Pre-publication copy.

Published in: *The Journal of Experimental Psychology: Animal Learning & Cognition*, vol 46 (2020), pp 256-272.

©American Psychological Association, 2020. This paper is not the copy of record and may not exactly replicate the authoritative document published in the APA journal. Please do not copy or cite without author's permission. The final article is available, upon publication , DOI 10.1037/xan0000242

Pavlovian conditioning under partial reinforcement:
The effects of non-reinforced trials versus cumulative CS duration

Justin A. Harris
*School of Psychology,
University of Sydney*

&

Mark E. Bouton
*Dept of Psychological Science,
University of Vermont*

Address for correspondence:

Justin Harris
School of Psychology
The University of Sydney,
Sydney 2006
Australia

Running head: Trials versus Time under partial reinforcement

Acknowledgements. This research was supported by funding from the Australian Research Council, grant DP 150101274.

Abstract

A core feature of associative models, such as those proposed by Allan Wagner (Rescorla & Wagner, 1972; Wagner, 1981), is that conditioning proceeds in a trial-by-trial fashion, with increments and decrements in associative strength occurring on each occasion that the conditioned stimulus (CS) is present either with or without the unconditioned stimulus (US). A very different approach has been taken by theories that assume animals continuously accumulate information about the total length of time spent waiting for the US both during the CS and in the absence of the CS (e.g., Gallistel & Gibbon, 2000). Here we describe three experiments using within-subject designs that tested between trial-based and time-accumulation accounts of the acquisition of conditioned responding using magazine approach conditioning in rats. We found that responding was affected by the total (cumulative) duration of exposure to the CS without the US rather than the number of trials on which the CS occurred without the US. We also found that exposure to the CS without the US had the same effect on conditioning whether that exposure occurred shortly (60 s) before each CS-US pairing or whether it occurred long (240 s) before each pairing. These findings are more consistent with time-accumulation models of conditioning than trial-based models like the Rescorla-Wagner model and Wagner's (1981) Sometimes Opponent Process model. We discuss these findings in relation to other evidence that favours trial-based models rather than time-accumulation models.

Trial-based theories of associative learning, such as the highly influential model proposed by Rescorla and Wagner (Rescorla & Wagner, 1972), have enjoyed a great deal of success in explaining how animals acquire conditioned responding. However, because these models treat the trial as the unit of learning, they have difficulty accounting for the influence of time, both within and between trials, on conditioning. It has long been recognized that Pavlovian conditioning is affected by the length of the temporal interval between presentation of the conditioned stimulus (CS) and the unconditioned stimulus (US; e.g., Ellison, 1964; Gibbon, Baldock, Locurto, Gold, & Terrace, 1977; Schneiderman, 1966; Smith, 1968; Smith, Coleman, & Gormezano, 1969) and also by the interval separating each pairing of the CS and US (e.g., Barela, 1999; Barnet, Grahame, & Miller, 1995; Domjan, 1980; Gibbon et al., 1977; Gibbon, Farrell, Locurto, Duncan, & Terrace, 1980; Lattal, 1999; Li & Westbrook, 2008; Morris, Furlong, & Westbrook, 2005; Sunsay & Bouton, 2008; Sunsay, Stetson, & Bouton, 2004). Indeed, there is evidence that, under some conditions, the acquisition of Pavlovian responses is specifically determined by the ratio of these intervals (Gibbon et al., 1977). Trial-based models are not equipped to account for these effects of time on conditioning. This limitation has motivated development of associative models that update associative strength in a more-or-less continuous manner within and between trials (Desmond & Moore, 1988; Harris & Livesey, 2010; Kehoe, 1988; Killeen & Fetterman, 1988; Machado, 1997; McLaren & Mackintosh, 2002; Sutton & Barto, 1981, 1990; Vogel, Brandon, & Wagner, 2003; Wagner, 1981).

One of the most important examples of a continuous associative model is the Standard Operating Procedures (SOP) model proposed by Allan Wagner (1981). According to SOP, the presentation of a stimulus, S , activates a node of elements representing S from an inactive state to a primary activation state (A_1) with a probability p_1 . Each element making up the node will

decay from A1 into a secondary activation state (A2) with probability p_{d1} at each time t . From A2, elements will decay with probability p_{d2} back to the inactive state. Whenever the elements of two stimuli, S1 and S2, are active in A1 at the same time, an excitatory connection between them is strengthened by an amount proportional to the number of conjointly activated elements. Whenever the elements of S1 are active in A1 while the elements of S2 are in A2, an inhibitory connection from S1 to S2 is strengthened. If an excitatory connection exists from the elements of S1 to the elements of S2, activation of S1 into A1 or A2 can prime the elements of S2 directly into A2; the number of S2 elements primed into A2 will be proportional to the strength of the excitatory association between S1 and S2. Conversely, if there is an inhibitory connection from S1 to S2, the presence of S1 will reduce the probability that elements of S2 can be primed into A2.

The set of assumptions just described allows SOP to explain how conditioning is affected by the CS-US interval in both trace conditioning, when the US occurs some amount of time after the end of the CS presentation, and delay conditioning, when the US coincides with offset of the CS. Each onset of the CS causes a steep rise in the number of CS elements in A1, with the number gradually decreasing as elements decay from A1 to A2. Each presentation of the US likewise causes a steep rise in the number of US elements in A1 before they begin to decay to A2. The number of CS elements that acquire excitatory connections with the US will be proportional to the number of CS elements still in A1 at the time that US elements enter A1. Trace conditioning occurs as long as some CS elements are still in the A1 state when the US is presented, even if the CS itself had terminated before the US occurred. Nonetheless, a long trace interval will produce weaker conditioning than a short interval because, over the longer interval, more CS elements will have decayed from A1 before the US is presented. The same effect of CS-US interval is

predicted for delay conditioning. In that case, there will be fewer CS elements in A1 at the end of a long CS presentation than at the end of the short CS presentation because CS elements gradually decay from A1 to A2 even while the CS is present, and while those elements are in the A2 state they cannot be reactivated into A1. According to SOP, an additional process will reduce the strength of conditioning over longer CS-US intervals. Once some excitatory connection has been established between the CS elements and the US, during each trial while the CS is present, but before the US has occurred, the CS elements in A1 will prime the US into A2, creating an opportunity for inhibitory connections to be strengthened between the CS and US. Long CS-US intervals will provide more opportunity to strengthen inhibitory connections than short CS-US intervals, which will result in a lower asymptote of conditioning strength for long CS-US interval than short intervals.

SOP can also explain why long inter-trial intervals produce better conditioning than short intervals by appealing to the gradual decay of elements from A2 to the inactive state. The length of the interval will affect how many elements of the CS and US are still active in A2 from one trial to the next, with more stimulus elements remaining in A2 across shorter intervals. As a result, when the interval is short, fewer CS and US elements will be activated into A1 when those stimuli are presented on each subsequent trial, and this will limit the increase in excitatory strength between the CS node and US node on those trials. Across long intervals, CS and US elements decay back to the inactive state and therefore can be successfully activated into A1 on each subsequent trial. Long intervals also reduce (extinguish) the excitatory association between the conditioning context and both the CS and US, and as a result the context is less able to prime the CS and US elements into A2 which would otherwise reduce A1 processing of those stimuli (Sunsay & Bouton, 2008).

The evidence that conditioning is very sensitive to temporal intervals has also motivated a very different conceptualization of the learning process, one that represents a significant challenge to the approach characterized by models like SOP. This alternative view rejects the idea that learning involves the strengthening of excitatory or inhibitory connections, but instead argues that animals learn specifically about the temporal intervals between events and use their estimates of these intervals to make decisions about whether and when to respond in anticipation of a US (Balsam, Drew, & Gallistel, 2010; Balsam & Gallistel, 2009; Gallistel & Gibbon, 2000; Gibbon & Balsam, 1981; Gibbon, Church, Fairhurst, & Kacelnik, 1988). The best-known example of this class of model is Rate Estimation Theory (RET; Gallistel & Gibbon, 2000) which argues that animals keep track of how frequently the US occurs when the CS is present and compare this with an estimate of how frequently the US occurs when the CS is absent. This comparison can be expressed as the ratio of the accumulated time that the animal has waited for the US when the CS was absent (the inter-trial interval, I) and the accumulated time it has waited when the CS was present (T). When the I/T ratio exceeds a decision threshold, the animal will respond in anticipation of the US. This account simultaneously explains why conditioned responding is weaker when the CS-US interval (T) is long or when the I is short, and specifically predicts that responding will be invariant over changes in I and T as long as the I/T ratio is constant (Balsam et al., 2010; Balsam & Gallistel, 2009; Gallistel & Gibbon, 2000).

Bouton and Sunsay (2003) conducted a series of experiments designed to test between the predictions of associative models, such as SOP, and time-accumulation models like RET. Several of these experiments used what we will call the “split-trials” design to pit the impact of trials against accumulated time on the acquisition of conditioned responding (magazine approach in rats). In one such experiment (their Experiment 2) they compared the acquisition of conditioned

responding between two groups of rats: One group was trained with a 30-s CS that was paired with delivery of a food pellet US at the end of every presentation (consistently reinforced), and the other group was trained with a 10-s CS that was presented 3 times as often but followed by the US on only 1/3rd of trials (partially reinforced). The schedule of the 10-s CS is, in effect, the same as the 30-s CS but where each 30-s CS presentation has been split into 3 separate presentations without increasing the accumulated duration of CS time or number of trials that included the US. According to time-accumulation models like RET, this difference should have no effect on acquisition of conditioning because the I/T ratio is matched in each case. Bouton and Sunsay observed that the rats trained with the 10-s partially reinforced CS were slightly slower to acquire responding but their responding reached a higher asymptote than for the rats trained with the 30-s consistently reinforced CS, however these differences were not significant. More recently, Harris, Patterson and Gharaei (2015) replicated their procedure but using a within-subjects design (rats were trained with both a 30-s CS that was consistently reinforced and a 10-s CS that was partially reinforced). They also observed slower acquisition but higher responding at asymptote for the 10-s CS, and in this case the interaction between CS and session was significant.

