Extinction of Pavlovian Conditioning:
The Influence of Trial Number and Reinforcement History

C. K. Jonas Chan and Justin A. Harris

The University of Sydney

Address correspondence to:

Justin Harris

School of Psychology

The University of Sydney,

Sydney 2006

Australia

Justin.Harris@sydney.edu.au

Running head: Extinction of Pavlovian Conditioning

Key words: Rat; magazine approach; variable interval; partial reinforcement.

Abstract

Pavlovian conditioning is sensitive to the temporal relationship between conditioned stimulus (CS) and unconditioned stimulus (US). This has motivated models that describe learning as a process that continuously updates associative strength during the trial or specifically encodes the CS-US interval. These models predict that extinction of responding is also continuous, such that response loss is proportional to the cumulative duration of exposure to the CS without the US. We review evidence showing that this prediction is incorrect, and that extinction is trial-based rather than time-based. We also present two experiments that test the importance of trials versus time on the Partial Reinforcement Extinction Effect (PREE), in which responding extinguishes more slowly for a CS that was inconsistently reinforced with the US than for a consistently reinforced one. We show that increasing the number of extinction trials of the partially reinforced CS, relative to the consistently reinforced CS, overcomes the PREE. However, increasing the duration of extinction trials by the same amount does not overcome the PREE. We conclude that animals learn about the likelihood of the US per trial during conditioning, and learn trial-by-trial about the absence of the US during extinction. Moreover, what they learn about the likelihood of the US during conditioning affects how sensitive they are to the absence of the US during extinction.

In Pavlovian conditioning, responding is acquired to a conditioned stimulus (CS) that consistently signals the occurrence of an unconditioned stimulus (US); that responding will subsequently extinguish if the CS is repeatedly presented but no longer followed by the US. Theoretical accounts of extinction almost universally attribute it to learning based on the unfulfilled expectation of the US that had been acquired during conditioning (Gallistel & Gibbon, 2000; Pearce, 1994; Redish, Jensen, Johnson, & Kurth-Nelson, 2007; Rescorla & Wagner, 1972; Wagner, 1981). Evidence that extinction depends on the violation of an expectancy of the US, rather than the absence of a US *per se*, is provided by demonstrations that extinction does not occur when a CS is presented without its US if the CS is presented alongside a conditioned inhibitor on those non-reinforced trials (Lovibond, Davis, & O'Flaherty, 2000; Rescorla, 2003).

Traditional models of associative learning, as exemplified by the Rescorla-Wagner model (Rescorla & Wagner, 1972), assume that any discrete pairing of a CS with a US, otherwise known as a trial, serves as the basic unit of learning. Such models treat extinction learning as the output of a trial-based operation that computes the total expectancy of the US, based on the summed conditioning strengths of all stimuli present. The model is trial-based in the sense that it treats learning (including extinction) as a one-off update in conditioning strength each time the CS is presented.

The trial-based learning algorithm cannot deal with evidence that the acquisition of conditioned responding is sensitive to the time interval between CS onset and the US. For example, longer intervals produce weaker responding than shorter intervals (e.g., Harris & Carpenter, 2011). Moreover, when the US occurs at a consistent interval after CS onset, responding is temporally graded, with the response peaking at the time the US should occur (Davis, Schlesinger, & Sorenson, 1989; Drew, Zupan, Cooke, Couvillon, & Balsam, 2005; Harris, 2015; Kehoe & Joscelyne, 2005; Pavlov, 1927; Roberts, 1981; Smith, 1968; Williams, Lawson, Cook, Mather, & Johns, 2008). To accommodate these features of conditioning, "real-time" models have been developed that describe learning as a more-or-less continuous process of updating associative strength during the trial

(Desmond & Moore, 1988; Harris & Livesey, 2010; Machado, 1997; Sutton & Barto, 1981, 1990). This means that extinction should also occur continuously when the CS is present. For example, in Wagner's SOP model (Wagner, 1981), extinction depends on the strengthening of an inhibitory CS-US association each moment that elements of the CS are directly activated by the presence of the CS while elements of the US are indirectly activated by the CS (and not directly activated by the US). A very different time-based approach is proposed by Rate Estimation Theory (RET) which describes extinction as the outcome of a decision to cease responding to a CS once sufficient evidence has accumulated that there has been a change to the rate at which the CS is reinforced (Gallistel, 2012; Gallistel & Gibbon, 2000). Evidence for the change in reinforcement is continuously accumulated when the CS is present and US absent.

The real-time associative models of conditioning, and time-based models like RET, make a clear prediction about the course of extinction. Since extinction develops continuously whenever the CS is present without the US, long trials (in which the CS is presented for an extended duration) should produce more extinction than short trials. Indeed, according to RET, extinction is directly related to the total (cumulative) amount of time that the CS has been exposed since it was last reinforced with the US. Therefore, the extent that responding to a CS will extinguish should be exactly the same after either many short non-reinforced exposures or few long exposures, as along as the total time of exposure is matched in each case. This prediction is not specific to models mentioned above, but can also be derived from Rescorla's (1997) suggestion that response loss resulting from an extinction treatment is directly proportional to the amount of responding that had been elicited by the CS during that treatment. Rescorla proposed this rule based on evidence that manipulations which suppress responding reduce evidence for extinction, and manipulations that inflate responding facilitate extinction.

Most empirical evidence has been counter to the prediction that extinction will be proportional to cumulative CS exposure. Shipley (1974) reported that loss of conditioned fear was a direct function

of the total duration of CS exposure during extinction. But more recent studies (Drew, Yang, Ohyama, & Balsam, 2004; Golkar, Bellander, & Öhman, 2013; Harris & Andrew, 2017; Haselgrove & Pearce, 2003) found that response loss is not determined by total CS exposure. Rather, in each of those studies, there was a clear effect of number of trials—responding decreased as the number of extinction trials increased. Among these studies, there was also some evidence for an effect of trial length (Drew et al., 2004; Haselgrove & Pearce, 2003), although interpretation of that evidence was complicated by the fact that differences in trial length introduced discrepancies in how well the CS duration during extinction was matched to the timing of the US during training. For example, differences in responding during long rather than short CS extinction trials could reflect differences in the timing of peak responding during each trial based on when the US was expected.

