data. For the Continuous data set, the BIC score is 50 units lower when μ was free and w was fixed than when both μ and w were free. Applying the exponential function, $e^{x/2}$, to this difference (x) gives the odds ratio of the evidence for each of the two models. In this case, the difference in BIC equates to very large odds (more than $10^{10}:1$) favoring the fixed- w model over the free- w free- μ model. The fixed- w model was also 858 units lower than the fixed- μ model, which equates to extremely large odds (more than $10^{186}:1$) favoring the fixed- w model. The BIC values tell the opposite story for the 2-state data. Here, the BIC is 36 units lower when w was free and μ was fixed than when both μ and w were free, which equates to very large odds (more than $50,000,000:1$) favoring the fixed- μ model over

the free- w free- μ model. Not surprisingly, there is much stronger evidence for the fixed- μ model than for the model that fixes w but varies μ . The difference in these BIC scores is 770, which equates to extremely large odds (more than $10^{167}:1$) favoring the fixed- μ model over the fixed- w model.

In sum, the analysis developed here is sensitive to differences in the form of the frequency distribution of response rates. Therefore, the approach should allow us to test between two different descriptions of timed CRs, those in which responses emerge in a graded fashion across each trial versus those in which responding appears abruptly within each trial.

Methods

Subjects. A total of 24 experimentally naïve rats were run in Experiment 1. The experiment was first run with 16 male Hooded Wistar rats (8 per group), 7 to 8 weeks of age at the start of the experiment, obtained from the Laboratory Animal Services breeding unit at The University of Adelaide, South Australia. Due to defective wiring in three conditioning chambers, six rats (three per group) did not receive any white noise stimulus during all or some large part of the experiment. The data of these rats were omitted from all analyses. Subsequently, another eight rats were run (four per group). As a result of interim changes in the supply of animals, the new rats were female albino Sprague Dawley rats, obtained from the Animal Resources Centre, Perth, Western Australia. Therefore, usable data were collected from 18 rats, with each group comprised of five male hooded Wistar rats and four female Sprague Dawley rats.

During the initial run with 16 male rats, the rats were housed in groups of eight in large white plastic boxes, measuring 59 x 37 x 26 cm (length x width x height). The eight rats run subsequently were housed in groups of four in split-level ventilated plastic boxes (Techniplast™ S.p.A., Buguggiate, Italy), measuring 40 x 46 x

40cm (length x width x height). All rats were located in the animal colony maintained by the School of Psychology at the University of Sydney. They had unrestricted access to water in the home tubs. Three days prior to commencement of the experiment, they were placed on a restricted food schedule. Each day, half an hour after the end of the daily training session, each box 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 box. This amount is approximately equal to their required daily energy intake (Rogers, 1979), and took at least 2 h to be eaten (but was usually finished within 3 h). This meant that all rats in the tub had access to food for an extended period, which should reduce differences between rats in their levels of hunger.

Apparatus. Rats were trained and tested in 8 Med Associates™ conditioning chambers measuring 30 x 25 x 28.5 cm (length x depth x height). The end walls of each chamber were made of aluminum; the sidewalls and ceiling were Plexiglas™. The floor of the chamber consisted of stainless-steel rods, 0.5 cm in diameter, spaced 1.5 cm apart. Each chamber had a recessed food magazine in the center of one end wall, with an infra-red LED and sensor located just inside the magazine to record entries by the rat. A small metal cup measuring 3.5 cm in diameter and 0.5-cm deep was fixed on the floor of each food magazine. Attached to the food magazine was a dispenser delivering 45-mg food pellets (purified rodent pellets; Bioserve, Frenchtown, NJ). Each chamber was enclosed in a sound- and light-resistant wooden shell. Throughout all sessions, fans located in the rear wall provided ventilation; the operation of these created a background noise level measuring 70dB. Experimental events were controlled and recorded automatically by computers and relays located in the same room. White noise (78dB) was presented from a speaker mounted on the

wall of each conditioning chamber above and to the left of the food magazine.

Procedure. Prior to the start of conditioning, rats received a single 20-min magazine training session during which 20 food pellets were presented on a VT 1-min schedule, with no stimulus presentations. Rats that ate fewer than half of the pellets were given a second session of magazine training the following day. After all rats completed magazine training, they commenced daily conditioning sessions 5 days per week for a total of 40 days. Each session contained 60 trials, in which the noise CS was presented for exactly 60 s. On 30 of these trials, a food pellet was delivered during the CS presentation; no pellet was delivered on the remaining 30 trials. These reinforced and non-reinforced trials were randomly intermixed, with the constraint that each block of 12 trials included six reinforced and six non-reinforced trials. For Group FT30, the food pellet was delivered midway through the CS presentation (exactly 30 s after CS onset). For Group VT30, the pellet was delivered at a random time between 1 s and 59 s after CS onset (mean = 30 s). The inter-trial interval varied randomly between 40 s and 140 s (mean = 90 s). Photo-beam interruptions by head entry into the magazine were recorded during each CS and each 20-s pre-CS period. A single response was recorded per beam break, and time stamped at the first moment that the beam was broken. Sessions lasted approximately 2.5 h.

Results

Only magazine activity during non-reinforced trials was analyzed, so that response rates were not affected by the arrival of the food pellet. Figure 3A shows response rates during the CS and the pre-CS period across training days for both groups. Response rates during the CS rose quickly, and remained relatively stable over the second half of the experiment (Days 21 to 40). Figures 3B and 3C also show how, for each rat, response rates changed as time elapsed during the CS presentations, averaged for 5-s time bins over all non-reinforced trials from Days 21 to 40.

For Group FT30, response rates increased steadily over the first 30 s, and then decreased again, indicative of responses being timed to coincide with the expected time of the US. For Group VT30, response rates were elevated from the start of the CS presentation, and remained relatively elevated over the 60 s, consistent with the uniform distribution of reinforcement times in this group. A mixed-model ANOVA, with Greenhouse-Geisser corrections to the degrees of freedom when the data violated assumptions of sphericity, showed that there was no significant difference in response rates overall between the two groups, $F(1, 16) = 0.72$, $p = .410$. However, there was a significant main effect of Time, $F(1.6, 25) = 11.80$, $p < .001$, $\eta^2_p = 0.43$ (0.36 to 0.67), and a significant interaction between Time and Group, $F(1.6, 25) = 9.81$, $p = .001$, $\eta^2_p = 0.38$ (0.31 to 0.71). Trend analyses showed that there were significant linear and quadratic trends across time, $F(1, 16) = 5.05$ and 21.09 , $ps = .039$ and $<.001$, $\eta^2_p = 0.24$ and 0.57 , both of which interacted significantly with Group, $F(1, 16) = 7.36$ and 14.26 , $ps = .015$ and $.002$, $\eta^2_p = 0.32$ and 0.47 . Follow-up ANOVAs on each group individually showed that there was a significant quadratic trend, $F(1, 8) = 19.16$, $p = .002$, $\eta^2_p = 0.71$, but not linear trend, $F(1, 8) = 0.12$, $p = .743$, for Group FT30; whereas, for Group VT30, there was a significant linear trend, $F(1, 8) = 11.63$, $p = .009$, $\eta^2_p = 0.59$, but not quadratic trend, $F(1, 8) = 1.93$, $p = .202$.