The results reported by Bouton and Sunsay (2003) and Harris et al. (2015) are inconsistent with the prediction by RET that the split-trial design should not affect conditioning. However, Harris et al. also provided evidence from many other experiments using the split-trials design that was more consistent with RET's prediction. In those experiments, Harris et al. found that the acquisition of rats' responses to a short partially reinforced CS was virtually identical to that of a long consistently reinforced CS if those CSs had the same overall reinforcement rate. These particular experiments differed from the experiments described in the preceding paragraph in at

least one important way: They randomly varied the CS-US interval from trial to trial (e.g., a CS with mean duration of 30 s could vary between 2 and 58 s, while a 10-s CS could vary between 2 and 18 s), whereas CS-US intervals were fixed in the experiments reported by Bouton and Sunsay as well as the one experiment by Harris et al. which found an effect of the split-trial design. As discussed next, this difference between fixed and variable intervals would have affected the pattern of the rats' responding within each trial.

There is a great deal of evidence that conditioned responding changes across the course of a fixed CS-US interval, with response rates peaking close to the end of the interval (the expected time of US delivery; Davis, Schlesinger, & Sorenson, 1989; Drew, Zupan, Cooke, Couvillon, & Balsam, 2005; Harris, 2015; Kehoe & Joscelyne, 2005; Roberts, Cheng, & Cohen, 1989; Smith, 1968; Williams, Lawson, Cook, Mather, & Johns, 2008), an effect confirmed in Experiment 3 of Harris et al. (2015). Therefore, any comparison of the mean response rates to short and long CSs with fixed durations could be affected by differences in the timing of the responses. This should not be a problem if timing accuracy and timing precision scale exactly with the CS-US interval. However, while scalar invariance has been reported for the timing of conditioned responses (Gibbon, 1977; Gibbon & Church, 1990), it is not always observed. For example, Gallistel, King, and McDonald (2004) found that the onset of a bout of nose poke responses in mice did not scale uniformly with the CS-US interval. When the CS-US interval was short (5 or 15 s), mice began responding much sooner, as a proportion of CS-US interval, than when the interval was long (45 s). To test for such differences in the experiments described above, we have re-analysed the response times to the fixed duration CSs in Experiment 3 of Harris et al. (2015). We fitted Gaussian curves to the distribution of response rates (in 1-s bins) during the 10-s and 30-s fixed CSs, using data from the last 10 sessions (Sessions 21-30 of conditioning) which is when mean

responding to the 10-s CS was higher than to the 30-s CS. This analysis revealed significant deviations from scalar invariance, which is itself problematic for timing-based models like RET. The timing of the peak of the Gaussian curve, as a proportion of the CS-US interval, was significantly earlier for the 10-s CS (mean = 0.88, or 8.8 s into the 10-s interval) than for the 30-s CS (mean = 1.1, or 33 s)¹, $t(15) = 3.59$, $p = .003$. Responding also began to emerge at an earlier point during the 10-s CS than the 30-s CS: responding reached 50% of the maximum rate at a significantly earlier time point, relative to the CS-US interval, for the 10-s CS (mean = 0.3) than for the 30-s CS (mean = 0.4), $t(15) = 2.68$, $p = .017$. This difference in timing would contribute to make the mean response rate to the 10-s CS higher than the mean response rate to the 30-s CS.

While response timing with fixed CS-US intervals may complicate interpretation of differences in the mean response rate between short and long CSs, this problem can be avoided by using variable CS-US intervals. Variable intervals produce uniform levels of responding across the length of each trial (Harris & Carpenter, 2011; Harris, Gharaei, & Pincham, 2011), presumably because the probability of the US is uniform at each moment in the trial. For this reason, Harris et al. (2015) argued that variable CS-US intervals should be used in preference to fixed intervals to test the role of accumulated time of CS exposure on conditioning. Given the results from their experiments that used variable CS-US intervals, they concluded that rats do learn about the overall reinforcement rate of a US, based on the accumulated duration of CS exposure per US, rather than learning about individual trials.

¹ In most studies of response timing with fixed CS-US intervals, responding peaks slightly before the expected US time, rather than occurring slightly after the US time as estimated here for the 30-s CS. One reason for this may be that, in the experiment we have analysed here, all presentations of the 30-s CS were reinforced and the CS never continued beyond the 30-s CS-US interval. By contrast, studies showing response rates that peak before the US time use a procedure (the peak procedure) in which response timing is measured on a subset of test trials in which the US is not delivered and the CS is maintained well beyond the expected time of the US. From these long unreinforced trials, the animal learns to quit responding at some point after the expected time of US delivery, and this learning may shift the peak response rate to be slightly earlier than the true US time. In this regard, it is notable that responding did peak before the expected US time for the 10-s CS that was partially reinforced.

While the results of Harris et al.'s (2015) experiments using variable CS-US intervals are problematic for simple trial-based associative models like the Rescorla-Wagner (1972) model, their implications for a continuous associative model like SOP are less clear. According to SOP, associative loss (an increase in strength of an inhibitory connection between the CS and US) will accrue continuously during presentations of the CS whether the trials are long or split into many short trials. On the face of it, this would suggest there should be no difference in conditioning using the split-trials design. However, SOP also assumes that the associability of the CS changes across the length of any trial, first increasing as CS elements are recruited into A1, then decreasing as the elements decay from A1 to A2 (as described earlier). Therefore, the split-trials design should produce a difference in conditioning, but the direction of the difference will depend on how quickly elements are recruited into A1 and decay out of A1. If the rate of recruitment into A1 is very slow (i.e., $p_1 < 0.1$ per second) such that the number of elements in A1 continued to rise beyond 10 s, then the mean associability of the CS would be higher over one 30-s trial than over three 10-s trials, and there would be greater inhibitory strength acquired over the single long trial than three short trials. Conversely, if initial recruitment were more rapid ($p_1 \geq 0.2$), such that A1 activity peaked before 10 s, the mean associability of the CS should be higher over multiple short trials than one long trial, and more inhibitory strength would be acquired for the short CS than the long one. In other words, depending on the rate of recruitment into A1, the net loss of associative strength (increase in inhibitory strength) could be smaller or larger across split trials, and conditioning of the short partially-reinforced CS could be stronger or weaker than the long consistently-reinforced CS in the split-trials design.

SOP's predictions about conditioning strength in the split trials design are further complicated by the fact that the difference in associability between short and long trials also

means that the associability of the short and long CSs will differ at the time of the US. If recruitment is rapid, such that associability is lower at the end of long trials than short trials, then increments in excitatory strength will be smaller on reinforced trials of the long CS than the short CS. If recruitment is very slow, such that there are more elements in A1 at the end of a long trial than at the end of a short one, the US would produce larger increments in excitatory strength on long trials than on short trials. Note these effects will serve to negate the differences in amount of inhibitory conditioning that accrues across multiple short trials versus a few long trials (described in the preceding paragraph). In other words, according to SOP, splitting a long trial into several short trials could increase inhibitory conditioning during the CS while also increasing excitatory conditioning at the time of US delivery, or decrease inhibitory conditioning during the CS while also decreasing excitatory conditioning by the US. Exactly how the split-trials procedure will affect these two processes depends on the specific choice of model parameters, and therefore SOP doesn't make a committed prediction about the impact of the split-trial design on the acquisition of conditioning.

Some of the difficulty in specifying what SOP predicts in a split-trials design can be circumvented using another of the designs introduced by Bouton and Sunsay (2003). Rather than testing the effect of splitting a long reinforced trial and into several short trials, this design tests the effect of splitting a long non-reinforced trial into several short trials. For example, in their Experiment 3, Bouton and Sunsay trained three groups of rats with 10-s reinforced presentations of a CS intermixed among non-reinforced presentations of the same CS. These non-reinforced trials were 40-s long for two groups or split into four short (10-s) trials for the third group. As in the previous split-trial design, the accumulated time of CS exposure is matched between the three groups, which means that time-accumulation models like RET once again predict no

difference in the acquisition of responding. This design also provides a clearer test of SOP's prediction because it removes any impact of trial length on increments in excitatory strength during reinforcement. In this case, differences in CS processing when trials are split into many short presentations will only affect increases in inhibitory strength during the CS when the US is absent. For example, if recruitment into A1 is sufficiently rapid that the number of elements in A1 peaks within the first 10-s of a CS presentation, then SOP predicts that four 10-s non-reinforced trials will reduce conditioning more than a one 40-s non-reinforced trial. Consistent with this prediction, but not with the prediction derived from time-accumulation models, Bouton and Sunsay found that responding was weaker in the group given many split trials, even when, for all groups, they only analysed response rates during the 10-s CS presentations that ended in reinforcement.

