Pre-publication copy.

Published in: The Journal of Experimental Psychology: Animal Learning & Cognition, vol 45, pp 43-58, 2019.
©American Psychological Association, 2019. 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, <http://dx.doi.org/10.1037/xan0000190>.

The Partial Reinforcement Extinction Effect: The Proportion of Trials Reinforced During Conditioning Predicts the Number of Trials to Extinction

C. K. Jonas Chan, Justin A. Harris

The University of Sydney

Author Note

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

Correspondence concerning this article should be addressed to Justin A. Harris,
School of Psychology, Griffith Taylor Building, The University of Sydney, Sydney, 2006,
Australia. E-mail: Justin.Harris@sydney.edu.au

Abstract

Four experiments compared the extinction of responding to a continuously reinforced (CRf) conditioned stimulus (CS) consistently reinforced on every trial, with extinction of responding to a partially reinforced (PRf) CS that had been inconsistently reinforced. To equate the acquisition of responding between the two CSs, the average duration of the CRf CS was extended so that it scheduled the same overall rate of reinforcement per unit time as the PRf CS. Experiment 1 used a within-subjects design to compare the rates of extinction for a 10-s PRf CS reinforced on 33% of trials versus a 30-s CRf CS. Experiment 2 made the same comparison but using a between-subjects design. Experiment 3 compared extinction in a group trained with a 10-s PRf CS reinforced on 20% of trials and a group trained with a 50-s CRf CS. Experiment 4 compared the rates of extinction following two partial reinforcement schedules, a 10-s PRf CS reinforced on 33% of trial versus a 20-s CRf CS reinforced on 66% of trials. In each experiment, responding took longer to extinguish for the CS that scheduled a lower per-trial probability of reinforcement. Modelling of individual extinction curves using Weibull functions indicated that the latency to initiate extinction was directly related to the per-trial probability of reinforcement learned during acquisition. For example, compared to training with a CRf CS, rats reinforced on 33% of trials took approximately three times as many trials to initiate extinction, and rats reinforced on 20% of trials took five times as many trials to initiate extinction. These results provide support for trial-based accounts of extinction (e.g. Capaldi, 1967), whereby rats learn about the expected number of trials per reinforcer, and extinction depends on the number of expected reinforcers that have been omitted rather than on the number of extinction trials per se.

Keywords: Partial reinforcement extinction effect, timing models, associative models, trial-wise models, Weibull function.

A well-known phenomenon in Pavlovian associative learning is the Partial Reinforcement Extinction Effect (PREE). It refers to the fact that conditioned responses (CRs) extinguish more slowly if the conditioned stimulus (CS) has been conditioned using a “partial” reinforcement schedule (i.e., the CS was inconsistently paired with the unconditioned stimulus, US) than if it had been conditioned using a “continuous” reinforcement schedule (i.e., the CS was consistently paired with the US on every trial). Theoretical accounts for the PREE often appeal to the differences in generalisation between acquisition and extinction stages (Mackintosh, 1974). According to Amsel's (1958) frustration theory, when animals have been trained on a continuous reinforcement schedule, they experience frustration at the absence of the expected US during the extinction phase. This reduces generalisation of conditioning from the acquisition phase, and thus promotes the loss of responding under extinction. If animals are instead trained on a partial reinforcement schedule, frustration at the absence of the US is part of the conditioning context and thus supports the generalisation of conditioning to the extinction phase, thereby reducing response loss. Capaldi's (1966, 1994) sequential theory offers a related account of the PREE. It argues that animals trained on a partial reinforcement schedule learn that the US follows a sequence (or “chunk”) of non-reinforced trials, and this makes their experience of the CS during acquisition more similar to what they experience during extinction, which serves to perpetuate responding.