To obtain frequency distributions of response rates, the data from each CS presentation for each rat from Days 21 to 40 were divided into 12 5-s bins. Frequency distributions of pre-CS response rates were obtained from the final 5-s bin of the 20-s pre-CS interval of each trial. Examples of the frequency distribution of pre-CS response counts, and of response counts from two time bins during the CS (0-5 s, and 25-30 s), for one rat (Rat 2) are shown in Figures 3D, 3E and 3F. For the data from each rat, $F(x)$, as

defined in Equation 1, was fitted, using the method of least squares, to the pre-CS response distributions and the response distributions from each of the 12 time bins during the CS. The response distributions were well fitted by $F(x)$; the average R^2 across all functions was 0.98 for Group FT30 and 0.96 for Group VT30. Three examples of the function, as well as the Poisson pdf in each case, fitted to data from Rat 2 are shown in Figures 3D, 3E and 3F. The functions were used to analyze how the mean and weight of the Poisson pdf changed across time bins.

Figures 4A and 4B show the mean and weight of the Poisson pdfs for responding during the last 5-s of the pre-CS period and each of the 12 5-s time bins during the CS for Groups FT30 and VT30. When both parameters of Equation 1 were free to vary, the data fitting operation identified functions in which both the mean and the weight of the Poisson varied systematically across time bins in a manner that tracked the change in the rats' response rates. In other words, this step of the analysis suggested that both a change in response strength, modeled by the Poisson mean, and a shift in probability of being in the response state, modeled by the weight, jointly accounted for the change in response rate over time during the CS. To tease apart these contributions, the same data fitting analysis was performed again, but in these analyses one of the two parameters was fixed. One analysis fixed the mean of the Poisson pdf, so that any change in response distribution across time could only be accounted for by a change in w . That is, while w was allowed to vary across each of the 13 time bins, the analysis selected only one value for μ that optimized the fits to the response distribution data across all time bins. Fixing μ in this manner had little influence on the values obtained for w compared to those obtained when both μ and w were free to vary in the initial analysis.

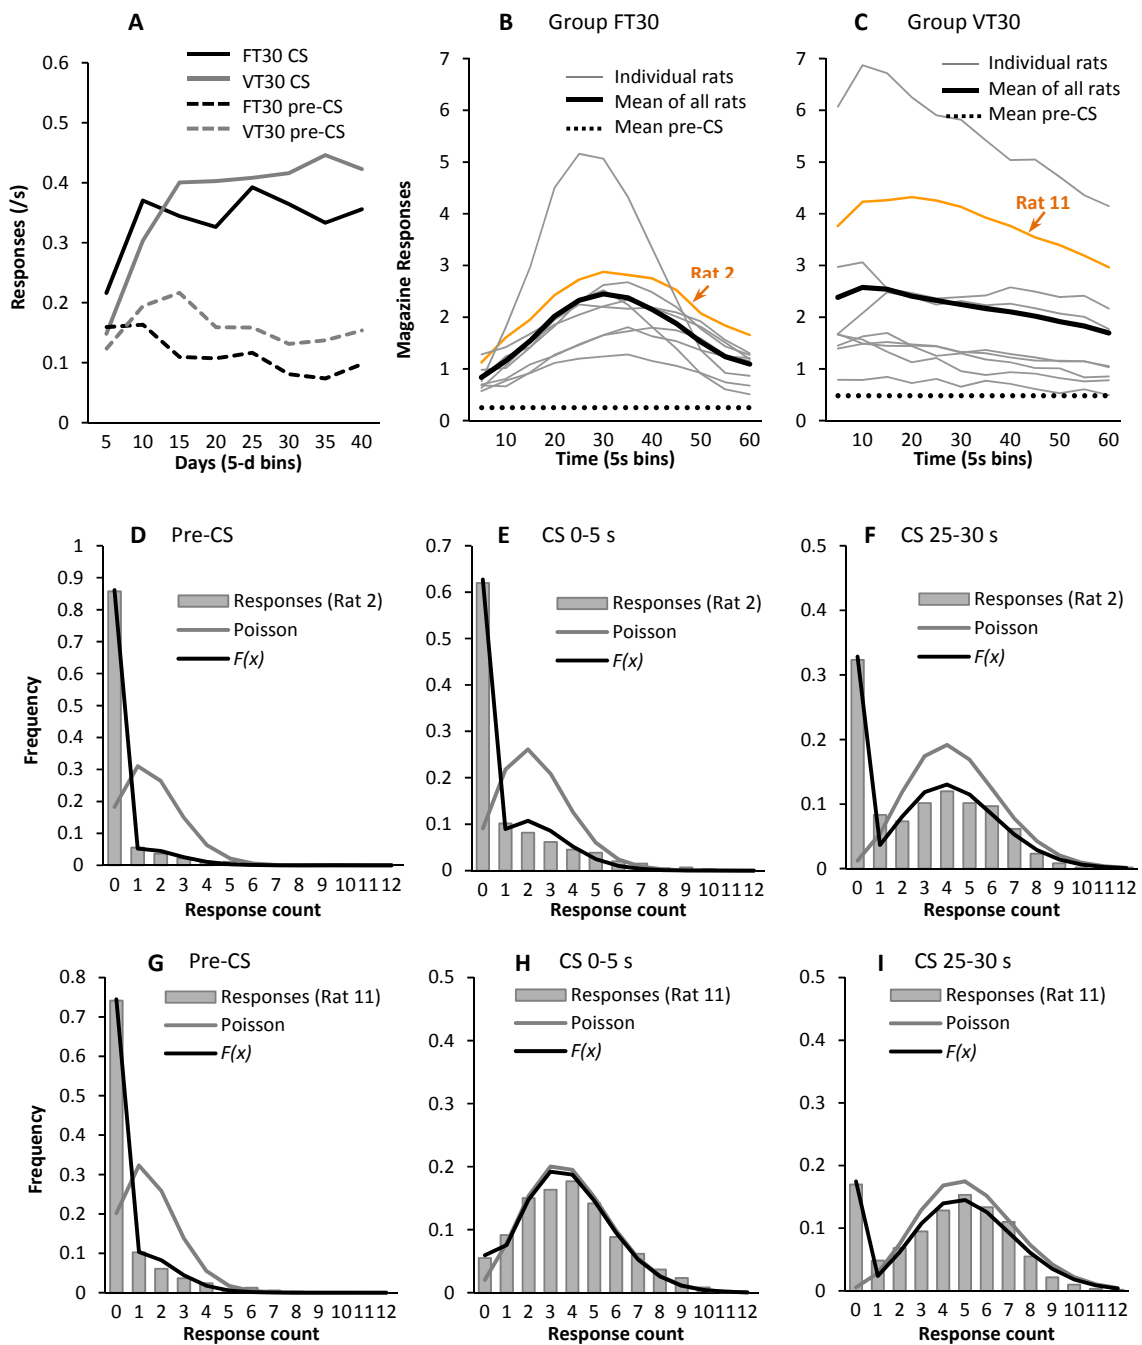


Figure 3. **A:** Response rates during the CS and pre-CS intervals in each 5-day bin for rats in Group FT30 and VT30 in Experiment 1. **B** and **C** show the response rates, in 5-s bins, across the 60-s non-reinforced presentations of the CS for each rat in Groups FT30 and VT30, averaged over Days 21 to 40. These plots also show the mean response rates for each group (thick black lines) and the mean pre-CS response rate (dotted black lines). The six lower plots show the frequency distributions of response counts for data from two representative rats (Rats 2 and 11, highlighted in panels B and C) from Days 21 to 40, for 5-s bins during the pre-CS period (**D** and **G**), during the first 5 s of the CS (**E** and **H**), and during the peak response interval from 25-30 s after CS onset (**F** and **I**). The plots also show the best-fitting function for Equation 1 (black line), as well as the unweighted Poisson pdf (gray line) from that best-fitting function.