Recent experiments from our lab demonstrate clearly that response loss is related to how often the CS is presented during extinction (i.e., number of trials) and not the length of time the CS is presented for (Harris & Andrew, 2017). These experiments avoided the confounding effect of response timing (when the response rate peaks at the expected time of the US) by conditioning with variable CS-US intervals. In keeping with the earlier results (Drew et al., 2004; Golkar et al., 2013; Haselgrove & Pearce, 2003), when total duration of exposure during extinction was matched between two CSs, a CS given many short presentations extinguished faster than a CS given fewer but longer presentations. Indeed, if two CSs received the same number of exposures during extinction, responding extinguished at the same rate even if those extinction exposures were much longer for one of the CSs than the other. Therefore, extinction appears to be event-based, in that responding extinguished with each non-reinforced trial, and not time-based, in that extinction does not develop with each non-reinforced moment of the CS (Gallistel & Gibbon, 2000; Wagner, 1981) or each response to the CS (Rescorla, 1997).

Any discussion of what determines the rate of extinction must take account of how extinction is affected by the reinforcement schedule used during conditioning. One of the best-known

phenomena in associative learning is the Partial Reinforcement Extinction Effect (PREE). This refers to the observation that extinction is slower following partial reinforcement, when the US has been delivered on only some trials, than after continuous reinforcement, when the US has occurred on all trials. The effect has been widely reported both for instrumental and Pavlovian conditioning (see Mackintosh, 1974). It presents a paradox for most accounts of conditioning because, in theory, conditioning should be weaker after partial reinforcement than continuous reinforcement, and therefore responding should extinguish more quickly following partial reinforcement. The fact that partial reinforcement produces conditioning that is more resistant to extinction has inspired a number of theoretical accounts of the PREE specifically. Most emphasise that partial reinforcement increases the opportunity for generalisation of conditioning from training to extinction (e.g., Amsel, 1962; Capaldi, 1966). That is, the presence of non-reinforced trials during the partial schedule increases the similarity between that training and subsequent extinction (when all trials are non-reinforced), and so supports responding that persists for longer under extinction.

One obstacle for experiments investigating the PREE is that response rates themselves are usually very different during training with partial versus continuous reinforcement (Nevin, 1988). Differences in the level of responding at the start of extinction will confound interpretation of differences in the rate of extinction. Because of this, strong claims for a PREE rely on demonstrating that, across the course of extinction, responding to the continuously reinforced CS falls significantly below the level of responding to the partially reinforced CS (e.g., Bouton, Woods, & Todd, 2014). This means that researchers cannot determine from their data how much slower extinction is following partial reinforcement, because calculating the true rate of response loss must take into account the starting level of responding. Some researchers have attempted to bypass this problem by expressing response rates during extinction as a proportion of response rate at the end of conditioning (Nevin, 1988). But this operation relies on a strong assumption that the response scale itself is linear, in order to treat the data as on a ratio scale. This assumption cannot usually be tested and is in some circumstances unlikely to be correct.

One solution to the above problems is to compare continuous and partial reinforcement schedules that produce matched levels of responding. We recently adopted this strategy in a magazine approach paradigm with rats (Chan & Harris, submitted). Using both within- and between-subject designs, we trained rats with a long continuously reinforced CS (reinforced on every trial after a variable interval with an average duration of 30 s) and/or a short partially reinforced CS (reinforced on 33% of trials after a variable interval averaging 10 s). The choice of these reinforcement schedules was based on extensive evidence that conditioned responding is acquired in identical fashion when the total rates of reinforcement are matched in this way (Harris, Patterson, & Gharaei, 2015). As expected, both the acquisition of responding and asymptotic levels of responding were equal for the two CSs. These conditioning schedules were followed by 15 sessions of extinction, in which both CSs were repeatedly presented without reinforcement for an average duration of 17 s. (This duration was selected as the geometric mean of 10 and 30 s, making the CS duration in extinction equally similar to the two CS durations during conditioning). We found that the short partially reinforced CS extinguished at a slower rate than did the long continuously reinforced CS. The two CSs extinguished at different rates even though they had been reinforced at the same overall (cumulative) rate which appeared to produce equally strong conditioning, insofar as the rats acquired equal levels of responding to the two CSs.

The findings just described have specific implications for an explanation of the PREE offered by RET (Gallistel, 2012; Gallistel & Gibbon, 2000). As mentioned earlier, according to that theory, extinction occurs when an animal identifies that the rate of reinforcement has changed from what it had been during training. When the reinforcement rate was high during continuous reinforcement, the animal can quickly recognise that the rate has changed during extinction because it only takes a small amount of cumulative exposure to the CS before the animal notices that many expected USs have been missed. By contrast, when the reinforcement rate is low during partial reinforcement, it takes longer for the animal to identify that the reinforcement rate has changed because more time must pass during the CS in order for the animal to have missed many expected reinforcements.