In the experiment just described, Bouton and Sunsay (2003) also observed a difference in conditioning between the two groups given the 40-s non-reinforced trials. In the design, the two groups differed in when those non-reinforced trials occurred. For one group, each non-reinforced trial occurred early in the interval between reinforced trials (60 s after the previous trial, and 240 s before the next reinforced trial) whereas for the other group each non-reinforced trial occurred late in that interval (240 s after the previous trial and 60 s before the next trial). Responding was lower for the second group given the non-reinforced presentation late in the interval. A similar effect was obtained in Bouton and Sunsay's final experiment. This difference in responding between the groups as a function of when the non-reinforced trial occurred is also consistent with SOP but not with a time-accumulation model like RET. According to SOP, if each presentation of the CS leaves residual activation of that stimulus in A2, this will diminish processing of the CS (by reducing access to A1) on the subsequent presentation. As a result, non-

reinforced presentations that occur close to the reinforced trial will reduce conditioning on the reinforced trial more than non-reinforced trials that occur well before the reinforced trial. Again, this prediction depends on parameters—if decay from A2 is very fast, such that there is no difference in the number of elements in A2 after 60 s versus 240 s, then SOP would not be able to explain the effect reported by Bouton and Sunsay. Conversely, if decay from A1 is very slow, there could be a greater accumulation of elements in A2 across the 240-s interval than the 60-s interval, leading to a reverse prediction.

Further evidence for SOP's account was provided by Sunsay, Stetson, and Bouton (2004). In their Experiment 1, rats were trained with two CSs, a tone (T) and a light (L), both reinforced on 50% of trials. In one group ("Same"), each reinforced trial of T or L was preceded 60-s earlier by a non-reinforced presentation of the same stimulus (i.e., T before T, and L before L). In a second group ("Different"), each reinforced trial was preceded by a non-reinforced trial of the other stimulus (i.e., L before T, T before L). Even though reinforced trials were preceded by non-reinforced at the same time point in the US-US interval for both groups, conditioning to the two CSs proceeded more slowly in the Same group than in the Different group. This result is consistent with the CS-priming mechanism proposed by SOP because, in the Same group, A1 activation of the CSs would have been reduced on each reinforced trial due to residual activation of that CS in A2 from the preceding non-reinforced trial. RET, on the other hand, does not anticipate any differential impact of the sequencing of the non-reinforced trials on conditioning.

In this paper, we present three experiments that revisit the two findings reported by Bouton and Sunsay (2003) that are inconsistent with time-accumulation models but are consistent with the temporally continuous associative process described in Wagner's SOP model. Experiments 1 and 2 address an alternative explanation of Bouton and Sunsay's (2003) finding

that four 10-s nonreinforced trials had a stronger negative impact on conditioning than a single 40-s nonreinforced trial. This alternative explanation focuses on the fact that Bouton and Sunsay used fixed CS-US intervals, which means that the rats would have learned to time their expectation of the US based on the fixed CS-US interval. Given that the reinforced trials were all 10-s long, the rats would have expected the US to occur 10 s after CS onset, and their expectation of the US would wane during the last 30 s of a 40-s non-reinforced trial. In other words, only part of a 40-s non-reinforced trial would impact on the rat's expectation of the US within the first 10-s of a trial. This type of analysis could explain why conditioning strength was less affected by one 40-s non-reinforced trial than four 10-s non-reinforced trials without invoking the mechanisms described in SOP. It is worth mentioning that, as the authors noted, this explanation of Bouton and Sunsay's finding is not based on any mechanism described in RET which, despite its emphasis on temporal intervals, is unable to account for the finding. Since this explanation depends on the rats being able to time the arrival of the US as a result of training with fixed CS-US intervals, it would not apply to a case where rats had been conditioned using variable CS-US intervals, such that their expectation of the US was constant across each moment of the trial. Experiment 1 and 2 of the present series address this point. Experiment 3 then revisits the other finding reported by Bouton and Sunsay, that conditioning is affected by the temporal proximity of a non-reinforced trial to the next reinforced trial.

Experiment 1

Experiment 1 compared rats' acquisition of magazine responding to two CSs that each signaled food on a partial reinforcement schedule and had the same overall (cumulative) rate of reinforcement. Reinforced trials with the two CSs were the same, but the CSs differed in the number and length of non-reinforced trials, along the lines of the split-trials design introduced

by Bouton and Sunsay (2003) in their Experiment 3. For one CS, B, each session included 5 48-s non-reinforced trials, and for the other CS, C, each session included 15 16-s non-reinforced presentations. To prevent rats from timing when the US would occur, all reinforced trials varied randomly in length. According to time-accumulation models like RET, the acquisition of responding to B and C should not differ because the 15 non-reinforced trials of C would be equivalent to the 5 non-reinforced trials of B. According to a trial-based associative model like SOP, rats should respond less to C than to B because the many short non-reinforced trials of C will produce greater inhibitory conditioning than the few long non-reinforced trials given to B.

Trials with CSs B and C were intermixed among trials with two other CSs, A and D, and all four CSs had five reinforced trials per session with a mean CS-US interval of 24s (range 2-46s). The schedule of non-reinforcement of the four CSs conformed to a 2x2 factorial design. There were either five (CSs A and B) or 15 (C and D) non-reinforced trials per session; non-reinforced trials were either 16 s long (A and C) or 48 s long (B and D). As already noted, the key comparison is between B and C since they received identical reinforced trials, and the same cumulative amount of non-reinforced exposure each session (=240 s), but they differed in the number of non-reinforced trials (5 long vs 15 short). The design is summarized in Table 1.

CS	Reinforced trials	Non-Reinforced trials
Experiment 1		
A	5 x ~24 s [2 – 46 s]	5 x 16 s
B	5 x ~24 s [2 – 46 s]	5 x 48 s
C	5 x ~24 s [2 – 46 s]	15 x 16 s
D	5 x ~24 s [2 – 46 s]	15 x 48 s
Experiment 2		
A	8 x ~30 s [2 – 58 s]	8 x ~30 s [2 – 58 s]
B	8 x ~10 s [2 – 18 s]	40 x ~10 s [2 – 18 s]

Table 1. Summary of the design of Experiments 1 and 2. Rats were trained with 4 CSs (A to D) in Experiment 1 and with 2 CSs (A and B) in Experiment 2. The table shows the number and mean duration [and range] of reinforced trials and of non-reinforced trials of each CS within each session.

Methods

Subjects

Sixteen 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 (Techniplast™), measuring 40 x 46 x 40cm (length x width x height), located in the School of Psychology animal colony. The rats had unrestricted access to water in their home boxes. 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), and took at least 2 h to be eaten (but was usually finished within 3 h).

Apparatus

Rats were trained and tested in 16 Med Associates™ conditioning chambers (28.5 x 30 x 25 cm: height x length x depth). The end walls of each chamber were made of aluminum; the sidewalls and ceiling were Plexiglas™. The floor of the chamber consisted of stainless-steel rods, 0.5 cm in diameter, spaced 1.5 cm apart. Each chamber had a recessed food magazine in the center of one end wall, with an infra-red LED and sensor located just inside the magazine to record entries by the rat. A small metal cup measuring 3.5 cm in diameter and 0.5 cm deep was fixed on the floor of each food magazine. Attached to the food magazine was a dispenser delivering 45 mg food pellets (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 and created background noise (between 61 and 66 dB, depending on the chamber). Experimental events were controlled and recorded automatically

by computers and relays located in the same room. White noise (19 dB above background) was presented from a speaker mounted on the wall of each conditioning chamber above and to the left of the food magazine. A high-pitched tone (2.9 kHz, 17 dB above background) was provided by a piezo buzzer positioned on the floor of the sound-attenuating shell behind each conditioning chamber. A steady light (30cd/m²) was produced by an incandescent bulb mounted high on the back wall of the sound-attenuating shell. A flashing light (2 Hz; 3.0 cd/m²) was produced from a 3 × 3 array of white LED bulbs, positioned on the floor of the shell in front of conditioning chamber. The four stimuli were used as CSs A to D in counterbalanced fashion following a Latin Square design.

Procedure

Prior to the start of conditioning, rats received a 60-min magazine training session during which 20 food pellets were presented on a variable time (VT) 3-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 28 sessions. Each session contained five reinforced trials of each CS. These trials had a mean duration of 24 s, sampled from a uniform distribution between 2 s and 46 s, and termination of the stimulus coincided with delivery of a single food pellet. Each session also contained non-reinforced trials of each CS that had a fixed duration: There were five trials of A (16 s), five trials of B (48 s), 15 trials of C (16 s), and 15 trials of D (48 s). Trials of each type were randomly intermixed. The inter-trial interval varied randomly with a mean of 96 s, and the session lasted approximately 120 min. Photo-beam interruptions by head entry into the magazine were recorded during each CS and each 24-s pre-CS period.

Results and Discussion

All CSs received the same number and duration of reinforced trials per session, but differed in the number and length of non-reinforced trials. Therefore, we only analysed response rates recorded during reinforced trials to make sure the data for each CS was collected under comparable conditions. Figure 1A shows response rates to each CS across all 28 sessions. Both the number of non-reinforced trials per session (5 vs 15) and the duration of those trials (16 s vs 48 s) had an impact on the response rates. The most relevant comparison is between responding to B and to C given that these were matched for total duration of non-reinforced exposure. The figure suggests that responding to B was acquired more quickly than to C.

----- Insert Figure 1 about here -----

A 2x2x28 ANOVA was conducted on the data shown in Figure 1A. It confirmed that there was a significant main effect of number of non-reinforced trials per session, $F(1,15) = 45.46$, $p < .001$, $\eta^2_p = 0.75$, and length of non-reinforced trials, $F(1,15) = 26.96$, $p < .001$, $\eta^2_p = 0.64$, as well as a significant main effect of Session, $F(4.6,69) = 5.90$, $p < .001$, $\eta^2_p = 0.28$. Session interacted with both the number, $F(7.9,119) = 3.76$, $p = .001$, $\eta^2_p = 0.20$, and length of non-reinforced trials, $F(7.1,107) = 3.89$, $p = .001$, $\eta^2_p = 0.21$, but these factors did not interact with one another, $F(1,15) = 2.28$, $p = .152$, nor was there a significant 3-way interaction, $F(7,106) = 1.46$, $p = .189$.