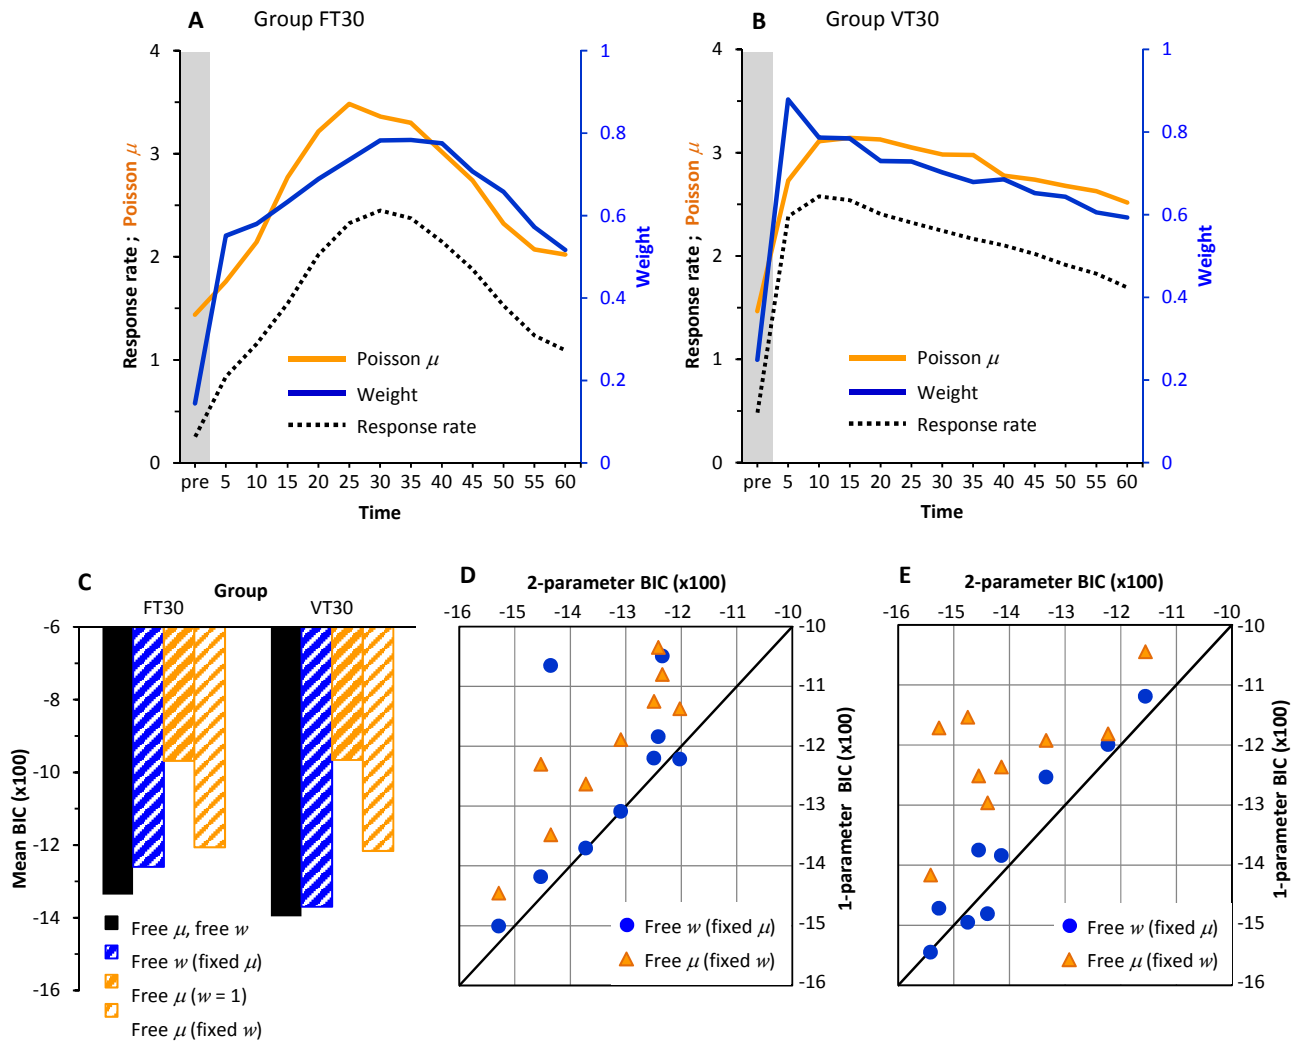


Figure 4. **A** and **B** show response rates (black dotted lines) as a function of time during the CS for Groups FT30 and VT30 in Experiment 1. The plots also show the values obtained for the mean (μ , orange line) and weight (blue line) of the Poisson pdf when Equation 1 was fitted to the frequency distributions of response rates in each 5-s time bin. For both groups, the μ and weight of the Poisson changed over time bins, in close step with the change in response. **C** shows the mean Bayesian Information Criterion (BIC) values obtained when four different versions of Equation 1 were fitted to the frequency distribution of response rates across time bins. In one analysis, both μ and the weight (w) varied as free parameters during the fitting operation; in another analysis, μ was fixed and w varied as a free parameter; in the other two analyses, w was fixed, either at 1 or at a single value that optimized the fit across all time bins, while μ varied as a free parameter. Note that more negative values for BIC indicate better evidence for that function given the data. **D** and **E** plot, for each rat in Groups FT30 (**D**) and VT30 (**E**), the BIC value for the 2-parameter model (free m , free w) against the BIC values for two of the 1-parameter models (free w , fixed m ; and free m , fixed w). Points above the diagonal indicate that the BIC for the 2-parameter model is more negative than the BIC for the 1-parameter model.

Two further analyses were conducted to assess the unique contribution of the Poisson mean in accounting for changes in the distribution of response rates. The first of these analyses set w to 1, so that any change in response distribution could only be accounted for by changes in μ . The second analysis also fixed w , but selected one value across all time bins that optimized the fit of the model to the response distributions. In both of these analyses, the mean of the Poisson pdf changed over time bins, but it appeared to track the change in response rate less well than it had done when w was free to vary.

BIC values were used to compare each of the analyses just described. For each rat, a single BIC was computed that combined all 13 functions fitted to the 13 time bins. The mean BICs, across the 9 rats in each group, are shown in Figure 4C. Even though the BIC penalizes the initial analysis most because it has two free parameters, the BIC values are lower for this analysis than any of the three analyses in which one parameter was fixed across time bins. For both groups, the model with fixed μ came closest to the model with free w and μ : the difference in mean BIC was 77 for Group FT30 and 27 for Group VT30. Nonetheless, these differences equate to large odds favoring the model with two free parameters (more than $10^{16}:1$ for Group FT30; more than $700,000:1$ for Group VT30)¹. Figures 4D and 4E plot each rat's BIC score for the model with μ and w free against its BIC score for the model with μ fixed (but $\neq 1$) and the model with w fixed. These plots show that, for some rats, the model with two free parameters and the model with fixed μ (free w) performed equivalently (and indeed the model with fixed μ was superior for 3 rats), but in the majority of cases the 2-parameter model was superior. The plots also show that the 2-parameter model was superior to the model with fixed w (free μ) for

every rat. Therefore, information about the frequency distribution of response rates is lost when either w or μ is fixed, and therefore both parameters account for a significant amount of the data.