RET's account of the PREE makes two strong predictions. One prediction is that there will be no PREE if CSs are matched for their cumulative reinforcement rate during conditioning. That is, the rate of extinction should not be affected by partial reinforcement if the length of the CS is reduced by the same proportion (e.g., if both the length and proportion of reinforced trials are reduced to 50%), since this keeps constant the overall rate of reinforcement per unit time. Our evidence for a PREE (Chan & Harris, submitted), despite matching CSs for total reinforcement rate during conditioning, contradicts that prediction. Our results replicate findings reported by Haselgrove, Aydin, and Pearce (2004) and Bouton, Woods and Todd (2014). Both of these studies compared extinction between groups that had been trained with a CS that was either reinforced on every trial during conditioning, or reinforced on 50% of trials. The total rate of reinforcement was matched between the two groups by either presenting two USs on every trial for the second group, or by doubling the length of the CS for the first group. Even though reinforcement rates had been matched between the groups, there was consistent evidence for a PREE—slower extinction in the group for which the CS had been reinforced on 50% of trials.

The second prediction for the PREE made by RET is that the difference in the rate of extinction between a continuously reinforced and a partially reinforced CS will correspond exactly to the difference in their rates of reinforcement during training. For example, a CS that was reinforced on 1 in 3 trials will take 3 times longer to extinguish than a CS reinforced on every trial. Therefore, the PREE should disappear if the partially reinforced CS receives more extinction exposure than the continuously reinforced CS, either as a consequence of more trials or longer trials. Bouton et al. (2014) recently tested this prediction. They trained rats with a 10-s CS that was followed with food on either 25% of trials or 100% of trials. Responding to the CS was then extinguished by an alternating schedule of 30-s and 10-s long presentations of the CS without food. A trial-for-trial comparison of responding during extinction revealed a clear PREE between the groups—responding to the 100% CS extinguished sooner than responding to the 25% CS, even though the response rate to the 100% CS had been higher during training. The alternating 30-s/10-s schedule of extinction

trials allowed the researchers to perform an ingenious analysis. They were able to compare responding to the CSs at specific points across the whole extinction session. In one comparison, the 25% group had received 4 times as much cumulative CS exposure as the 100% group, but only twice as many trials. In the other comparison, the 25% group had received 4 times as many trials as the 100% group, and 8 times as much cumulative exposure. According to RET, extinction should be matched between the two groups in the first comparison, but not the second comparison (in which the 25% group should appear to extinguish more rapidly). Instead, the analysis showed that there was a difference in rate of extinction, in the form of a PREE, when the 25% group had received 4 times the amount of cumulative exposure, but there was no difference (no PREE) when the 25% CS had received 4 times as many trials. Therefore, they concluded, trials are more important than time for extinction and the PREE.

-------------------- *Figure 1 about here* --------------------

Notwithstanding the evidence described above, the duration of CS exposure may still play an important role in extinction after partial reinforcement. Bouton et al. (2014) trained both groups of rats with a fixed (10-s) CS-US interval, and this may have contributed to the difference they observed in the rate of extinction when the 25% group received 4 times as much CS exposure, but only twice as many trials, compared to the 100% group. This particular analysis specifically compared responding during each of the 10-s extinction trials in the 25% group (i.e., excluding all of the 30-s trials) with responding during every 10-s time window across all trials in the 100% group. Thus the data for the 100% group included response rates measured during the late portion of the 30-s trials, but no part of these long trials were included in the data for the 25% group. If the rats had learned that the US always occurred 10 s after the CS onset during conditioning, we should expect that their response rates would peak across those first 10 s of the trial, and then wane across the extended portion of the 30-s trials. If this is correct, then the lower levels of responding in the 100% group compared with the 25% group, using this analysis, could be the result of differences in response

timing within the trial, rather than differences in rate of extinction. Given this issue, we have sought to re-investigate the question posed by Bouton et al., as described below.

We conducted two experiments designed to compare the relevance of time versus trials in the PREE, as Bouton et al. (2014) had done, while avoiding the problem that response timing might differentially affect response rates measured during long versus short trials. Both experiments used a within-subjects design that had rats trained with both a partially reinforced CS (PRf, reinforced on 33% of trials) and a continuously reinforced CS (CRf). The two CSs were matched for their overall rates of reinforcement by making the CRf CS 3 times longer than the PRf CS (mean length of 30 s versus 10 s). This procedure reliably produces a PREE when the two CSs are extinguished using the same number and length of trials (Chan & Harris, submitted). The two CSs were then given extinction treatments, but the PRf CS was exposed 3 times more than the CRf CS. In Experiment 1, the length of each presentation was the same for the two CSs, but there were 3 times as many presentations of the PRf CS as the CRf CS. In Experiment 2, the two CSs were presented the same number of times per session, but each presentation of the PRf CS was 3 times longer than the CRf CS (30 s versus 10 s). If the difference in rate of extinction following partial versus continuous reinforcement is a function of the accumulated time of CS exposure, then we should see equivalent rates of extinction per session (i.e., no PREE) in both experiments. However, if differences in rates of extinction are specifically related to differences in the number of trials required to extinguish responding, as concluded by Bouton et al., then we should see a PREE in Experiment 2 but not in Experiment 1. Summaries of the designs of each phase in the two experiments are shown in Figures 1 and 2.

Experiments 1 and 2

Method

Subjects. For each experiment, 16 experimentally naive female Sprague Dawley rats (*Rattus norvegicus*), 8 to 10 weeks of age, were obtained from the Animal Resources Centre, Perth, Western

Australia. They were housed in groups of four in ventilated plastic tubs, measuring 40 x 40 x 46 cm (height x length x depth), located in the animal colony maintained by the School of Psychology at the University of Sydney. The animal colony room was maintained on a reverse light-cycle. The lights were switched on from 8pm - 8am. One week prior to the commencement of the experiment, rats were placed on a restricted food schedule. They had unrestricted access to water in the home tubs. Each day, each tub of rats received dry chow equal to 5% of the total weight of all rats in the tub, half an hour after the end of their daily training session. All food was consumed within 3 hr. The data of one rat in Experiment 1 were removed from the experiment due to equipment failure.