We ran a separate 2x28 ANOVA to compare responding to B and C over the 28 sessions. This identified a significant difference overall between responding to B and responding to C averaged across sessions, $F(1,15) = 14.84$, $p = .002$, $\eta^2_p = 0.50$. There was also a significant main effect of session, $F(5.1,76) = 5.11$, $p < .001$, $\eta^2_p = 0.25$, that did not interact with CS, $F(7.9,117) = 1.05$, $p = .404$. As predicted by SOP, rats acquired less responding to the CS with three times as

many non-reinforced trials per session than to the CS with non-reinforced trials that were three times as long.

Figure 1B plots response rates to B and C in 2-s bins spanning the maximum length of the reinforced trials averaged over the last 25 sessions. We predicted that response rates would be uniform across the trial because we had scheduled the CS-US intervals to vary randomly from trial to trial. However, as is clear from Figure 1B, this did not occur. Instead, the rats showed different patterns of responding to the two CSs: response rates to C were initially low and rose gradually across the length of the trial, while response rates to B rose to a high level early in the trial and then decreased across the length of the trial. A 2x24 ANOVA on these data showed that there was a significant difference overall between responding to B and responding to C, $F(1,16) = 5.10$, $p = .039$, $\eta^2_p = 0.25$, and a significant main effect of Time, $F(4.6,69) = 4.42$, $p = .002$, $\eta^2_p = 0.23$, as well as a significant Time by CS interaction, $F(4.9,72) = 4.82$, $p = .001$, $\eta^2_p = 0.24$. This difference in the pattern of responding across the trial was very likely due to the difference in length of non-reinforced trials for B and C. That is, because there were many short non-reinforced trials with C, the probability that a trial would end without reinforcement was high early on in a presentation of C and decreased as the length of the trial extended beyond 16 s. Conversely, because non-reinforced trials with B were all long, the probability of a trial ending without reinforcement was low in the initial portion of any presentation but increased as the trial extended out towards 48 s. This interpretation implies that rats do not simply learn about the timing of the US relative to the onset of the CS, since this was the same for all CSs, but that they also learn about the timing of non-reinforcement. When the timing of the US is variable, such as in the present experiment, there is evidence that animals detect the absence of the US only at the end of a non-reinforced trial (Harris & Kwok, 2018). Thus, when non-reinforced trials are uniformly short relative to the average CS-US interval, as for CS C, animals may remember

that non-reinforcement occurs relatively close to the start of a trial and this reduces their response rate early in the trial but not later in a long trial. Conversely, when non-reinforced trials are uniformly long, as for CS B, animals may remember detecting non-reinforcement only after many seconds have elapsed in the trial, and this memory suppresses responding late in a long trial but not earlier.

The difference in pattern of responding across the length of the trial is problematic for our interpretation of the overall difference in response rates acquired to B and C. The problem arises because these reinforced trials varied randomly around a mean of 24 s, which means that some trials were short and others were long. If one averages across all reinforced trials, as done in Figure 1A, then the measured response rate will be heavily represented by responding that occurs in the first part of the trial because the early part of the trial is present on *all* trials, whereas the late portion of long trials is relatively under-sampled. At the extremes, the imbalance is clearest—all trials last at least 2 s whereas almost none last for 46 s, and therefore response rates in the first 2 s will be much more strongly represented in the overall mean. Therefore, the difference in mean response rates shown in Figure 1A might reflect a difference in associative strength or it could be entirely due to the difference in response timing shown in Figure 1B.

In sum, the split trials design in this experiment indicated that the acquisition of responding was more negatively affected by non-reinforced trials if they were presented as many short trials rather than a smaller number of long trial, consistent with a trial-based model like SOP rather than a time-based model like RET. However, this evidence was confounded by differences in the way responding to the CSs was distributed over time within the trial. The reinforcement schedule that included many short non-reinforced trials established a pattern of responding that steadily increased across the length of a trial, whereas the schedule that

included a few long non-reinforced trials established the opposite pattern of responding—an early peak of responding that steadily declined across the length of the trial. While this difference in the pattern of responding could be explained by changes in likelihood that a trial would be reinforced as time in the trial increased, it complicated interpretation of the overall difference in responding created by the split trials design.

Experiment 2

Experiment 1 found evidence that many short non-reinforced trials led to weaker conditioning than a few long non-reinforced trials, but this evidence was confounded with differences in response rate across the length of the trial. Experiment 2 attempted to correct for this confound. Rats were trained with two CSs. One CS, A, had 16 presentations per session, half of which were reinforced, and all trials (reinforced and non-reinforced) had a mean duration of 30 s. The other CS, B, had 48 presentations per session, each with a mean duration of 10 s, and one in six trials (17%) were reinforced. This means that the two CSs were matched for total reinforcement rate per second, as in Experiment 1, but the likelihood that a trial would be reinforced or non-reinforced was uniform across the length of each trial.

The comparison between CS A and CS B is similar to the comparison made by Harris et al. (2015) between a 30-s CS reinforced on 100% of trials and a 10-s CS reinforced on 33% of trials, except that the reinforcement probabilities of the two CSs has been halved (50% versus 17%, respectively). This means that the present experiment has added eight long (mean 30 s) non-reinforced presentations of A, and 24 short (mean 10 s) non-reinforced presentations of B to each session. We know from the various experiments reported by Harris et al. that, without these added non-reinforced trials, responding to A and B is acquired in more or less identical fashion.

Thus, the question posed by this experiment is whether adding 24 short non-reinforced trials will have the same effect on acquisition as adding eight long non-reinforced trials. For a time-based model like RET, both schedules of non-reinforced trials will have the same effect on the overall reinforcement rate and, therefore, should have the same effect on conditioning. In contrast, SOP makes an unambiguous prediction that eight long non-reinforced trials of A should produce less loss of associative strength than 24 short trials of B (despite having the same total duration) because the average associability of B would be higher than that of A. The design of Experiment 2, unlike that in Experiment 1, matches the two CSs for the likelihood of reinforcement per second as time elapses during the CS, which should prevent any changes in responding as a function of time in the trial.

Methods

Subjects and Apparatus

Sixteen experimentally-naïve female albino Sprague Dawley rats (8 to 10 weeks of age) were obtained from the same source and were housed in the same manner as for Experiment 1. They were trained and tested in the same conditioning chambers as in Experiment 1. The steady light and white noise described in Experiment 1 were used as CSs A and B in counterbalanced fashion.

Procedure

Rats received magazine training before commencing daily conditioning sessions 5 days per week for a total of 32 sessions. Each session contained 16 presentations of CS A, with a mean duration of 30 s (range 2-58 s), and 48 presentations of CS B, with a mean duration of 10 s (range 2-18 s). Eight of the presentations of each CS were reinforced with a single food pellet delivered at offset of the stimulus. The remaining trials were all non-reinforced. The inter-trial interval varied randomly with a mean of 100 s, and the session lasted approximately 120 min. Photo-

beam interruptions by head entry into the magazine were recorded during each CS and each 20-s pre-CS period.

Results and Discussion

Mean response rates to CSs A and B across the 32 acquisition sessions are shown in Figure 2A. The unbroken black line in the plot shows the response rate averaged over the full length of each trial of CS A. For the purpose of comparing response rates measured across equivalent trial durations for the two CSs, we also calculated the mean response rate to CS A measured from the first third of each trial (regardless of its length). This is shown as the dashed black line in Figure 2A. It is clear from this plot that responding to the two CSs was acquired at very similar rates.

----- Insert Figure 2 about here -----

A 2x32 repeated measures ANOVA, with Greenhouse-Geisser correction, compared responding to CS B with responding during the first third of each trial of CS A across the 32 conditioning sessions. This showed there was no overall difference between the two CSs, $F(1,15) = 0.71$, $p = .794$, no significant main effect of Session, $F(1.6,24) = 2.69$, $p = .098$, and no interaction between CS and Session, $F(3.9,59) = 1.55$, $p = .202$. A Bayesian paired sample t-test (JASP_Team, 2018) compared the strength of evidence that rats responded more to A than to B, averaged over the 28 sessions, versus the strength of evidence that there was no difference in responding to the two CSs. The Bayes Factor, with Cauchy prior set to 0.707, provided moderate support for the null hypothesis, $BF_{01} = 4.7$.

Figure 2B shows the response rates to CS A and CS B per second averaged over the last 16 sessions. It is clear that the pattern of responding to the two CSs was very similar across the first 18 s of the trial (the trial length that was common to both CSs). It is also clear that responding

only reached the peak level after approximately 10 s into the trial for both CSs. The fact that responding was lower for the first 10 s than for time points longer than 10 s would explain why response rates averaged over the full length of trials with CS A (unbroken line in Figure 2A) appeared to be slightly higher than response rates to CS B. This is because the first 10 s, when responding is still rising, represents a much larger proportion of B's total length than of A's total length. Therefore, when we just look at responding during the first third of A, the two CSs come out very similar.

In sum, the impact on conditioning of many short non-reinforced trials appears to be indistinguishable from the impact of a few long non-reinforced trials if the total duration of those trials is the same. This result is consistent with time-based models like RET but is inconsistent with trial-based models, even a real-time model like SOP which predicts that many short non-reinforced trials should produce greater associative loss than a few long trials.