It is clear in Figure 4C that the BIC score is lower (more negative) when the mean of the Poisson was fixed and w was free compared to either of the analyses in which w fixed and μ varied. The difference is largest when w was set to 1: for both groups, the mean BIC score was more than 290 higher than the BIC score obtained when μ was fixed and only w varied. Allowing w to be fixed to a value that optimized the fits to the distributions across all time bins reduced the BIC scores considerably—by more than 230 for both groups. This shows that, even if w itself is not free to vary across time bins, there is a very substantial gain when it is set below 1. Therefore, the distribution of response rates is better accounted for as the sum of two distributions, one of which has all its mass at zero, than as a single distribution with mean greater than one. Nonetheless, even when w was fixed at a value below 1, the best fit to the data achieved by varying μ was still considerably worse than when μ was fixed but w varied. The difference in mean BIC scores was 54 for Group FT30 and 153 for Group VT30, which equate to odds ratio more than $10^{11}:1$ and $10^{33}:1$ in favor of the model in which w was free. Indeed, the fixed- μ model was superior to the fixed- w model for almost every rat, as revealed by the difference in heights of the two sets of symbols (circles and triangles) in Figures 4D and 4E. Therefore, it is clear that the data are accounted for better when the mean of the Poisson is fixed and changes in response distribution are accounted for by w alone, than when w is fixed and the changes in response distribution are accounted for by μ .

¹ These odds are calculated on the difference in mean BIC values, and therefore represent a comparison of the average evidence for each model per rat. A comparison of the total amount of evidence for the

whole group would be in the same direction but much larger (the difference in overall BIC values would be almost nine times greater than the difference in mean BIC values).

The BIC comparisons between models was repeated but excluding the first 5-s bin of the CS presentations. This was done in case response rates in that bin were affected in a particular way by the CS onset which might distort the analyses of response distributions. However, the new analyses, without data from the first bin, revealed a pattern of BIC scores among the models that was very similar to the pattern reported above.

In the analyses described so far, the functions fitted to the response distributions described the weighted average of a Poisson pdf and a zero distribution. This assumes that the rats were either in a state to produce some measurable level of responding, captured by a single distribution, or in a state that produced no responses. A more complex analysis would characterize the rats' responses as comprising a high response state and a low response state. I have conducted such an analysis that identifies the low response state with the baseline response rate during the inter-trial interval. The first step of this analysis was to fit Equation 1 to the frequency distributions of pre-CS response rates, in order to obtain a function describing baseline responding. This baseline function, $B(x)$, was then substituted for $Z(x)$ in Equation 1, and the new function was fitted to the data from each time bin during the CS. This analysis produced marginally better fits to the CS response data than were obtained from the previous analyses using the zero distribution, but did not allow direct comparisons to be made between pre-CS and CS response distributions. Nonetheless, when considering the response distributions during the CS, the results of both sets of analyses were virtually identical. All functions in which both μ and w were free provided better fits to the response distributions than functions that fixed one or other parameter, and functions with free w and fixed μ produced better fits than functions with fixed w and free μ .

Discussion

Different patterns of responding were observed when the CS-US interval was fixed versus when it varied. A fixed 30-s interval between CS and US led to accurately timed responding, such that response rates rose across the first 30 s of the CS, and then decreased again for the next 30 s. A variable CS-US interval led to a sharp rise in responding at CS onset that remained elevated for the full length of the CS. Both findings replicate previous demonstrations (Church et al., 1994; Harris & Carpenter, 2011; Harris et al., 2011; Smith, 1968; Williams et al., 2008), and show that conditioned responses track the expected time of US arrival. Nonetheless, there was a modest fall in response rate across the CS in Group VT30, evident in Figure 3C, that has not been observed in our previous experiments using variable CS-US intervals. It is likely that the decline in response rate observed here reflects the peculiarities of the partial reinforcement schedule in which the probability that a trial will be a reinforced one—as opposed to a non-reinforced trial—decreases as the trial progresses without reinforcement.

The detailed analyses using Equation 1 to model the frequency distribution of response rates show that more information about response rates is accounted for by changes in the relative weight of two response distributions than by a progressive shift in the mean of one response distribution. In other words, there is more evidence for a model in which changes in responding across a trial are attributed to changes in the probability of being in one of two response states than for a model in which changes in responding are attributed to the strength of responding when in the response state. However, the most information about the response distributions was accounted for by a combination of both the weight and mean of the response distributions. Therefore, the available evidence indicates that changes in responding across a trial are jointly explained by both a change in the probability of being in one of two response states and a change in strength of responding when in the response state.

In this experiment, 50% of trials were not reinforced so that response rates could be measured across the full 60 s of the CS without the rats' behavior being contaminated by the delivery of the food pellet. Using a 50% reinforcement schedule would have reduced response rates overall, compared to a 100% schedule, and may have had the effect of inflating the proportion of trials in which the rats did not respond. This raises the possibility that the shift between the zero response distribution and the positive response distribution might only exist in the form found here because of the high proportion of non-reinforced trials. However, other evidence indicates that this is unlikely. For example, I have performed the analysis described here on data from an experiment in which rats were trained with a 10-s CS reinforced on termination of 33% of trials and a 30-s CS reinforced on termination of 100% of trials (Group Fixed of Experiment 3 in Harris, Patterson, & Gharaei, 2015). Both the mean of the Poisson and the relative weight of the Poisson versus the zero distribution rose across 5-s bins during both CSs, just as was observed in the present experiment for time bins before the expected time of reinforcement. Moreover, the shift in weight occurred at an earlier time bin for the partially reinforced CS than for the 100% CS, although the weights in the final bin of each CS were very similar. This shows that the changes in weight and magnitude of the response distribution are tied to the timing of reinforcement within a trial rather than the percent of reinforcement across trials.

Experiment 2

Experiment 1 examined changes in the distribution of rats' response rates as time elapsed within a trial, and compared response rates conditioned using a single fixed CS-US interval of 30 s with responses conditioned using a variable CS-US interval with a mean of 30 s. The present experiment sought to extend the analysis by conditioning two groups of rats with fixed CS-US intervals of either 20 s, for Group FT20, or 40 s, for Group FT40. In these cases, we

should again expect to see that response rates rise across the initial portion of each CS presentation, and then decline again, as we saw in Group FT30 of Experiment 1. However, if these changes are accurately tracking the timing of the US itself, then the time point at which responses peak should differ in each group—responding should peak at approximately 20 s into the CS for Group FT20 and should peak at approximately 40 s into the CS for Group FT40. This experiment also differed from Experiment 1 in that each CS presentation was much longer, lasting 120 s (compared with 60 s in Experiment 1). This provided a greater opportunity for response rates to return to baseline, in order to give a more complete picture of how the distribution of responses changes both as responding increases from baseline to its peak, and then decreases again down to that baseline. To anticipate the results of the experiment, while response rates did not return completely to baseline for Group FT40, they did in fact fall below baseline during the last 60 s of the CS in Group FT20. This presented an additional challenge for each theoretical account, to explain not only the rise and fall in response rates during the CS but also the suppression of responding below baseline when rats learned about the negative correlation between food and the extended portion of the CS.

Methods

Subjects and apparatus. A total of 16 experimentally-naïve female albino Sprague Dawley rats were run in Experiment 2. They were obtained from the Animal Resources Centre, Perth, Western Australia. They were housed in groups of four in split-level ventilated plastic boxes (Techniplast™), measuring 40 x 46 x 40cm (length x width x height), with unrestricted access to water. All rats were located in the animal colony maintained by the School of Psychology at the University of Sydney. Three days prior to commencement of the experiment, they were placed on a restricted food schedule described for Experiment 1. The rats were trained and tested in 16 Med Associates™

conditioning chambers, as described for Experiment 1.