Apparatus. Thirty two Med Associates conditioning chambers were used to train and test the rats. Twenty four of the chambers (all 16 chambers in Experiment 1, and 8 chambers in Experiment 2) had the following specifications. They measured 29 x 30.5 x 24 cm (interior height x length x depth). The side walls of all chambers were constructed of Plexiglas, as was the ceiling. The end walls were constructed of aluminium. The floor of the chambers consisted of steel rods measuring 0.5 cm in diameter, spaced 1.5 cm apart parallel to the end walls. In the centre of one end wall was a recessed food magazine, with an infrared LED and sensor located at the entrance of the magazine in order to detect entries by the rat. A small metal cup (3.5 cm diameter, 0.5cm deep) was attached to the floor in the centre of the magazine; 45 mg purified rodent food pellets (Bioserve, Frenchtown, NJ) were delivered into the cup from a dispenser attached to the back of the magazine. Each chamber was enclosed in a sound- and light-resistant wooden shell. The other 8 chambers, used in experiment 2, were the same to the 24 just described with the following exceptions. They were slightly shorter, measuring 21 x 30.5 x 24 cm (interior height x length x depth), and the food cup was positioned to the left of centre, and to the right was a well that could be backfilled with liquid but was never used in this experiment. They were housed inside light and noise-attenuating PVC shells.

A fan was installed in the rear wall of the wooden shell to provide ventilation. The fan produced a background noise measuring 70 dB during operation. Experimental events were controlled and recorded automatically by computers and relays in the same room.

One visual stimulus (a steady light) and two auditory stimuli were presented from spatially separated sources. A speaker on the wall of the conditioning chamber located above and to the right of the food magazine produced white noise at 78 dB. A clicker module (Med Asssociates, product ENV 135M) was located on the end wall of the chamber, delivering four clicks per second (24 dB above background). In 24 chambers (all 16 in Experiment 1, 8 in Experiment 2), the light was provided by an incandescent bulb located on the back wall of the wooden shell produced steady light (30 cd/m^2). In the other 8 chambers of Experiment 2, the light was produced by a LED strip along the ceiling of the wooden outer shell. The white noise and steady light stimuli were counterbalanced across all rats such that one stimulus, PRf, signalled partial reinforcement during training, while the other stimulus, CRf, signalled continuous reinforcement. The clicker served as a continuously reinforced CS throughout the experiments for all rats.

Procedure. All procedures were approved by the Animal Ethics Committee at the University of Sydney (Project number 2013/6026). One day prior to the beginning of the experiment, all rats received a single magazine training session during which 24 food pellets were delivered within 60 min according to a variable-time (VT) 150-s schedule, without any stimulus presented. The rats then received daily conditioning sessions 5 days per week for a total of 35 sessions. The experiments were conducted in two phases.

Each of the first 20 sessions consisted of 18 presentations of the PRf CS, and 6 each of the CRf CS and the clicker. The duration of each CS presentation varied from trial-to-trial according to a uniform distribution. A reinforced trial consisted of the delivery of a single food pellet at stimulus offset. The PRf CS had a mean duration of 10 s (2 to 18 s) and termination of the CS coincided with delivery of food on one-third of trials (6 of 18) according to a random schedule. The CRf CS and clicker

terminated with food on every trial. The mean duration of the CRf CS was 30 s (2 to 58 s). The mean duration of the clicker was 17 s (2 to 32 s), selected as the geometric mean of 10 s and 30 s. All trial types were randomly intermixed, with a mean inter-trial interval (ITI) of 120 s.

From Sessions 21 to 35, responding to the PRf and CRf CSs was extinguished. In Experiment 1, each session consisted of 18 presentations of the PRf CS, 6 presentations of the CRf CS, and 12 presentations of the clicker. The duration for both PRf and CRf CSs was fixed at 10 s. In Experiment 2, each session consisted of 6 presentations each of the PRf and CRf CSs. The duration of the PRf CS was 30 s, while the duration of the CRf CS was 10 s. No food pellets were delivered at termination of either CS. The clicker continued to be reinforced on each trial as in the previous 20 sessions. The clicker was included in the experiments so that rats continued to receive signaled deliveries of food pellets during the extinction phase. We did this so that responding did not simply cease due to extinction of the magazine itself, which might otherwise obscure any difference in evidence for extinction of the two CSs of interest. All trial types were randomly intermixed, with a mean ITI of 100 s.

Across all sessions, the timing of each photo-beam interruption by head entry into the magazine was registered during each CS. Entries were also recorded during the 20-s pre-CS period.

Results

Experiment 1. Panel C in Figure 1 shows the mean response rates to the three CSs, and the pre-CS period, over the 20 sessions of conditioning. Rats acquired very similar levels of responding to the PRf and CRf CSs across the conditioning phase. A 2 (CS) X 20 (Session) repeated measures ANOVA, with Greenhouse-Geisser correction to degrees of freedom where necessary, showed that there was no significant overall difference in responding between the two CSs, $F(1, 14) = 1.13$, $p = .305$. There was a significant main effect of Session, $F(2.9, 40.8) = 8.20$, $p < .001$, $\eta_p^2 = .369$, but this did not interact with CS type, $F(4.8, 67) = .69$, $p = .628$. Equivalent rates of acquisition for the partially and

continuously reinforced CSs indicate that the rate of learning during acquisition is determined only by the overall reinforcer rates, as reported by Harris et al (2015).

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

The mean response rates to the three CSs over 15 sessions of extinction are shown in Panel D of Figure 2. The key comparison is between responding to the PRf CS and the CRf CS. A 2 (CS) X 15 (Session) repeated measures ANOVA showed that there was no overall difference in responding between the two CSs, $F(1, 14) = .20$, $p = .658$. There was a significant main effect of Session, $F(2.5, 35) = 10.53$, $p < .001$, $\eta_p^2 = .429$, but no significant interaction between CS and Session, $F(4.0, 56) = 1.97$, $p = .113$. To test the strength of evidence that responding to the two CSs was equal, we calculated a Bayes Factor (BF) using the method described in Rouder, Speckman, Sun, Morey, and Iverson (2009). This provided support for the null hypothesis, JZS BF = 3.51, meaning that, based on the data, it is 3.5 times more likely that there was no difference between the CSs than that there was a difference.