Experiment 3

Experiment 3 was designed to replicate Bouton and Sunsay's (2003) demonstration that conditioning was more negatively affected by non-reinforced presentations if those presentations occurred 60 s before the reinforced presentation (i.e., late in the interval separating reinforced trials) than if the non-reinforced presentations occurred 240 s before the reinforced presentation (early in the interval between reinforcements) (see also Sunsay et al., 2004). This result can be explained by SOP by assuming that A1 processing of the CS on reinforced trials would be reduced by the preceding non-reinforced presentation, and that a more recent presentation will reduce A1 processing more than a less recent presentation. However, there is an alternative explanation for Bouton and Sunsay's finding that stems from the fact that their experiment used fixed inter-reinforcement intervals (350 s) which would have allowed the rats

to learn to expect delivery of the reinforcer towards the end of that interval and not early in the interval. There is good evidence that rats learn to anticipate the time of a US across fixed US-US intervals, with their response rate gradually increasing across the length of that interval (Kirkpatrick & Church, 2000, 2003). Given this, when Bouton and Sunsay presented the CS without reinforcement early in the US-US interval (60 s into a 350 s interval), the rats' expectation of the US would have been low compared to their expectation when the CS was presented late in that interval (240 s into the 350 s interval). The fact that late presentations of the CS led to weaker conditioning than early presentations could be attributed to the fact that the rats' expectation of the US was higher during the late presentations, with the result that late presentations had a more negative impact on conditioning. This timing-based mechanism does not apply to the result reported by Sunsay et al. (2004), described earlier, which showed that conditioning was affected by the identity of the stimulus that was presented before each reinforced trial rather than the timing of that presentation. Moreover, the explanation depends on the rats learning about the fixed US-US interval, and therefore it would not predict any difference in the effect of a non-reinforced CS presentation early versus late in the inter-US interval if those intervals were variable such that the rats' expectation of the US was matched at the early and late presentation times. This possibility is addressed in Experiment 3.

The design of Experiment 3 is shown in Figure 3. Rats were conditioned with three CSs, A, B, and X. Each 10-s reinforced (+) trial of A or B was preceded by a 40-s non-reinforced (–) trial of the same CS that occurred either 60 s (for A) or 240 s (for B) earlier. The A– and B– trials occurred after a variable interval following the preceding + trial. These intervals were 180-s longer for A than for B, such that the timing of A+ and B+ trials since the last reinforcement was matched. Reinforced trials of CS X were intermixed among the A– then A+ and the B– then B+

sequences to ensure that the likelihood of reinforcement was matched for all intervals since the last +. Each session included an additional six X+ trials after an interval of 80 s, 170 s, and 260 s from the previous +. This ensured that reinforcement was equally likely from 90 s up to 630 s after the last reinforcer.

Time (s)									
0	90	180	270	360	450	540	630		
X+									
+	[80]	X+							
+	[80]	X+							
+	[170]		X+						
+	[170]		X+						
+	[260]			X+					
+	[260]			X+					
+	[250]			A–	[60]	A+			
+	[340]				A–	[60]	A+		
+	[430]					A–	[60]	A+	
+	[520]						A–	[60]	A+
+	[70]	B–	[240]		B+				
+	[160]		B–	[240]		B+			
+	[250]			B–	[240]		B+		
+	[340]				B–	[240]		B+	

Figure 3. Summary of the design of Experiment 3, showing the structure of the 15 trial-sequences of each session. Rats were trained with 3 CSs, A, B, and X. On all reinforced (+) trials, the CS duration was 10 s. On non-reinforced trials (–) the CS duration was 40 s. The first trial of every session was X+ (Row 1), after which the order of the remaining 14 sequences was randomly shuffled. The grey + in the left-most column marks the delivery of reinforcement at the end of the preceding sequence. Light grey cells represent inter-stimulus intervals within a trial-sequence, with length of the interval shown in square brackets. The dark grey cells are part of the following sequence.

Methods

Subjects and Apparatus

Sixteen experimentally-naïve female albino Sprague Dawley rats (8 to 10 weeks of age) were obtained from the same source and were housed in the same manner as for Experiment 1. They were trained and tested in the same conditioning chambers as in Experiment 1. The steady light and white noise described in Experiment 1 were used as CSs A and B in counterbalanced fashion. In addition, a clicker stimulus (4 clicks per second, each click 24 dB above background) was emitted from a module (Med Associates, product ENV 135M) located on the wall of the conditioning chamber above and to the right of the magazine. The clicker served as CS X for all rats.

Procedure

Rats received magazine training before commencing daily conditioning sessions 5 days per week for a total of 32 sessions. The first trial of every session was a 10-s presentation of X followed immediately by delivery of a food pellet. Each session contained six more 10-s X+ presentations, four 10-s A+ presentations, and four 10-s B+ presentations. The six X+ trials commenced either 80 s, 170 s, or 260 s after the previous reinforced trial (two trials at each interval). Each A+ presentation followed a 40-s A– presentation, separated by a 60-s inter-stimulus interval. Each B+ presentation followed a 40-s B– presentation, separated by a 240-s inter-stimulus interval. The four A– presentations occurred either 250 s, 340 s, 430 s, or 520 s after the previous reinforced trial; the four B– presentations occurred either 70 s, 160 s, 250 s, or 340 s after the previous reinforced trial. The six X+ trials, the four A– then A+ sequences, and the four B– then B+ sequences were randomly intermixed. The full list of sequences is shown in

unrandomized order in Figure 3. Photo-beam interruptions by head entry into the magazine were recorded during each CS and each 20-s pre-CS period. The session lasted approximately 90 min.

Results and Discussion

Mean response rates during the 10-s reinforced trials of both CSs are shown in Figure 4A. Responding to the two CSs was very similar across the 32 sessions of conditioning. A 2 x 32 repeated measures ANOVA compared response rates between the two CSs across the 32 sessions. It showed that there was a significant main effect of Session, $F(5.9,89) = 9.23, p < .001$, $\eta^2_p = 0.09$, but no main effect of CS, $F(1,15) = 1.50, p = .239$, nor was there a significant interaction between CS and Session, $F(7.7,116) = 1.38, p = .084$. To test the strength of evidence against our hypothesis that responding to CS A would be lower than responding to CS B, we conducted a one-tailed Bayesian t-test (JASP) comparing response rates to A and B averaged over the 32 sessions. The Bayes Factor indicated that there was moderately strong evidence against our hypothesis, $BF_{01} = 7.7$.

----- Insert Figure 4 about here -----

The experiment was designed so that the likelihood of food was matched for each time point at which CS A or CS B was presented without reinforcement. To check that this manipulation was effective, we have calculated the mean response rate over the first 10 s of each 40-s non-reinforced trial of A and B as a function of the time when that trial occurred after the previous reinforced trial (see Figure 4B). As predicted, there was little difference in responding as a function of time since last US. A 2 x 4 repeated measure ANOVA showed that there was no significant main effect of CS, $F(1,16) = 0.10, p = .755$, or Time, $F(2.1,32) = 2.63, p = .084$, nor was there a significant interaction between CS and Time, $F(2.8,42) = 2.13, p = .115$.

In sum, it appears that the impact of non-reinforced presentations of a CS on the acquisition of responding to that CS did not vary as a function of the interval between the non-reinforced and subsequent reinforced presentations. That is, in contrast to the prediction of SOP, recent non-reinforced presentations of a CS do not reduce conditioning more than do remote presentations.

General Discussion

We have presented three experiments designed to test whether conditioning is affected by characteristics of trial structure in ways that can be explained by the trial-based account of conditioning offered by Wagner (1981) but not by a time-accumulation account like RET (Balsam et al., 2010; Gallistel & Gibbon, 2000). The experiments followed two designs used by Bouton and Sunsay (2003). In the first design, those authors had shown that non-reinforced exposure to a CS during conditioning reduced the acquisition of responding more if that exposure was split across many short non-reinforced trials than if it occurred as a few long trials. In the second design, they had shown that non-reinforced trials during conditioning reduced acquisition more if each non-reinforced trial occurred 60 s before each reinforced trial than if it occurred 240 s before each reinforced trial (see also Sunsay et al., 2004). The experiments we present here sought the same evidence but using within-subjects rather than between-subjects designs and, perhaps more importantly, using variable CS-US intervals when replicating Bouton and Sunsay's first design and using variable US-US intervals when replicating their second design. We introduced variable intervals into the experiments so that the rats could not learn when to expect delivery of the food reinforcer which could otherwise affect what rats learned from short versus long non-reinforced trials or from non-reinforced trials that occurred close to or far from the expected time of reinforcement. Considering the results of all three experiments presented here,

our findings differ from those reported by Bouton and Sunsay, and thus have not supported the predictions made by SOP (Wagner, 1981) about the way that non-reinforced trials can affect conditioning.

Experiment 1 trained rats with four CSs on a partial reinforcement schedule using variable CS-US intervals (between 2 and 46 s) on reinforced trials. The key comparison was between responding to a CS, B, that had a few long non-reinforced trials (5 trials of 48 s) in each session and a CS, C, that had many short non-reinforced trials (15 trials of 16 s) in each session. When we compared responding to these CSs on reinforced trials (these trials were chosen because they were identical in length and number for both CSs) we observed that the rats responded more on average to B than to C, consistent with SOP's prediction that many short non-reinforced trials would produce greater net loss of associative strength. However, even though we had used variable CS-US intervals on reinforced trials to prevent rats from learning when to expect the US, we still observed a difference between the two CSs in the time course of the rats' responding within the trial. The rats showed lower levels of responding early in each presentation of CS C but their response rates steadily increased across the length of the trial. The rats showed high levels of responding early in each presentation of CS B and their response rate steadily declined as the length of the trial increased. This difference in pattern can be explained by changes in the likelihood that a trial would end without reinforcement. This likelihood decreased across trials for CS C because non-reinforced trials were all short, so the longer a trial lasted the more likely it would be reinforced, consistent with the steady increase in the rats' response rates. Conversely, the likelihood of non-reinforcement increased for CS B because all non-reinforced trials were long, consistent with the steady decline in the rats' response rates across the length of the trial. This difference in pattern of responding across the length of the CS complicated the

evidence that conditioning of CS C was weaker than conditioning of CS B. That evidence depends on comparing response rates averaged over all trials, both long and short. However, in this average there will be greater representation of responding that occurs close to CS onset than responding that occurs at longer intervals after CS onset (because responding soon after CS onset will be recorded in all trials, but responding later in the CS can only be recorded on long trials).