Procedure. The rats were divided into two groups of eight. Prior to the start of conditioning, they received a single 20-min magazine training session during which 20 food pellets were presented on a VT 1-min schedule, with no stimulus presentations. Rats that ate fewer than half of the pellets were given a second session of magazine training the following day. After all rats completed magazine training, they commenced daily conditioning sessions 5 days per week for a total of 45 days. Each session contained 40 trials, in which the noise CS was presented for exactly 120 s. On 20 of these trials, a food pellet was delivered during the CS presentation; no pellet was delivered on the other 20 trials. These reinforced and non-reinforced trials were randomly intermixed, with the constraint that each block of eight trials included four reinforced and four non-reinforced trials. For Group FT20, the food pellet was delivered 20 s after CS onset. For Group FT40, the pellet was delivered 40 s after CS onset. The inter-trial interval varied randomly between 40 s and 120 s (mean = 80 s). Photo-beam interruptions by head entry into the magazine were recorded during each CS and each 20-s pre-CS period. Sessions lasted approximately 2¼ h.

Results

Only magazine activity during non-reinforced trials was analyzed, so that response rates were not affected by the arrival of the food pellet. As can be seen in Figure 5A, response rates during the CS (averaged across the full 120-s presentations) increased and response rates during the pre-CS period decreased across days. Response rates within 5-s bins during the CS, and in the pre-CS period, averaged over the last 20 days of the experiment, are shown in Figures 5B (Group FT20) and 5C (Group FT40). For both groups, response rates rose over the initial portion of the CS presentations, and then decreased again. The peak of responding was very close to 20 s for Group FT20, and close to 40

s for Group FT40. Indeed, the mean response rate peaked at the fourth bin (15-20 s) for Group FT20 and peaked at the eighth bin (35-40 s) for Group FT40. A Mann-Whitney U test confirmed that the rats in Group FT20 reached their peak response rate at a significantly earlier time bin than the rats in Group FT40, $z(7) = 3.31, p < .001$.

As is evident in Figure 5C, response rates in Group FT40 did not return fully to baseline (the pre-CS rate) during the CS presentations. A paired *t*-test showed that the average response rate over the last 60 s of the CS (mean = 0.69 responses per 5-s bin) was significantly greater than the baseline (mean = 0.45 responses per 5-s bin), $t(7) = 3.40, p = .011$, Cohen's *d* = 1.20 (95% confidence intervals: 0.25 to 2.11). By contrast, response rates in Group FT20 fell below the pre-CS baseline across the second half of the CS presentation (Figure 5B). Averaged across this 60-s period, response rates during the CS (mean = 0.38 responses per 5-s bin) were significantly below baseline (mean = 0.47 responses per 5-s bin), $t(7) = 3.61, p = .009$, Cohen's *d* = 1.28 (0.3 to 2.21). Therefore, while rats in Group FT20 learned to respond in anticipation of the food pellet during the first half of each CS presentation, they also learned to suppress responding during the second half of the CS presentation.

Frequency distributions of response rates for each rat were obtained by dividing the data from each non-reinforced CS presentation from Days 26 to 45 into 24 5-s bins. The frequency distribution for pre-CS response rates was based on responses during the final 5-s of the 20-s pre-CS interval from each trial. $F(x)$, as defined in Equation 1, was fitted to the pre-CS response distributions and the response distributions from each of the 24 time bins during the CS. Figures 6A and 6B show the mean and weight of the Poisson pdfs for the 25 time bins for Groups FT20 and FT40. In Group FT20, both the mean and weight varied systematically over the time bins, tracking the changes in response rates over those bins. However, the weight did not rise across the first four bins of the CS, as response

rate did, but instead showed a sharp jump from the pre-CS bin to the first CS bin and then slightly decreased up to the 5th bin. The mean and weight of the Poisson pdfs also varied for Group FT40, although changes in w seemed to track the change in response rate better than did the change in the mean of the Poisson.

As was done for Experiment 1, three additional analyses fitted Equation 1 to the response distributions, but one parameter was fixed across time bins. In one analysis, μ was fixed across time bins; in the other two analyses, w was fixed— w was set to 1 or was fixed to a value that optimized the fit across all time bins. Figure 6C shows the mean BIC scores for these three analyses, as well as the initial analysis in which μ and w were free when Equation 1 was fitted to the data. For Group FT20, the three additional analyses produced poorer fits to the data, shown as higher (less negative) BIC scores, compared to the fits obtained when both parameters varied. The analysis in which w was free and μ was fixed came closest to the model with both parameters free: the difference in BIC

was 252, which equates to odds of more than 10^{50} :1 favoring the analysis with two free parameters. As evident in Figure 6D, which plots BIC scores for the 2-parameter model against the 1-parameter models for each rat, the model with two free parameters was superior in all but one rat (shown as the one blue circle below the diagonal). For Group FT40, the analyses in which w was fixed and μ was free produced poorer fits to the data compared to the fits obtained when both parameters varied (the difference in BIC was at least 200). However, the analysis in which w was free and μ was fixed was almost indistinguishable from the model with both parameters free: the difference in their BIC scores was just 0.58, which equates to odds of 1.34:1 favoring the analysis with two free parameters. This conclusion is also born out when comparing the BIC scores of the two models for each individual rat—as shown in Figure 6E, where all cases (blue circles) are close to the diagonal and as many sit below the line as above it.

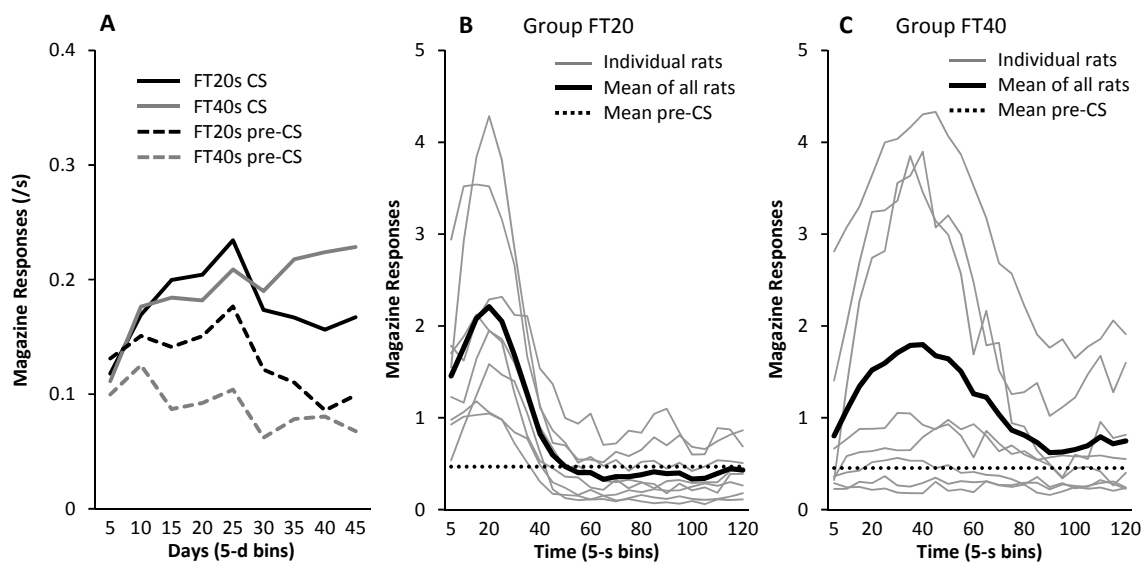


Figure 5. **A:** Response rates during the CS and pre-CS intervals in each 5-day bin for rats in Group FT20 and FT40 in Experiment 2. **B** and **C** show the response rates, in 5-s bins, across the 120-s non-reinforced presentations of the CS for each rat in Groups FT20 and FT40, averaged over Days 26 to 45. These plots also show the mean response rates for each group (thick black lines) and the mean pre-CS response rate (dotted black lines).