Experiment 2. Mean response rates to the three CSs, and during the pre-CS period, over the 20 conditioning sessions in Experiment 2 are shown in panel C of Figure 3. As was observed in Experiment 1, the rats acquired responding to the PRf CS and CRf CS at the same rate. A 2 (CS) X 20 (Session) repeated measures ANOVA showed no significant main effect of CS, $F(1, 15) = 1.03$, $p = .362$. Therefore, rats were responding equally to both the partially reinforced and the continuously reinforced CSs during acquisition. There was a significant main effect of Session, $F(1.67, 25) = 6.85$, $p < .001$, $\eta_p^2 = .313$, but no interaction between CS and Session, $F(4.3, 65) = .90$, $p = .579$.

As for Experiment 1, the comparison of interest is between responding to the PRf CS and CRf CS during extinction (Sessions 21-35). Response rates across the full 30-s presentations of the PRf CS and across the 10-s presentations of the CRF CS, as shown in Panel D of Figure 3, were subjected to a 2 (CS) X 15 (Session) repeated measures ANOVA. There was a significant main effect of CS, $F(1, 15) = 5.13$, $p = .039$, $\eta_p^2 = .255$, and a significant main effect of Session, $F(1.9, 29) = 12.28$, $p < .001$, $\eta_p^2 = $

.450, and the interaction between Session and CS fell short of the adjusted level significance, $F(1.6,$ 25) = 3.23, p = .065.

----------------- *Insert Figure 3 about here* -----------------

The analysis described in the preceding paragraph compares responding to short (10-s) presentations of the CRf CS with responding to long (30-s) presentations of the PRf CS. This difference in trial length may have introduced differences in response rates between the CSs that result from changes in responding across the length of the trial, rather than being due to differences in rates of extinction. For example, rats may respond at the same rate to both CSs during the first 10-s after CS onset, but their responding to the CRf CS may continue to increase beyond that time. This would produce a difference in mean response rate between the CSs that reflected a difference in how the response data were collected rather than differences in extinction. To address this possibility, we have performed a second analysis of the extinction data from Experiment 2. We have analysed responding during the first 10 s of each 30-s presentation of the PRf CS so that it could be compared with responding across the same time window as the CRf CS (which was always 10-s long). These data are shown in panel E of Figure 3. A 2 X 15 ANOVA confirmed that there was a significant difference in responding between the PRf CS and the CRf CS, $F(1, 15)$ = 11.15, p = .004, η_p^2 = .426. There was a significant main effect of Session, $F(3.2, 48)$ = 14.44, p < .001, η_p^2 = .490, and a significant interaction between Session and CS, $F(2.8, 43)$ = 2.96, p = .045, η_p^2 = .165.

Comparison of Experiments 1 and 2. The analyses reported above identified a significant difference in responding to the two CSs across extinction sessions in Experiment 2 but not in Experiment 1. However, the separate analyses run on each experiment do not test whether the differences between the CSs are significantly different between the two experiments. For this reason, we have run a 2x(2)x(15) ANOVA on the combined extinction data, with experiment as a group factor. (This between-experiment analysis can be justified on the grounds that the two experiments were run contemporaneously.) This analysis showed there was no overall difference (averaging across the two

CSs) in responding across extinction between the two experiments, $F(1,29) = 1.77$, $p = .194$, but there was a significant difference overall (combined across experiments) between responding to the PRf CS and responding the CRf CS, $F(1,29) = 6.77$, $p = .014$, $\eta_p^2 = 0.19$. Of most importance for the current purposes, there was a significant interaction between CS (CRf vs PRf) and Experiment, $F(1,29) = 6.92$, $p = .014$, $\eta_p^2 = 0.19$, meaning that the difference between responding to the CRf and PRf CSs was significantly greater in Experiment 2 than Experiment 1. Finally, there was a significant main effect of Session, $F(3.2,92) = 23.84$, $p < .001$, $\eta_p^2 = 0.45$, which interacted with the main effect of CS, $F(4.0,116) = 3.48$, $p = .010$, $\eta_p^2 = 0.11$, but there was not a significant 3-way interaction with the factor Experiment, $F(4.0,116) = 1.30$, $p = .273$.

Discussion

The per-session rates of extinction for the PRf and CRf CSs were equal in Experiment 1 but not in Experiment 2. Therefore, when comparing extinction of a partially reinforced CS and a continuously reinforced CS, evidence for a PREE disappears when the PRf CS is presented three times as often as the CRf CS. By contrasts, the PREE is maintained when the PRf CS is presented for three times longer than the CRf CS. In other words, the rats required more presentations, but not more time, of the PRf CS to extinguish responding at the same rate as for the CRf CS.

General Discussion

The two experiments presented here examined whether the PREE could be reduced by increasing the amount of exposure given to the PRf CS, compared to the CRf CS, during each extinction session. The two experiments differed in how that extra exposure was provided. Experiment 1 tested whether tripling the number of presentations of the PRf CS would facilitate its extinction to the point where responding did not differ between it and the CRf CS. Experiment 2 tested whether tripling the duration of each presentation of the PRf CS would facilitate its extinction such that responding to the two CSs would not differ. We observed a difference in extinction (i.e., we found a PREE) in

Experiment 2 but not in Experiment 1. Therefore, tripling the duration of each extinction exposure to the PRf CS did not increase its rate of extinction to match that of the CRf CS. On the other hand, tripling the number of extinction exposures to the PRf CS did accelerate its extinction to the point that it extinguished at the same rate as the CRf CS. In the latter case, a Bayesian analysis confirmed there was moderate support for the null hypothesis, showing it was three and a half times more likely that responding to the two CSs was equal than that they differed.