The PREE is often regarded as a special case, and indeed the prevailing theories described above assume that partial reinforcement introduces an additional element into what is learned, and this element is responsible for promoting the continuation of responding under extinction. A somewhat different stance has been adopted more recently by Gallistel and Gibbon (2000) as part of their Rate Estimation Theory. RET denies that the PREE is a special case of learning, arguing that it is a consequence of different overall rates of reinforcement

commonly used in partial reinforcement schedules. Typically, experiments studying the PREE create partial reinforcement schedules by either removing reinforcers from an existing continuous schedule, or by inserting additional non-reinforced CS trials into an existing continuous schedule. Both of these methods increase the total time that the CS is present per USs. This means that the expected time to reinforcement is longer under the partial schedule than the continuous schedule, and therefore, according to RET, animals should take longer to extinguish because more CS time must pass before each expected US is missed. Moreover, in this model the amount of time that the CS must be presented in order for responding to extinguish should be directly proportional to the total length of time that the CS was present per US during acquisition, and this should be true whether the reinforcement schedule was partial or continuous. For example, if USs were scheduled to occur twice as frequently for one CS than for another CS, responding to the second CS should take twice as long to extinguish. As such, RET argues that the PREE exists only because most experiments schedule a lower overall rate of reinforcement for their partial reinforcement schedules than their continuous reinforcement schedules.

One solution to the problems outlined above when comparing rates of extinction after partial and continuous reinforcement schedules is to create reinforcement schedules that produce comparable rates of responding at the end of acquisition. Haselgrove, Aydin and Pearce (2004) compared one group of rats trained with a continuously reinforced CS that was either 10-s or 20-s long, with a group trained with a partially-reinforced CS that was either 10-s or 5-s long, reinforced on 50% of trials. With this arrangement, both CSs were reinforced every 15 s, and responding was comparable between the two CSs at the end of acquisition. During extinction, a PREE was found despite identical overall rates of reinforcement. These results contradict the prediction of RET that extinction should progress

equally for two CSs reinforced at equal overall rates of reinforcement regardless of whether they are partial or continuously reinforced.

Other recent experiments have tested the second prediction of RET—that the difference in the rates of extinction between a continuously reinforced and a partially reinforced CS will be equal to the difference between the overall rate of reinforcement during acquisition. An implication of this claim is that the PREE should disappear if a partially reinforced CS is subjected to more extinction. Other models of the PREE (Capaldi, 1966; Amsel, 1958) also predict that evidence for the PREE will be reduced if the partially-reinforced CS receives extra extinction trials. However, RET specifically predicts that the PREE will be reduced if the partial CS receives extra extinction trials or longer extinction trials, and the PREE will disappear entirely if the extra exposure to the partial CS matches the extra time that it was presented per US during conditioning (Gallistel, 2012). Bouton, Woods, and Todd (2014) tested RET's prediction. They compared extinction in a group of rats trained with a 10-s CS on a 25% reinforcement schedule with a group trained with a 10-s CS on a 100% reinforcement schedule. Responding in both groups was then extinguished using an alternating schedule of 30-s and 10-s non-reinforced presentations of the CS. Not surprisingly, when compared trial-for-trial, responding was slower to extinguish in the partially-reinforced group than the continuously reinforced group. The alternating schedule during extinction allowed the researchers to make more targeted analyses of the progress of extinction. One analysis allowed the partially-reinforced group to accumulate 4 times the length of CS exposure as the continuous group, but only twice as many trials. According to RET, this analysis should yield equivalent rates of extinction between the group, but the results showed otherwise—there was still less extinction in the partial group than the continuous group. A second analysis allowed the partial group to accumulate 4 times as many extinction trials as the continuous group, at which point they had also accumulated 8 times

the length of CS exposure. According to RET, this analysis should yield more evidence for extinction in the partial group, yet the results showed very similar rates of extinction between the two groups.