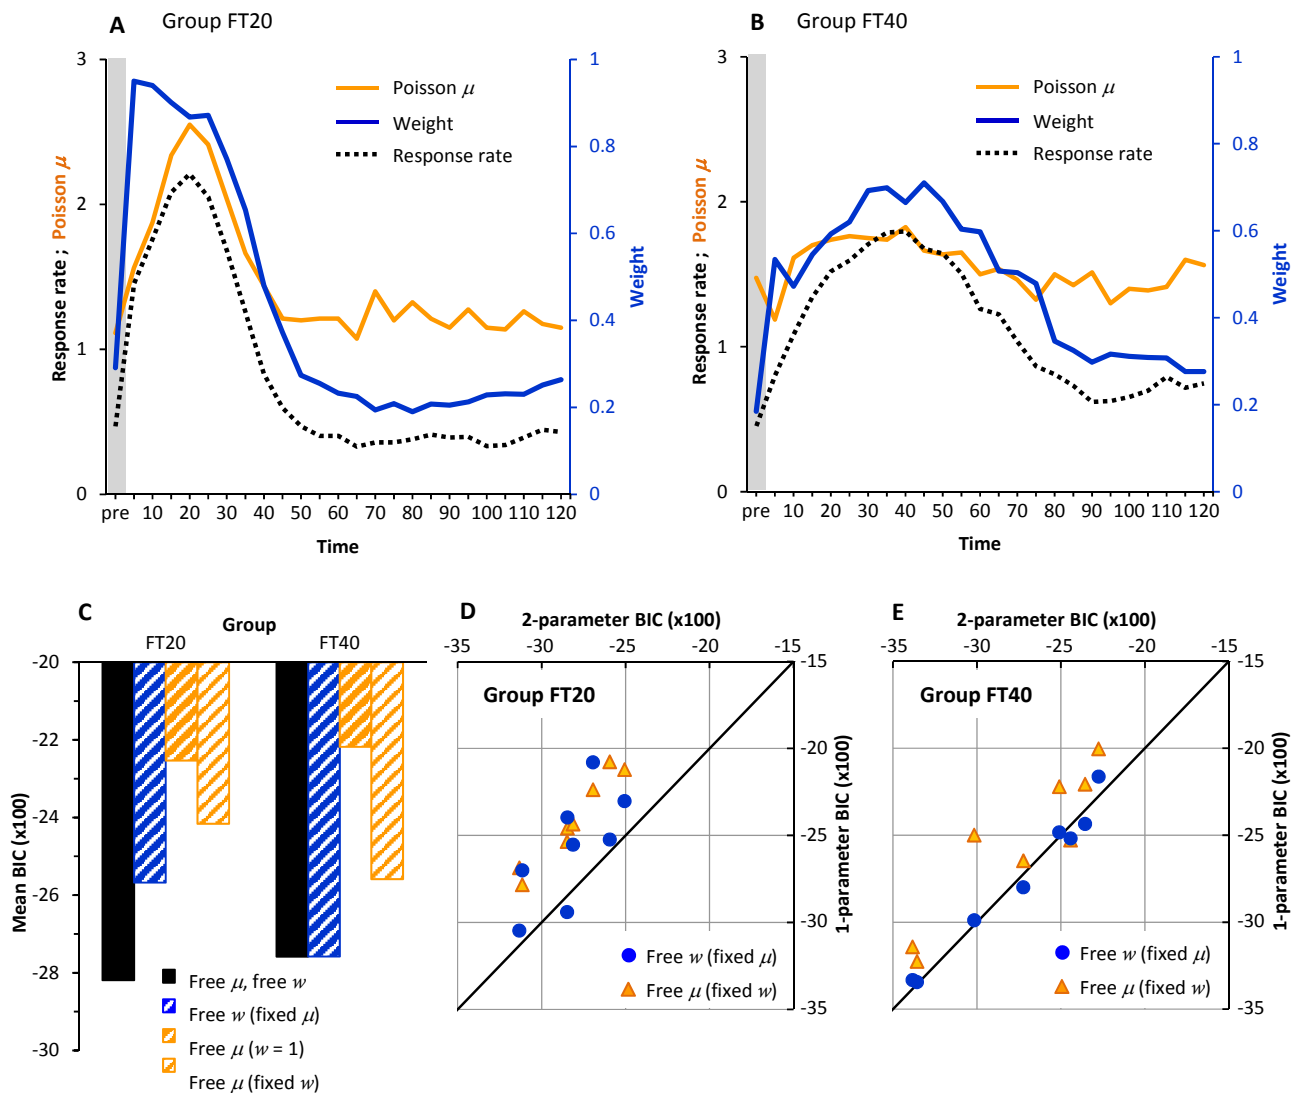


Figure 6. **A** and **B** show response rates (black dotted lines) as a function of time during the CS for Groups FT20 and FT40 in Experiment 2. The plots also show the values obtained for the mean (μ , orange line) and weight (blue line) of the Poisson pdf when Equation 1 was fitted to the frequency distributions of response rates in each 5-s time bin. For both groups, both the mean and weight of the Poisson changed over time bins, in close step with the change in response. **C** shows the mean Bayesian Information Criterion (BIC) values obtained when four different versions of Equation 1 were fitted to the frequency distribution of response rates across time bins. In one analysis, both μ and the weight (w) varied as free parameters during the fitting operation; in another analysis, μ was fixed and w varied as a free parameter; in the other two analyses, w was fixed, either at 1 or at a single value that optimized the fit of Equation 1 across all time bins, while μ varied as a free parameter. Note that more negative values for BIC indicate better evidence for that function given the data. **D** and **E** plot, for each rat in Groups FT20 (**D**) and FT40 (**E**), the BIC value for the 2-parameter model (free m , free w) against the BIC values for two of the 1-parameter models (free w , fixed m ; and free m , fixed w). Points above the diagonal indicate that the BIC for the 2-parameter model is more negative than the BIC for the 1-parameter model.

Among the analyses in which one parameter was fixed, there were sizeable differences in BIC favoring the model in which w was free and μ was fixed over the analyses in which μ was free and w was fixed. When w was set to 1, the BIC scores were 315 units higher for Group FT20, and 540 units higher for Group FT40, than when w was free. These differences equate to very large odds (more than 10^{68} :1) favoring the analysis in which w was free. When w was fixed at a value that optimized the fits to the response distribution data, the BIC scores were 152 (Group FT20) and 200 (Group FT40) higher than when w was free. These differences also equate to very large odds (more than 10^{32} :1) favoring the analysis in which w was free. The difference in BIC scores between the two types of model can be seen for each individual rat in Figures 6D and 6E—comparing the height of the orange triangles and the blue circles shows that the model with fixed w and free μ was superior (lower) in the majority of cases.

All of the comparisons described above were repeated using BIC scores calculated from the same data but excluding the first 5-s bin after CS onset. This was done in case response rates in that bin were affected in a specific way by the CS onset. As in Experiment 1, the pattern of BIC scores from this new analysis was virtually identical to that obtained from the full set of data.