The results confirm the conclusions reached by Bouton et al. (2014). Those authors found that evidence for a PREE could be eliminated if they compared responding on every fourth extinction trial in the partially reinforced group with responding on every extinction trial in the continuously reinforced group. This is clearly consistent with our result in Experiment 1. Bouton et al. also found that the PREE was present if they compared responding after the partial group had received 4 times as much cumulative exposure to the CS as the continuous group (see Figure 1). This result is consistent with what we found in Experiment 2.

The consistency in findings between the present experiments and those described by Bouton et al. (2014) are notable because there are important differences in how the experiments were run. Bouton et al. used identical extinction schedules for their CRf and PRf groups, which allowed them to make direct trial-for-trial comparison between the two groups and show an overall PREE. However, in order to compare responding in the continuous group with responding in the partial group after the latter had received more exposure to the CS, they had to sample from trials at very different periods within the long extinction session, and compare responding at different time points within a trial. These differences between the groups could have contributed to the results of those particular analyses. As Drew et al. (2004) showed, responding within an extinction session is affected by the length of the trials, especially when the trial length differs from the CS-US interval during conditioning. In contrast, our experiments gave rats more, or longer, extinction trials with the PRf CS than with the CRf CSs, which meant that the key comparisons between the two CSs used response

data collected concurrently across the full length of each session. Moreover, we were able to compare responding between the PRf and CRf CSs sampled from the same portion of each trial. Nonetheless, because our analyses were based on whole-session averaging of response rates, we cannot distinguish between effects of trial length or number on the extinction of responding that develops within a session versus extinction that carries over between sessions (cf. Drew et al., 2004).

That fact that more trials, but not longer trials, resulted in faster extinction to the PRf CS is also consistent with the findings reported by Harris and Andrew (2017). They trained rats to respond to two identical continuously reinforced CSs, after which both were extinguished using either many short extinction trials or fewer long extinction trials, matched for total duration. They found that responding extinguished faster with many short extinction trials. Indeed, extinction was unaffected by large differences in CS duration during extinction, as long as the number of trials was held constant. This result, along with the results of the current experiments, suggests that extinction is event-based rather than continuous. That is, each presentation of a CS during extinction appears to be treated as a single episode, regardless of the trial duration, and when a trial is not reinforced, the animal re-evaluates responding towards the CS.

An implication of the present findings, as well as those reported by Bouton et al. (2014), is that the rate of extinction is strictly determined by the per-trial reinforcement rate learned during acquisition. In Experiment 1, there were 3 times as many PRf trials as CRf trials per food pellet during conditioning. Subsequently, those CSs extinguished at the same overall rate when there were 3 times as many trials of the PRf CS as trials of the CRf CS. In similar vein, Bouton et al. found equivalent extinction of their two CSs when they compared responding on every fourth trial of the PRf CS with responding on every trial of the CRf CS, which matched the difference between those CSs in number of trials per reinforcement during conditioning. An intriguing implication of this is that extinction tracks how often the animal should expect to receive the US. If, during conditioning, the US only occurs every n presentations of the CS, then responding will take n times as many trials

to extinguish. The evidence for this claim is at this stage somewhat circumstantial, and would require a larger investigation into the effect of reinforcement rate on the rate of extinction. Nonetheless, this conclusion is consistent with Capaldi's (1966, 1994) Sequential account which argues that animals learn the sequence of trial outcomes, both reinforced and non-reinforced, during conditioning, and therefore the unit of learning during extinction will be measured in chunks of trials that make up a sequence. The challenge for this account is to explain how these processes can support a PREE in a within-subjects design, such as used here (see also Chan & Harris, submitted; Rescorla, 1999). In a within-subjects design, sequences of reinforced and non-reinforced trials of a partially reinforced CS are randomly intermixed with trials of other consistently reinforced CSs. In order to show a PREE, the animal must keep track of the separate sequences of all CSs, which would require very substantial working memory capacity. An alternative view is that, at least in these designs, animals are not learning about sequences of trials, but are learning in a trial-by-trial fashion about the probability of reinforcement per trial for each CS. In this case, extinction would depend on the animal detecting that the per-trial rate of reinforcement of the CS has changed, which would take more trials if the per-trial rate were low. This logic is similar to Gallistel's (2012) description of extinction as a rational process tuned to detecting a change in the rate of reinforcement, except that we argue extinction depends on knowledge of the probability of reinforcement per trial rather than per unit of cumulative time.

References

Amsel, A. (1962). Frustrative non-reward in partial reinforcement and discrimination learning: Some recent history and a theoretical extension. *Psychological Review, 69*, 306-328.

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.

Capaldi, E. J. (1966). Partial reinforcement: A hypothesis of sequential effects. *Psychological Review, 73*, 459-477.

Capaldi, E. J. (1994). The sequential view: From rapidly fading stimulus traces to the organisation of memory and the abstract concept of number. *Psychonomic Bulletin & Review, 1*, 156-181.

Chan, C. K. J., & Harris, J. A. (submitted). The role of trials versus reinforcement rate in the acquisition and extinction of conditioned responding.

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

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

Drew, M. R., Yang, C., Ohyama, T., & Balsam, P. D. (2004). Temporal specificity of extinction in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes, 30*, 163-176.

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

Gallistel, C. R. (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.

Golkar, A., Bellander, M., & Öhman, A. (2013). Temporal properties of fear extinction - Does time matter? *Behavioral Neuroscience, 127*, 59-69.

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., & Andrew, B. J. (2017). Time, trials, and extinction. *Journal of Experimental Psychology: Animal Learning and Cognition, in press*.

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., & 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.