Bouton et al.'s (2014) conclusions have been corroborated by a recent experiment that we conducted (Chan & Harris, 2017). We trained two groups of rats to respond to two CSs, one reinforced continuously, with a mean duration of 30 s. The second CS was reinforced on 33% of trials, with a mean duration of 10 s. These schedules ensured that responding was acquired at very similar rates and to the same asymptotic level for the two CSs. During extinction, one group received an equal number of extinction trials of the two CSs, but extinction trials for the partially reinforced CS were 3 times longer than for the continuously reinforced CS. The other group received extinction trials that were the same duration for both CSs, but there were 3 times as many trials of the partially reinforced CS than the continuous CS. According to RET, there should be no evidence for a PREE in either group when responding is averaged across each session. However, we still found evidence for a PREE in the group that received the same number of extinction trials for each CS (even though presentations of the partial CS were 3 times longer), but there was no PREE in the group that received 3 times the number of extinction trials with the partial CS. Thus, like Bouton et al., our data shows clearly that the number of trials is more important than the cumulative exposure of the CS during extinction (see also Harris & Andrew, 2017, for the same conclusion about extinction after continuous reinforcement). While these results are inconsistent with RET's claim that the rate of extinction is directly proportional to the rate that the CS was reinforced per unit time, they do nonetheless support the idea that animals track during extinction how often per trial they should be expecting the US. For example, the rats trained by Chan and Harris (2017) would have learned to expect reinforcement on one in

every 3 trials for the partial CS, and therefore required 3 times as many trials to extinguish responding to that CS.

The experiments by Bouton et al (2014) and Chan & Harris (2017) are consistent with the claim that animals learn how many non-reinforced trials occur for every reinforced trial during conditioning, as proposed in Capaldi's (1966, 1994) Sequential Theory, and that this determines how often they expect the US and therefore how many trials it takes for their responding to extinguish when all reinforcement is ceased. However, that evidence is based on a null result, that there is no longer a significant PREE after adding in the requisite number of extinction trials of the partially-reinforced CS. Stronger evidence would require that a partial and a continuously reinforced CS receive identical extinction treatments and that the partial CS takes exactly X times as many trials to extinguish if it had been reinforced on one in every X trials during conditioning. The current experiments were designed to provide this sort of evidence.

One challenge when attempting to compare the rates of extinction between a partially reinforced and a continuously reinforced CS is that the asymptotic response rates at the end of conditioning can differ greatly between the two schedules (e.g. Harris et al. 2015). As a result, interpretation of the difference in rates of extinction is confounded by differences in the rates of responding prior to extinction. Previous demonstrations of the PREE have often relied on the initially higher rate of responding to the continuously reinforced CS dropping below the response rate of the partially reinforced CS during extinction. This is a conservative estimation of the strength of the PREE, and it precludes us from determining how much more slowly the partially reinforced CS is extinguishing because the initial levels of responding are not comparable. Some have compared rates of extinction by first normalising the response rates—expressing response rates during extinction as proportions of baseline response rates at the end of acquisition (e.g., Gibbon, Farrell, Locurto, Duncan, &

Terrace, 1980; Nevin, 1988). However, this method implicitly assumes that the response scale is linear, such that, for example, a 50% decrease in responding is equivalent from any initial rate of responding. In many cases, this is not a safe assumption. To circumvent this problem, the current experiments adopted a procedure that we have used in previous experiments to establish equivalent levels of conditioned responding between a partially and a continuously reinforced CS (e.g. Chan & Harris, 2017; Harris et al., 2015). To achieve this, the length of each CS presentation varied randomly from trial to trial (so that the rats could not time their responding to the expected moment of the US), and the mean length of the continuous CS was longer than the mean length of the partial CS such that the two CSs were matched for their overall rates of reinforcement per unit time. By equating the acquisition of responding to the two CSs, we were able to compare their rates of extinction without the confounding problem that the initial levels of responding differed between the two CSs.