Finally, the data from both groups were re-analyzed using a function that describes the weighted average of a high response state and a baseline (pre-CS) response state. As for Experiment 1, Equation 1 was first fitted to the frequency distributions of the pre-CS response rates. The resultant baseline function, $B(x)$, was then substituted for $Z(x)$ in Equation 1, before fitting $F(x)$ to the data from each time bin during the CS. This analysis produced mostly similar results to the analyses already described using $Z(x)$, but with two notable differences. First, for the data from Group FT40, there was now stronger evidence for the analysis with free w but fixed μ than for the analysis with free μ and

free w . For this comparison, the BIC was 17 units lower when w was free and μ was fixed than when both w and μ were free. This difference equates to odds of more than 4,000:1 in favor of the fixed w analysis. Second, for Group FT20, there was now weaker evidence for the analysis with free w but fixed μ than for the analysis with free μ but fixed w . In this case, the BIC was 88 units higher when μ was fixed and w free than when w was fixed and μ was free, which equates to odds more than 10^{19} :1 in favor of the latter model.

Discussion

The analyses presented here are consistent with the analyses from Experiment 1 in showing that the response distributions are better accounted for by changes in the relative weights of a positive response distribution and a zero response distribution than by changes in the mean of the positive response distribution. Nonetheless, as was also found for Experiment 1, the data provide the strongest support for a model in which both the weights and the mean of the response distributions vary when accounting for changes in responding across a trial. Some evidence at odds with these conclusions was obtained from the final analysis that used baseline (pre-CS) response distributions instead of the zero distribution when modeling the two response states. When this analysis was applied to the data from Group FT40, the evidence was strongest for the model in which all changes in responding across the CS were accounted for exclusively by a change in the weight of the two response distributions—the evidence for this model was even greater than the evidence for the model in which changes to both the weight and the mean of the distributions were used to account for the data. However, this improvement in the evidence for the former model in Group FT40 was offset by a decrease in evidence for that same model in Group FT20. In that group, the evidence was weakest for the model in which only the relative weights of the two response distributions changed across the trial, weaker even than the

evidence for the model in which the weight did not change at all but only the mean of the response distribution changed. This decrease in evidence for the free-weight fixed-mean model can be traced directly to its failure to adequately account for the very low response rates in the second half of the CS presentation in Group FT20. These response rates were below the pre-CS response rates, and therefore no matter how much the relative weighting shifted from the high response distribution to the baseline response distribution across the trial, the model could not produce response rates below the pre-CS rate.

An interesting detail in the present data is that response rates in Group FT20 fell below the pre-CS response rate for the later portion of the CS presentation, beyond the expected time of the US, suggesting that that portion of the CS had become inhibitory. Evidence of inhibition is a notable challenge for timing models, particularly in cases such as this when the evidence is of both inhibition and excitation for the same CS (see also Williams et al., 2008). Moreover, the finding had very significant implications for the analyses that used the pre-CS period to model the low response state (rather than the zero response distribution). In this case, the model that performed most poorly was the one that allowed only the weight parameter to vary while fixing the mean of the high response distribution. Indeed, that two-state model could not account at all for the sub-baseline response pattern during the later part of the CS because, in such a model, response rates can only lie within the range between the baseline rate and the high response rate. This result highlights complexities regarding the choice of an appropriate baseline response state, and argues in favor of using a zero response distribution for the baseline response state. An implication of this is that pre-CS responding is not a true baseline state but reflects occasional transitions into a positive response state, indicative of some weak but positive predictive value of the pre-CS period.

General Discussion

The present experiments used the peak procedure to measure changes in responding arising from conditioning protocols with fixed or variable CS-US intervals. Food was delivered either 20 s (Group FT20, Experiment 2), 30 s (Group FT30, Experiment 1), 40 s (Group FT40, Experiment 2), or at a variable time (Group VT30, Experiment 1) after onset of the CS; response rates were measured continuously during non-reinforced presentations of the CS that extended well beyond the time of food delivery on the reinforced trials. When food was delivered at one of the three fixed times, the rats' response rates rose steadily during the initial portion of the CS presentation, peaking at the expected time of food delivery, and then decreased again as the CS presentation extended past that time. This timing of the CR replicates many previous demonstrations of response timing (Church et al., 1994; Davis et al., 1989; Kehoe & Joscelyne, 2005; S. Roberts, 1981; W. A. Roberts et al., 1989; Smith, 1968; Williams et al., 2008), but contrasts with the pattern of responding observed when the CS-US interval varied from trial to trial. In the latter case, response rates rose sharply at the start of the CS, and remained elevated as time elapsed during the CS. This pattern is consistent with the fact that the rats given training with variable CS-US intervals cannot learn to anticipate the US at a specific time, and therefore maintain a uniform level of responding across the CS (Harris & Carpenter, 2011; Harris et al., 2011).

In addition to measuring how mean response rates change as time elapses during a trial, the work presented here examined how the distribution of response rates changes at different time points within a trial. The motivation behind this analysis was to test two contrasting descriptions of response timing. One approach assumes that response rates change in a continuously graded manner during the CS, as might reflect the continuous change in strength of an underlying association between the CS and US (Desmond & Moore, 1988; Sutton & Barto,

1981, 1990). A simple version of this account predicts that responding should be described by a single frequency distribution of response rates that shifts during a trial, such that, at any time during the trial, the mean of the frequency distribution equals the mean response rate. When the frequency distributions of response rates were analyzed in the present data, this model received the least support compared to other models considered here. In the large majority of instances, even when response counts were sampled in the time bin when responding was at its peak (such as shown in Figure 3F), the distribution of response counts did not conform to a single unimodal distribution. Rather, responses counts were bimodally distributed, with one mode centred on zero responses and the other described by a positive distribution with mean and variance greater than one.

The bimodal nature of the response distributions observed here suggest that the rats' behavior is best characterized using two response states, such that rats spend some proportion of their time in a no-response state and some proportion of their time in a positive response state. One such model contends that the probability of being in each state does not change across the length of a trial, but the strength of the positive response state does change. Changes in the strength of the response state would produce a systematic shift in the position of the higher response distribution, and as such this model often provided good fits to the observed response distributions.

However, the particular two-state model just described was consistently outperformed by another two-state model that made the opposite assumptions—that the strength of the positive response state does not change within a trial, but the probability of being in that state versus the no-response state does change. This latter model is in keeping with an abrupt switch between responding and not responding that is typically assumed by timing models (Gibbon, 1977; Guilhardi et al., 2007). It is also consistent

with the assumption made by these models that the relative weight of each component to the bimodal frequency distribution changes at different time points within the trial. Therefore the results presented here offer more support for that account of how responses change within a trial than for an account that explains timing as a continuous change in associative strength.

The goal of the present work was to compare descriptions of response timing in terms of either a shift in the location of a single response distribution versus a shift in weighting between two static distributions. Nonetheless, the analytic approach used here allowed those two descriptions to be combined in a hybrid model. Indeed, the model that received the greatest support is one that attributes the timing of responses within the trial to changes in both the probability of being in the response state and the strength of that response state. Not only did this combined model provide better fits to the response distributions in each group of rats, but consistently won out when differences in the Bayesian Information Criterion (BIC) were used to compare the amount of evidence for each model. Even for the comparison that produced the smallest difference in mean BIC values, the evidence for the combined model was stronger than either of the alternative models.