Haselgrove, M., & Pearce, J. M. (2003). Facilitation of extinction by an increase or a decrease in trial duration. *Journal of Experimental Psychology: Animal Behavior Processes, 29*, 153-166.

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.

Lovibond, P. F., Davis, N. R., & O'Flaherty, A. S. (2000). Protection from extinction in human fear conditioning. *Behaviour Research and Therapy, 38*, 967-983.

Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review, 104*, 241-265.

Mackintosh, N. J. (1974). *The psychology of animal learning*. London, UK: Academic Press.

Nevin, J. A. (1988). Behavioral momentum and the partial reinforcement effect. *Psychological Bulletin, 103*, 44-56.

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

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

Redish, A. D., Jensen, S., Johnson, A., & Kurth-Nelson, Z. (2007). Reconciling Reinforcement Learning Models With Behavioral Extinction and Renewal: Implications for Addiction, Relapse, and Problem Gambling. *Psychological Review, 114*, 784-805.

Rescorla, R. A. (1997). Response-inhibition in extinction. *Quarterly Journal of Experimental Psychology, 50B*, 238-252.

Rescorla, R. A. (1999). Within-subject partial reinforcement extinction effect in autoshaping. *Quarterly Journal of Experimental Psychology, 52B*, 75-87.

Rescorla, R. A. (2003). Protection from extinction. *Learning & Behavior, 31*, 124-132.

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

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

Shipley, R. H. (1974). Extinction of conditioned fear in rats as a function of several paramaters of CS exposure. *Journal of Comparative and Physiological Psychology, 87*, 699-707.