Experiment 2 was designed to eliminate the response-timing confound observed in Experiment 1 by making sure that the likelihood of a trial ending in reinforcement or non-reinforcement was uniform across the full length of the trial. It succeeded in this objective, in that the pattern of responding across the length of the trial was well matched between the two target CSs. However, there was also no evidence of any difference in the acquisition of responding to the two CSs. In other words, once we had removed the confounding effect of differences in timing of responding across the trial, there was no evidence that splitting long non-reinforced trials into many short non-reinforced trials had any differential effect on conditioning. This result is consistent with the predictions of RET which assumes that conditioned responding is affected by the total time of non-reinforced exposures to a CS regardless of how that time is accumulated across trials. It is inconsistent with simple trial-based models (e.g., Rescorla & Wagner, 1972) that ignore the effect of time within each trial and so assume that each non-reinforced trial, regardless of length, will have the same impact on conditioning. The result is even inconsistent with a “real-time” trial-based model like SOP that describes how time within the trial affects conditioning, and thus predicts that, when the number of extinction trials is held constant, long non-reinforced trials will reduce conditioning more than short non-reinforced trials. SOP nonetheless assumes that individual trials matter, rather than just cumulative time, such that it predicts a long non-reinforced trial will have a weaker impact on conditioning than

multiple short non-reinforced trials of equal cumulative duration. We have not seen evidence confirming that prediction here.

Experiment 3 followed the second design used by Bouton and Sunsay (2003) to test whether non-reinforced trials reduce conditioning more when they occur shortly before each reinforced trial than when they occur long before each reinforced trial (see also Sunsay et al., 2004). Whereas Bouton and Sunsay (and Sunsay et al.) observed such an effect—rats showed weaker conditioning when the interval between non-reinforced and subsequent reinforced trials was 60 s than when the interval was 240 s—we did not replicate that result in Experiment 3. One difference between those previous experiments and the present one is that the latter used a within-subjects design whereas the earlier experiments used between-subjects designs. It is thus possible, for example, that generalization between the two stimuli used here obscured any difference in responding, a possibility that is excluded in the between-subjects design used by Bouton and Sunsay and Sunsay et al.. However, the two stimuli used in the present experiment were from different modalities (one visual, one auditory) which should limit generalization, and, more tellingly, we did see evidence that the rats in Experiment 1 discriminated between these same stimuli. It is also possible that within-subject designs create more interaction (e.g., interference) among recently-activated memories that might affect their influence on associative learning or performance in some way.

The other key difference between Experiment 3 of the present series and the experiments reported by Bouton and Sunsay (2003) is that the interval between reinforced trials varied unpredictably in Experiment 3 whereas that interval was fixed in the experiments by Bouton and Sunsay. One way to explain the result reported by Bouton and Sunsay is that the rats learned to expect a reinforced trial only after a long interval since the last reinforcer. Therefore,

presentation of the CS early in the interval between reinforcers (i.e., 240 s before the next reinforced trial) would be accompanied by a relatively weak expectation of reinforcement compared to a presentation that occurred late in the interval between reinforcers (i.e., 60 s before the next reinforced trial). As a result, the rats would learn less about the absence of reinforcement on the early non-reinforced trials than on the late non-reinforced trials, and therefore show stronger conditioning in the former case, just as Bouton and Sunsay observed. This time-based expectancy mechanism would be eliminated in Experiment 3 because the likelihood of a trial being reinforced was equated for CS presentations that occurred 60 s or 240 s before the next reinforced trial. The fact that there was also no evidence for the effect of inter-trial interval on conditioning suggests that the result reported by Bouton and Sunsay could have in fact been due to time-based differences in expectation of reinforcement, rather than the mechanism used by SOP to account for the effects of inter-trial interval on conditioning (but see Sunsay et al., 2004).

What are the implications of these results for trial-based associative models of conditioning like SOP versus time-accumulation models like RET? The absence of an effect of splitting non-reinforced trials in Experiment 2, or of varying the interval between non-reinforced and reinforced trials in Experiment 3, is consistent with RET. According to RET, the emergence of conditioned responding depends on the overall rate at which the US occurs during the CS, and therefore it is affected by cumulative duration of non-reinforced exposures rather than the number of such exposures or the interval between exposures (as long as the overall rate of reinforcement is held constant). On the other hand, the results of Experiments 2 and 3 are problematic for SOP and weigh against the evidence presented by Bouton and Sunsay (2003) that had supported SOP. To provide a more thorough analysis of the implications of these

findings for SOP, we have run simulations of the model to demonstrate its predictions for these experiments.

SOP's predictions are sensitive to the dynamics of the element activation profiles and thus depend on the activation and decay parameters. Therefore, we have simulated the experimental data using different sets of assumptions about the rate at which CS elements are activated into A1 (p_1) as well as the rate at which they decay from A1 to A2 (p_{d1}) and from A2 to the inactive state (p_{d2}). When varying decay rates between simulations, we have followed Wagner's (1981) suggestion that the value for p_{d1} be 5 times larger than for p_{d2} . The plots in the top row of Figure 5 show the proportion of CS elements active in A1 and A2 states ($pA1$ and $pA2$) across time during a 58-s CS and during the subsequent 140-s inter-trial interval. The first plot, in the top left of the figure, shows CS activation for the case when elements are rapidly recruited into A1 ($p_1 = 0.8$) on presentation of the CS and then rapidly decay from A1 to A2 ($p_{d1} = 0.5$) and from A2 to become inactive ($p_{d2} = 0.1$). The second plot shows A1 and A2 activity with intermediate recruitment and decay rates, and the third plot (top right of the figure) shows slow recruitment and decay. To model learning and the strength of the conditioned response (CR) in the experiments described here, the simulations were run in 1-s time steps, and the CS duration and inter-trial interval varied from trial to trial as in the experiments themselves. The CR in each second was equal to the strength of US activation into A2, calculated by multiplying the net associative strength (V_{net} = excitatory strength, V_x , minus inhibitory strength, V_i) by the weighted sum of CS elements in A1 and A2 ($pA1 + 0.01 * pA2$; as suggested by Wagner), and this was averaged across the length of the trial. For each second that the CS was present and the US absent, V_i increased according to the equation: $\Delta V_i = \beta^0 * pA1 * (V_{net} * [pA1 + 0.01 * pA2])$, where β^0 is a learning rate parameter for each time step when the US is absent. In any second

that the CS and US are both present (at the end of a reinforced trial), V_x increased according to the equation: $\Delta V_x = \beta^+ * pA1 * (1 - V_{net} * [pA1 + 0.01 * pA2])$, where β^+ is a learning rate parameter when the US is present. We varied the values of β^+ and β^0 between simulations to compensate for the impact of differences in CS decay rate on overall learning rate, but they always maintained a consistent ratio as $\beta^+:\beta^0 = 5:1$. Our modelling updated V in the inter-trial interval following non-reinforced trials, but did not update V across the inter-trial interval after reinforced trials because it was assumed that the delivery of the US would quickly displace residual CS elements from A1.

The second row of plots in Figure 5 shows simulated CR strength during reinforced presentations of CS B and CS C across blocks of trials in Experiment 1. In all three simulations, CRs are higher for CS B (given a few long non-reinforced trials per session) than CS C (given many short non-reinforced trials). This difference arises for two reasons. First, mean associability (linked to $pA1$) is higher across the many short presentations of C than the few long presentations of B, resulting in greater inhibitory strength for C than B. Second, except when CS decay rates are very fast, net V will continue to decrease during the inter-trial interval after non-reinforced trials as a result of residual CS activation in A1. This post-trial loss of associative strength will be greater for CS C than CS B because there were many more non-reinforced trials of CS C. The results obtained in Experiment 1 (Figure 1) showed there was more responding to CS B than CS C, consistent with the SOP simulations. However, in the simulations the difference between B and C increases as responding approaches asymptote whereas the reverse is true of the experimental data (there, the difference was present during early sessions but disappeared as responding approached asymptote). Moreover, as already noted, any differences observed between B and C appeared to be linked to differences in the way that responding changed across the length of

the trial, rather than differences in overall level of responding. Therefore, it is unclear whether the experimental data show the difference in conditioning strength that is predicted by SOP.

The bottom row of plots in Figure 5 show simulated CRs for Experiment 2. In each case, simulated responding to the longer CS (A) is only taken from the first third of each trial in order to match it to the time window over which responding was measured for CS B. The empirical data from Experiment 2 showed no difference between responding to A and B over this matched time window (see Figure 2). In contrast, SOP predicts a clear difference, regardless of recruitment and decay rates. The reasons for this difference are the same as for the simulations of Experiment 1. It should be noted that the simulations presented in Figure 5 do not model the impact of carry-over activity from one trial to the next (this process is modelled in Figure 6, as discussed below). We have not attempted to model that effect for Experiments 1 and 2 because on the majority of trials the CS presented on a given trial was different from that presented in the previous trial, in which case the cross-trial effects examined in Experiment 3 and modelled in Figure 6 would not apply. That said, one would expect there to be occasions when the same CS was presented on consecutive trials, and this would happen more often for the more frequent CS (i.e., CS C in Experiment 1 and CS B in Experiment 2). In these instances, residual CS activity in A2 would reduce A1 activation on the subsequent trial, which would dampen responding and reduce learning. This would further increase the predicted difference between the CSs, making SOP's predictions even less consistent with what was obtained.