The theoretical approaches contrasted here describe changes in responding during a trial as either the function of a continuous strength variable or a binary decision variable, but do not offer a description of responding that combines both functions. How might timing theories of conditioning, or theories based on associative strength, accommodate the present evidence that changes in responding during a trial reflect a shift in both the probability of being in a response state and the strength of responding in that state? A combination of these factors has been incorporated in the modular extension of packet theory (Guilhardi et al., 2007). In this theory, bouts of responding are initiated as the output of a decision process. On any individual trial, a pattern memory that encodes the time of

reinforcement relative to CS onset on previous trials is compared with a subjective representation of time since CS onset in the current trial. When this comparison reaches a threshold, a binary change in response state is multiplied by a memory strength variable to produce the decision to initiate responding. Thus, the decision to respond is the product of both a binary state, related to the proximity of anticipated reinforcement, and a continuous strength variable. However, in this model, memory strength increases only when the reinforcer is delivered (according to a learning rule equivalent to that proposed by Rescorla & Wagner, 1972), and otherwise decreases within a trial. Therefore, the theory does not allow the memory strength variable to contribute to the continuous rise in responding across a trial leading up to the anticipated moment of reinforcement. From the perspective of timing models, the present results suggest that a comparison between currently elapsed time in the trial and the remembered time of reinforcement affects responding in two ways, via a binary decision to start/stop responding, and a continuous shift in response strength as a function of proximity to the expected time of reinforcement. That is, when remaining time to reinforcement reaches a decision threshold, the animal starts responding, and the vigor of its response continues to change as the time of reinforcement approaches.

From the perspective of associative strength models of conditioning, responding is a function of a graded strength variable which is itself a continuous function of distance from reinforcement. However, the evidence presented here indicates that the relationship between responding and associative strength must also be subject to a response threshold, such that responding only starts when associative strength exceeds a threshold, and stops when associative strength falls below a threshold. Arguments for the existence of response thresholds on conditioning strength go back to Hull (1943) and Spence (1956) when accounting for non-linearity in learning curves,

and could easily be used to account for evidence of abrupt non-linear changes in responding within a trial. However, an added complexity for any account of this sort is the requirement that the time at which the threshold is exceeded must vary from trial to trial, otherwise the average response rate should show a sharp discontinuity at the regular time when associative strength exceeds the threshold. This could be achieved either by assuming stochastic variation in the threshold value, or variation in the sequence of units that acquire associative strength (e.g., Machado, 1997).

References

- Brackney, R. J., Cheung, T. H. C., Neisewander, J. L., & Sanabria, F. (2011). The isolation of motivational, motoric, and schedule effects on operant performance: A modeling approach. *Journal of the Experimental Analysis of Behavior, 96*, 17-38.
- Cheng, K., & Westwood, R. (1993). Analysis of single trials in pigeons' timing performance. *Journal of Experimental Psychology: Animal Behaviour Processes, 19*, 56-67.
- Church, R. M., Meck, W. H., & Gibbon, J. (1994). Application of scalar timing theory to individual trials. *Journal of Experimental Psychology: Animal Behavior Processes, 20*, 135-155.
- Davis, M., Schlesinger, L. S., & Sorenson, C. A. (1989). Temporal specificity of fear conditioning: Effects of different conditioned stimulus-unconditioned stimulus intervals on the fear-potentiated startle effect. *Journal of Experimental Psychology: Animal Behavior Processes, 15*, 295-310.
- Desmond, J. E., & Moore, J. W. (1988). Adaptive timing in neural networks: The conditioned response. *Biological Cybernetics, 58*, 405-415.
- 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.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review, 107*, 289-344.
- Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84*, 279-325.

- Gibbon, J., & Church, R. M. (1990). Representation of time. *Cognition*, *37*, 23-54.
- Guilhardi, P., Yi, L., & Church, R. M. (2007). A modular theory of learning and performance. *Psychonomic Bulletin & Review*, *14*, 543-559.
- Harris, J. A. (2011). The acquisition of conditioned responding. *Journal of Experimental Psychology: Animal Behavior Processes*, *37*, 151-164.
- 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.
- 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.
- Harris, J. A., Patterson, A. E., & Gharaei, S. (2015). Pavlovian conditioning and cumulative reinforcement rate. *Journal of Experimental Psychology: Animal Learning and Cognition*, in press.
- Hull, C. L. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts.
- Kehoe, E. J., & Joscelyne, A. (2005). Temporally specific extinction of conditioned responses in the rabbit (*Oryctolagus cuniculus*) nictitating membrane preparation. *Behavioral Neuroscience*, *119*, 1011-1022.
- Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review*, *104*, 241-265.
- Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex*. (G. V. Anrep, Trans.). New York: Dover.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory*. (pp. 64-99). New York: Appleton-Century-Crofts.
- Roberts, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes*, *7*, 242-268.
- Roberts, W. A., Cheng, K., & Cohen, J. S. (1989). Timing light and tone signals in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *15*, 23-35.
- Rogers, A. E. (1979). Nutrition. In H. J. Baker, J. R. Lindsey & S. H. Weisbroth (Eds.), *The laboratory rat* (Vol. 1, pp. 123-152). New York: Academic Press.
- Schneider, B. A. (1969). A two-state analysis of fixed interval responding in the pigeon. *Journal of the Experimental Analysis of Behavior*, *12*, 677-687.
- Shull, R. L., Gaynor, S. T., & Grimes, J. A. (2001). Response rate viewed as engagement bouts: Effects of relative reinforcement and schedule type. *Journal of the Experimental Analysis of Behavior*, *75*, 247-274.
- Shull, R. L., Grimes, J. A., & Bennett, J. A. (2004). Bouts of responding: The relationship between bout rate and the rate of variable-interval reinforcement. *Journal of the Experimental Analysis of Behavior*, *81*, 65-83.
- Smith, M. C. (1968). CS-US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology*, *66*, 679-687.
- Spence, K. W. (1956). *Behavior theory and conditioning*. New Haven: Yale University Press.
- Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, *88*, 135-171. doi: 10.1037/0033-295X.88.2.135
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497-537). Cambridge, MA: Bradford Books/MIT Press.
- Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review*, *14*, 779-804.
- Williams, D. A., Lawson, C., Cook, R., Mather, A. A., & Johns, K. W. (2008). Timed Excitatory Conditioning Under Zero and Negative Contingencies. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*(1), 94-105.

Appendix

Two sets of simulated data were generated, one mimicking a low-high-low (2-state) discontinuous response pattern, the other created with a continuous change in the probability of responding across each trial. Both data sets comprised 1000 trials of 500 time steps, each step could take a value of either 0 (no response) or 1 (response).

For the Continuous data set, each step on each trial was initially set to 0 and then randomly changed to 1 with a probability that varied continuously across the length of the trial. The probability of a response at each time step was described by a normal (Gaussian) probability density function centered on the middle of the trial, with a standard deviation of 167, and scaled so that the minimum probability (at the first and last time steps) was 0.08, and the peak probability (at the 250th time step) was 0.25.

For each trial of the 2-state data set, two transition times, t_1 and t_2 , were selected randomly from each half of the same normal probability distribution described for the Continuous data set. All time steps before t_1 and

after t_2 were set to 0; all steps between t_1 and t_2 were randomly changed to 1 with a fixed probability of 0.25.

In both data sets, each trial was compressed to 50 time steps, with the value at each step being the sum of 10 consecutive steps binned from the original 500 steps. Therefore, the response rate at each of the 50 steps could take a value between 0 and 10.

Both data sets were analysed using an algorithm described by Church (1994). The algorithm performed an exhaustive search of each trial to find values of t_1 and t_2 that maximized the value, d , defined as

$$d = i_1(r - r_1) + i_2(r_2 - r) + i_3(r - r_3)$$

where i_1 is the interval from the start of a trial to t_1 , i_2 is the interval between t_1 and t_2 , i_3 is the interval from t_2 to the end of the trial, r is the mean response rate over the whole trial, and r_1 , r_2 and r_3 are the mean response rates during intervals i_1 , i_2 , and i_3 . This maximizes the difference between the high response state (r_2) and the low response states (r_1 and r_3), with each weighted by the time spent in that state.