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

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

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.
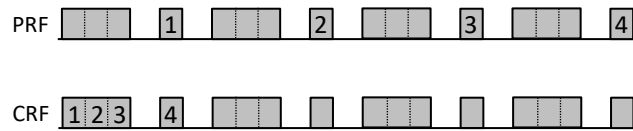
Figure 1. Trial structure during extinction for the partially reinforced (PRF) group and the continuously reinforced (CRF) group in Bouton et al. (2014). CS presentations (shown here as grey rectangles) alternated between long (30 s) and short (10 s) presentations. A "time expectancy" analysis compared responding on every short trial for the PRF group with responding on every 10-s segment of the CS in the CRF group. Numbers identify the corresponding CS segments between the two groups. This meant that the CRF group was compared with the PRF group after the latter had received 4 times the length of CS exposure. For example, the fourth 10-s segment of the CRF group was compared with the fourth 10-s trial of the PRF group, at which point the latter had received 160-s of cumulative exposure to the CS.
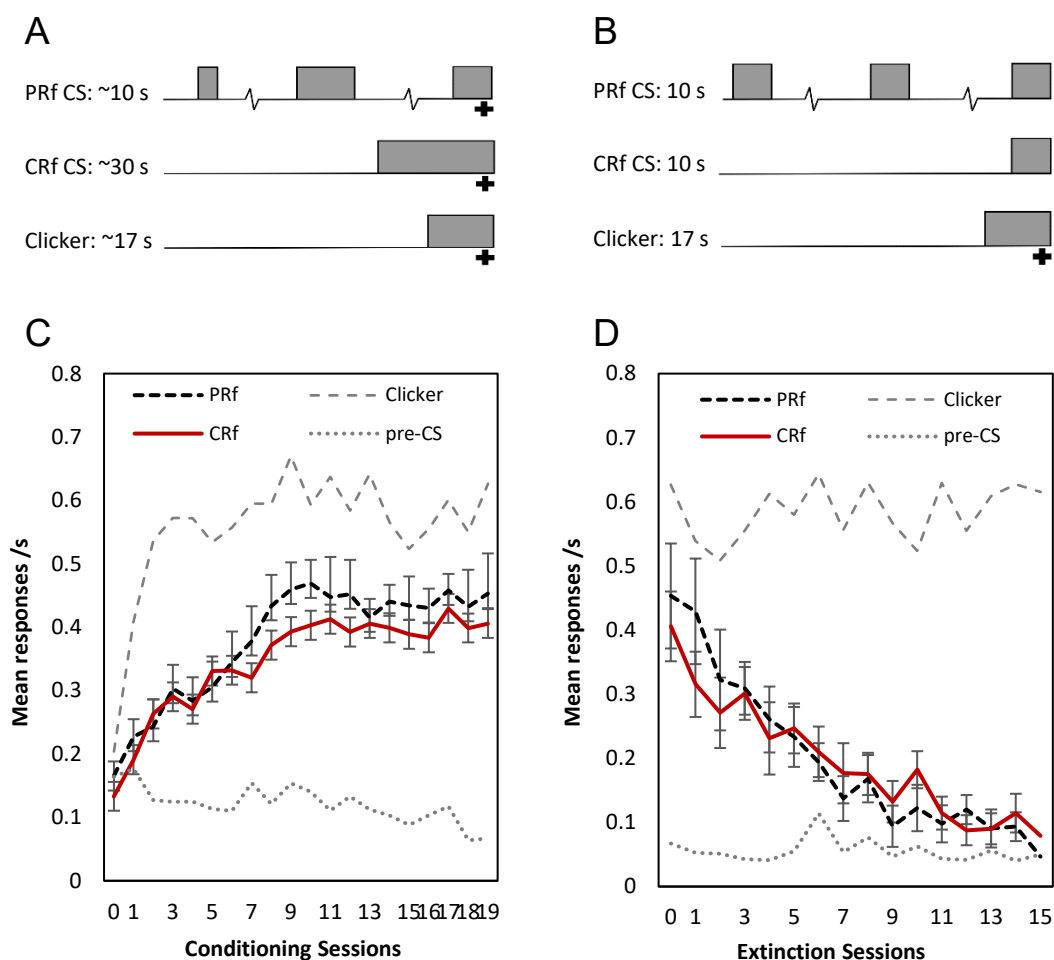
Figure 2. TOP: Summary of designs for acquisition (panel **A**) and extinction (panel **B**) in Experiment 1. The PRf CS was conditioned with food (**+**) on 33% of trials after a mean duration of 10 s; the CRf CS was reinforced on 100% of trials with a mean duration of 30 s; the clicker was reinforced on 100% of trials, with a mean duration of 17 s. During extinction, both the PRf CS and CRf CS were presented for a fixed duration of 10 s and neither CS was reinforced. There were 3 presentations of the PRf CS for every one of the CRf CS. The clicker continued to be reinforced during extinction of the PRf and CRf CSs. Bottom: Mean response rates during the PRf CS, CRf CS, and the clicker, as well as during the pre-CS interval, over 20 sessions of acquisition (panel **C**) and over 15 sessions of extinction (panel **D**). Vertical bars show the standard error of the mean difference between responding to the PRf CS and responding to the CRf CS.
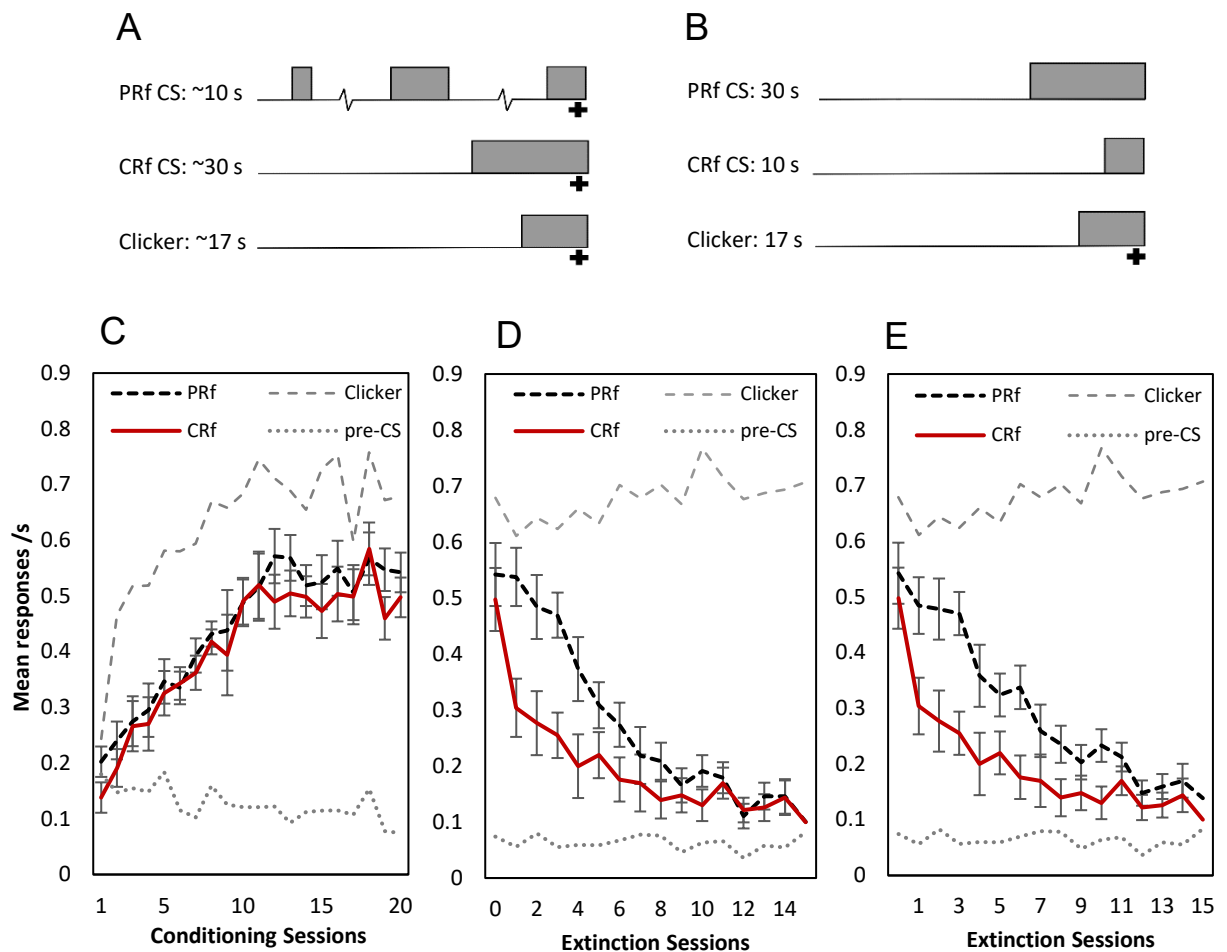
Figure 3. TOP: Summary of designs for acquisition (panel **A**) and extinction (panel **B**) in Experiment 2. As in Experiment 1, the PRf CS was conditioned (**+**) on 33% of trials after a mean duration of 10 s; the CRf CS was reinforced on 100% of trials after a mean duration of 30 s; the clicker was reinforced on 100% of trials after a mean duration of 17 s. During extinction, both the PRf CS and CRf CS were presented the same number of times without reinforcement, and presentations of the PRf CS were 3 times longer than presentations of the CRf CS (30 s versus 10 s). The clicker continued to be reinforced during extinction of the PRf and CRf CSs. Bottom: Mean response rates during the three CSs, as well as during the pre-CS interval, over 20 sessions of acquisition (panel **C**) and over 15 sessions of extinction (**D** and **E**). Response rates for the PRf CS are averaged across the full 30 s of each extinction trial in **D**, but are shown only for the first 10 s of each trial in **E** so that responding to this CS is compared across the same time window as for the CRf CS. Vertical bars show the standard error of the mean difference between responding to the PRf CS and CRf CS.