Figure 6 shows simulations of SOP for Experiment 3 that directly examined the impact of carry-over CS activation between trials. The top row shows A1 and A2 activation for two CSs, A and B, during the 40-s presentation of the CS, across the following inter-trial interval, and during the subsequent 10-s CS presentation either 60-s later, for CS A, or 240-s later for CS B. When CS

activation decays rapidly, there is no residual activity between the 40-s and 10-s presentations, and therefore there is no difference between CS A and CS B in their A1 activation on the 10-s reinforced presentation. As a result, the simulated CR strength did not differ between the two CSs, consistent with what was observed in Experiment 3. With an intermediate rate of recruitment and decay (middle plot), both CSs have decayed from A1 over their inter-trial intervals but there is much more residual A2 activation for CS A at the end of its short inter-trial interval than for CS B at the end of its long interval. As a result, during the 10-s reinforced presentation, there is less A1 activation of CS A than of CS B. This translates into dampened responding and weaker conditioning to A than B in SOP's simulation (lower middle plot in Figure 6), which is consistent with the result reported by Bouton and Sunsay (2003) but inconsistent with what was found here in Experiment 3. When the recruitment and decay rates are slow (top right plot in Figure 6), there are marked differences in both A1 and A2 activation over the inter-trial interval between CS A and CS B. At the end of the interval, there is greater residual A1 activation, and lower A2 activation, of A than B, and both of these differences mean there is stronger A1 activation of A than B during the 10-s reinforced trial of each CS. As a result, responding to A is predicted to be greater than responding to B (see bottom right plot in Figure 6). Once again this difference is inconsistent with the results obtained. These comparisons raise another possible explanation for the difference between the results we report here for Experiment 3 and those reported by Bouton and Sunsay—the different strains of rats used in each experiment may exhibit different stimulus decay rates, with the current rats showing much faster decay than Bouton and Sunsay's rats.

All of the simulations described above simultaneously varied recruitment (p_1) and both A1 and A2 decay, thus maintaining the relative rates of these processes. One could, of course,

allow these rates to vary independently. For example, in simulating Experiments 1 and 2, one could assume a rapid recruitment (e.g., $p_1 = 0.8$) with slow decay rates ($p_{d1} = 0.05$ and $p_{d2} = 0.01$, or $p_{d1} = 0.005$ and $p_{d2} = 0.001$), or slow recruitment ($p_1 = 0.2$) with relatively rapid decay ($p_{d1} = 0.05$ and $p_{d2} = 0.01$). However, our modelling (not shown) reveals these particular changes to the recruitment and decay rates do not materially affect SOP's predictions for the difference between CSs B and C in Experiment 1 or between CSs A and B in Experiment 2. Indeed, SOP's predictions only change, such that it predicts no difference between B and C in Experiment 1 and between A and B in Experiment 2, if the recruitment rate into A1 is reduced to the point where it is lower than the decay rate from A1 (e.g., if $p_1 = 0.02$, $p_{d1} = 0.05$, and $p_{d2} = 0.01$). Alternatively, simulations of Experiment 3 might be expected to change if one were to relax Wagner's (1981) proposal about relative decay rates from A1 and A2 (that p_{d1} be five times greater than p_{d2}). One might expect that a fast or intermediate decay rate from A1, accompanied by a very slow decay rate from A2, would produce a negligible difference in conditioning between CS A and CS B in Experiment 3 because the two CSs would both have similarly high levels of A2 activation at the start of the 10-s reinforced trial. Our own modelling of this possibility (not plotted here) confirms that the difference in CR strength between A and B could indeed be very small under these circumstances. However, these circumstances also result in very low levels of predicted responding (very weak activation of the US by the CS), such that the relative difference between A and B would still be substantial (on the order of a two-fold difference). In other words, as long as there is a measurable CR, SOP predicts there should be more responding to B than to A in Experiment 3.

While the present findings are more consistent with RET than SOP, there are other relevant findings to be considered when evaluating the two models. One finding concerns the

original results reported by Bouton and Sunsay (2003). The present experiments suggest that those results may have been due to time-based changes in expectation of reinforcement: rats expected reinforcement 10-s after onset of the CS and therefore learned relatively little about the absence of reinforcement across the last 30 s of a 40-s non-reinforced trial; the rats learned less about non-reinforced trials that occurred early in a fixed inter-reinforcement interval than late in the interval because their expectation of reinforcement was greater in the second case. However, this type of account is not built into RET. RET describes how knowledge of temporal intervals affects performance (the decision to respond) but not how it affects learning. That is, it does not specify any mechanism by which time in one portion of a trial would count more than time in another portion of a trial, or by which non-reinforced exposure at one point in an inter-reinforcement interval would count more than non-reinforced exposure at another point in that interval. Thus RET's description of the way animals learn about reinforcement rates cannot account for the results reported by Bouton and Sunsay.

Another relevant finding that is not accounted for by RET, but can be explained by SOP, was Sunsay et al.'s (2004) demonstration that conditioning of two CSs was impaired in a group for which reinforced trials of each CS were preceded by non-reinforced trials of the same CS relative to a group for which reinforced trials of each CS were preceded by non-reinforced trials of the other CS. The time-based US expectancy account offered here for Bouton and Sunsay's (2003) results do not apply to Sunsay et al.'s result because the timing of the non-reinforced trials, relative to US delivery, was matched across the two conditions being compared. According to RET, there should be no difference in conditioning between the two groups because the rate at which the CSs were reinforced, and the overall rate at which reinforcement occurred, was the same for each group. SOP, on the other hand, predicts that non-reinforced presentations of a CS

shortly before reinforced trials of the same CS will interfere with conditioning by reducing A1 processing of the CS on the reinforced trials, and that this interference will not occur if the reinforced trial of one CS follows a non-reinforced trial of a different stimulus. Thus the finding reported by Sunsay et al. does stand as evidence in favour of SOP rather than RET.

The third relevant finding in evaluating trial-based and time-based models concerns the manner in which responding extinguishes. Even though the acquisition of conditioned responding appears to be affected only by the cumulative duration of non-reinforced exposure to the CS and not the number of non-reinforced exposures (at least for the conditioning procedures used here), there is ample evidence using the same procedures that the number of non-reinforced trials during conditioning affects the subsequent extinction of that responding (Bouton, Woods, & Todd, 2014; Haselgrove, Aydin, & Pearce, 2004). Even when the addition of non-reinforced trials during conditioning is shown to have no effect on the acquisition of responding (as long as those trials do not increase the cumulative duration of CS exposure), partial reinforcement during conditioning does slow the rate of subsequent extinction (Chan & Harris, 2017, 2019; Harris & Kwok, 2018; Harris, Kwok, & Gottlieb, 2019; Harris, Seet, & Kwok, 2019). Moreover, the magnitude of this effect on extinction scales directly with the number of non-reinforced trials per conditioning session—more non-reinforced trials during conditioning leads to slower extinction without affecting the acquisition of responding (Chan & Harris, 2019). This evidence shows that rats do indeed learn about the number of non-reinforced trials they experience during conditioning, even though the acquisition of their responding is only affected by the cumulative duration of those trials. This is a challenge for time-based models that take no account of what animals may learn about individual trials. It suggests that any computational operation used by the animal to detect a change in reinforcement during extinction must depend

on an estimate of probability of reinforcement per trial rather than per unit time (Gallistel, 2012; Gallistel & Papachristos, 2020). The evidence that rats learn about non-reinforced trials during partial reinforcement is consistent with the emphasis on learning about individual trials captured in models like the Rescorla-Wagner model (Rescorla & Wagner, 1972) and SOP (Wagner, 1981). There is, nonetheless, a significant challenge for any model to explain why learning about non-reinforced trials only affects extinction of responding and not the acquisition of responding.

One way to understand this paradox is to consider the different temporal properties of reinforcement and non-reinforcement (Harris, 2019). The available evidence, at least from experiments using variable CS-US intervals, is that animals learn in a continuous (cumulative) fashion about reinforcement, and respond during conditioning based on the moment-by-moment probability of the US occurring when the CS is present, but at the same time they learn in a discrete manner about individual trials as single events when they are learning about non-reinforcement. Animals may learn about reinforcement and non-reinforcement in different ways because reinforcement is a temporally specific event—the US occurs at a specific moment within the trial, even when the timing of the US is variable. In contrast, the non-occurrence of the US is not something that happens at any specific moment (except, perhaps, when the timing of the US is fixed such that the animal has learned when the US occurs and therefore can detect its absence at a specific moment). It might be useful for animals to learn about a CS-US relation continuously throughout the trial so that they can anticipate the moment that the US arrives. However, when learning that the CS-US relationship has changed, it would be less useful to learn continuously during a trial that an expected US is absent since the omission of the US can only be established once the trial has ended.

References

- Balsam, P. D., Drew, M. R., & Gallistel, C. R. (2010). Time and associative learning. *Comparative Cognition and Behavior Reviews*, 5, 1-22.
- Balsam, P. D., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neurosciences*, 32, 73-78.
- Barela, P. B. (1999). Theoretical mechanisms underlying the trial-spacing effect in Pavlovian fear conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 25, 177-193.
- Barnet, R. C., Grahame, N. J., & Miller, R. R. (1995). Trial spacing effects in Pavlovian conditioning: A role for local context. *Animal Learning and Behavior*, 23(3), 340-348.
- 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.
- Bouton, M. E., Woods, A. M., & Todd, T. P. (2014). Separation of time-based and trial-based accounts of the partial reinforcement extinction effect. *Behavioural Processes*, 101, 23-31.
- Chan, C. K. J., & Harris, J. A. (2017). Extinction of Pavlovian conditioning: The influence of trial number and reinforcement history. *Behavioural Processes*, 141, 19-25.
- Chan, C. K. J., & Harris, J. A. (2019). The Partial Reinforcement Extinction Effect: The proportion of trials that are reinforced during conditioning predicts rate of extinction. *Journal of Experimental Psychology: Animal Learning and Cognition*, 45, 43-58.
- 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.
- Domjan, M. (1980). Effects of intertrial interval on taste-aversion learning in rats. *Physiology and Behavior*, 25, 117-125.
- 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.
- Ellison, G. D. (1964). Differential salivary conditioning to traces. *Journal of Comparative and Physiological Psychology*, 57, 373-380.
- Gallistel, C. R. (2012). Extinction from a rationalist perspective. *Behavioural Processes*, 90, 66-80.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107, 289-344.
- Gallistel, C. R., King, A., & McDonald, R. (2004). Sources of Variability and Systematic Error in Mouse Timing Behavior. *Journal of Experimental Psychology: Animal Behavior Processes*, 30, 3-16.
- Gallistel, C. R., & Papachristos, E. B. (2020). Number and time in acquisition, extinction and recovery. *Journal of the Experimental Analysis of Behavior*, in press. doi:10.1002/jeab.571
- Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, 84, 279-325.
- 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.

- 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., & Church, R. M. (1990). Representation of time. *Cognition*, 37, 23-54.
- Gibbon, J., Church, R. M., Fairhurst, S., & Kacelnik, A. (1988). Scalar expectancy theory and choice between delayed rewards. *Psychological Review*, 95, 102-114.
- 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.
- Harris, J. A. (2015). Changes in the distribution of response rates across the CS-US interval: Evidence that responding switches between two distinct states. *Journal of Experimental Psychology: Animal Behavior Processes*, 41, 217-231.
- Harris, J. A. (2019). The importance of trials. *Journal of Experimental Psychology: Animal Learning and Cognition*, 45, 390-404.
- 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., & Kwok, D. W. S. (2018). The probability of reinforcement per trial affects posttrial responding and subsequent extinction but not within-trial responding. *Journal of Experimental Psychology: Animal Learning and Cognition*, 44, 23-35.
- Harris, J. A., Kwok, D. W. S., & Gottlieb, D. A. (2019). The Partial Reinforcement Extinction Effect depends on learning about non-reinforced trials rather than reinforcement rate. *Journal of Experimental Psychology: Animal Learning and Cognition*, in press.
- Harris, J. A., & Livesey, E. J. (2010). An attention modulated associative network. *Learning & Behavior*, 38, 1-26.
- Harris, J. A., Patterson, A. E., & Gharaei, S. (2015). Pavlovian conditioning and cumulative reinforcement rate. *Journal of Experimental Psychology: Animal Learning and Cognition*, 41, 137-151.
- Harris, J. A., Seet, M. S., & Kwok, D. W. S. (2019). The Partial-Reinforcement Extinction Effect does not result from reduced sensitivity to non-reinforcement. *Journal of Experimental Psychology: Animal Learning and Cognition*, 45, 185-202.
- Haselgrove, M., Aydin, A., & Pearce, J. M. (2004). A Partial Reinforcement Extinction Effect despite equal rates of reinforcement during Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 30, 240-250.
- JASP_Team. (2018). JASP (Version 0.9, for Windows).
- Kehoe, E. J. (1988). A layered network model of associative learning: Learning-to-learn and configuration. *Psychological Review*, 95, 411-433.
- Kehoe, E. J., & Joscelyne, A. (2005). Temporally specific extinction of conditioned responses in the rabbit (*Oryctolagus cuniculus*) nictitating membrane preparation. *Behavioral Neuroscience*, 119, 1011-1022.
- Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review*, 95, 274-295.
- Kirkpatrick, K., & Church, R. M. (2000). Stimulus and temporal cues in classical conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 26, 206-219.
- Kirkpatrick, K., & Church, R. M. (2003). Tracking the expected time to reinforcement in temporal conditioning procedures. *Learning & Behavior*, 31, 3-21.

- Lattal, K. M. (1999). Trial and intertrial durations in Pavlovian conditioning: Issues of learning and performance. *Journal of Experimental Psychology: Animal Behavior Processes*, 25(433-450).
- Li, S. H., & Westbrook, R. F. (2008). Massed Extinction Trials Produce Better Short-Term but Worse Long-Term Loss of Context Conditioned Fear Responses Than Spaced Trials. *Journal of Experimental Psychology: Animal Behavior Processes*, 34, 336-351.
- Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review*, 104, 241-265.
- McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning and Behavior*, 30, 177-200.
- Morris, R. W., Furlong, T. M., & Westbrook, R. F. (2005). Recent exposure to a dangerous context impairs extinction and reinstates lost fear reactions. *Journal of Experimental Psychology: Animal Behavior Processes*, 31, 40-55.
- 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, 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.
- Schneiderman, N. (1966). Interstimulus interval function of the nictitating membrane response of the rabbit under delay versus trace conditioning. *Journal of Comparative and Physiological Psychology*, 62, 397-402.
- 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.
- Smith, M. C., Coleman, S. R., & Gormezano, I. (1969). Classical conditioning of the rabbit's nictitating membrane response at backward, simultaneous and forward CS-US intervals. *Journal of Comparative and Physiological Psychology*, 69, 226-231.
- Sunsay, C., & Bouton, M. E. (2008). Analysis of a trial-spacing effect with relatively long intertrial intervals. *Learning & Behavior*, 36, 104-115.
- Sunsay, C., Stetson, L., & Bouton, M. E. (2004). Memory priming and trial spacing effects in Pavlovian learning. *Learning & Behavior*, 32, 220-229.
- 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.
- Vogel, E. H., Brandon, S. E., & Wagner, A. R. (2003). Stimulus representation in SOP: II. An application to inhibition of delay. *Behavioural Processes*, 62, 27-48.
- Wagner, A. R. (1981). SOP: a model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: memory mechanisms* (pp. 5-47). Hillsdale, NJ: Erlbaum.
- 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.

Figure 1. Results of Experiment 1. **A:** Mean response rates per session during reinforced trials of each CS. **B:** Response rates to CSs B and C during each 2-s bin across the length of reinforced trials averaged over Sessions 4-28.

Figure 2. Results of Experiment 2. **A:** Mean response rates per session during presentations of CS A and CS B. Responding during CS A is shown both for the mean response calculated over the full length of each trial (unbroken line) and for the mean response during the first 1/3rd of each presentation (dashed line) corresponding to the mean length of each presentation of CS B. **B:** Response rates to CSs A and B during each 2-s bin across the length of the trials averaged over the final 15 sessions.

Figure 4. Mean response rates during CS presentations, and during the pre-CS period, in Experiment 3. Each reinforced presentation of CS A or CS B was preceded by a 40-s non-reinforced presentation of the same CS. This non-reinforced presentation always occurred either 60 s (CS A) or 240 s (CS B) before the reinforced trial. The data plotted in **A** show mean response rates during each 10-s reinforced trial across each of the 32 sessions of the experiment. The data in **B** show response rates during the first 10 s of each non-reinforced trial of CS A and CS B, averaged over Sessions 21 to 32, as a function of when that presentation occurred relative to the delivery of the US at the end of the preceding trial.

Figure 5. Simulations based on Wagner's (1981) SOP model. The top row shows the proportion of CS elements active in the A1 and A2 states across time both during a 58-s CS presentation (marked by an orange band in the online version, or a gray band in the print version) and across a 140-s inter-trial interval. The first plot (top left) shows fast recruitment of CS elements into A1 ($p_1=0.8$) and a rapid decay of elements from A1 to A2 ($p_{d1}=0.5$) and from A2 to the inactive state ($p_{d2}=0.1$). The second plot shows CS activation with intermediate recruitment and decay rates, and the third plot (top right) shows slow recruitment and decay. Each simulation followed Wagner's suggestion that p_{d1} be 5-times greater than p_{d2} . The middle row and bottom row show simulated conditioned response (CR) strength, calculated as the product of A1 activation and current net associative strength, for two CSs from Experiment 1 and Experiment 2. The simulations followed the experiments in making the CS duration vary randomly from trial to trial (creating variability in the simulated CR). For Experiment 1, the CR strengths are only calculated non-reinforced trials, as we did for analysis of the experimental data, because these trials were the same length for CS B and CS C. The simulations of Experiment 2 used all trials but averaged into bins, with the number of trials in a bin depending on the reinforcement schedule of the CS: the bin size was 2 trials for CS A (reinforced on 1 in 2 trials) and 6 trials for CS B (reinforced on 1 in 6 trials). For Experiment 2, the CR strength of the longer CS (A) was only taken from the first third of each simulated trial to match this duration with the duration of CS B (see Figure 2).

Figure 6. Simulations of Experiment 3 data using Wagner's (1981) SOP model. The top row shows the proportion of CS elements active in the A1 and A2 states across time during CS presentations and during the inter-trial interval (CS presentations are marked by orange bands in the online version, grey bands in the printed version). The bottom row shows simulated conditioned response (CR) strength, calculated as the product of A1 activation and current associative strength (V), for CS A and CS B (see Figure 3 for the design of the experiment). The two plots on the left show simulations that assume a fast recruitment of CS elements into A1 (p_1) and a rapid decay of elements from A1 to A2 (p_{d2}) and from A2 to the inactive state (p_{d1}). The two central plots are from simulations that assume intermediate recruitment and decay rates, and the two plots on the right assume intermediate recruitment but slow decay. Each simulation followed Wagner's suggestion that p_{d1} be 5-times greater than p_{d2} , and these simulations also assumed the same ratio for the difference in beta between reinforced trials and non-reinforced trials.