If the rats learn how many non-reinforced trials occur for each reinforced trial during conditioning, and if during extinction they keep track of the number of trials that should have been reinforced, then the ratios between rates of extinction should match the ratios of the per-trial probabilities of reinforcement in each experiment. In other words, the strength of the PREE should be determined precisely by the thinning factor applied to the partially reinforced schedule during acquisition. Furthermore, we should be able to observe a difference in rate of extinction when comparing between two partially reinforced schedules, but with different per-trial probabilities of reinforcement (as shown by Rescorla, 1999).

Experiment 1 was conducted using a within-subjects design that compared the rates of extinction for a 10-s CS, reinforced on 33% of trials, and a 30-s CS reinforced on every trial. Experiments 2, 3, and 4 were conducted using a between-subjects design, using different reinforcement schedules for the partial CS in each experiment, in order to ascertain more broadly how the rate of reinforcement affected the rate of extinction. Experiment 2 compared

extinction of a 10-s CS reinforced on 33% of trials with extinction of a 30-s CS reinforced on every trial (= 3:1 ratio). Experiment 3 compared a 10-s CS reinforced on 20% of trials with a 50-s CS reinforced on every trial (= 5:1 ratio). Experiment 4 compared between two partially reinforced CSs: a 20-s CS reinforced on 33% of trials and a 40-s CS reinforced on 66% of trials (= 2:1 ratio). Extinction curves derived from the experiment were fitted using a Weibull function in order to determine the rates of extinction for individual subjects.

Experiment 1

Experiment 1 compared the rates of extinction of a continuously reinforced CS against a partially reinforced CS. Based on our previous studies (Chan & Harris, 2017; Harris & Kwok, 2018), we should observe a partial reinforcement extinction effect (PREE). Specifically, responding to the continuously reinforced CS should extinguish faster than responding to the partially reinforced CS.

One notable feature of this experiment is that it uses a within-subjects design. It is sometimes claimed that the PREE cannot be found using within-subjects designs (e.g., Pearce, Redhead, & Aydin, 1997). Indeed, major theories of the PREE, such as Amsel's (1958) Frustration Theory, conclude that a PREE cannot be observed in within-subjects experiments because the experience of frustration following non-reinforced trials with the PRf CS would promote generalisation of responding during the extinction phase for both CSs equally. Our demonstrations of a PREE using a within-subjects design confirm similar demonstrations (e.g., Crawford, Steirn, & Pavlik, 1985; Mellgren & Dyck, 1972; Rescorla, 1999), and show that the PREE is governed by stimulus-specific learning processes. That is, whatever animals learn about non-reinforcement during conditioning, it must be specific to the partially-reinforced CS in order that it selectively impacts on responding to that CS alone during extinction. One key factor in Rescorla's design was that the animals continued to

receive reinforcement of another CS during the extinction phase. We, too, included an additional CS in our within-subjects designs, such that this CS was reinforced through the acquisition phase and continued to be reinforced when the target CSs were undergoing extinction. We reasoned that, by maintaining delivery of the US during trials with this extra CS, we prevented the magazine itself from being extinguished which might otherwise result in rapid and global response loss, and thereby obscure any difference in the rates of extinction between the PRf CS and the CRf CS. Accordingly, in Experiment 1 of the present series, we included a third CS, a clicker, that was reinforced throughout the entire experiment.

Method

Subjects. Thirty-two experimentally naive female Sprague Dawley rats (*Rattus norvegicus*) were obtained from the Laboratory Animal Services breeding unit at the University of Adelaide, South Australia. Subjects were 8 to 10 weeks of age at the beginning of the experiment. Rats 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.

Apparatus. Thirty-two Med Associates conditioning chambers were used to train and test the rats. Twenty-four of the chambers measured 29 X 30.5 X 24 cm (height X length X depth). The side walls and ceiling of the chambers were constructed of Plexiglas. 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. One end wall contained a recessed food magazine in the centre, 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 measuring 3.5 cm in diameter and 0.5cm deep was attached to the floor of the magazine. A dispenser was attached to the back of the magazine, which delivered 45 mg food pellets (purified rodent pellets; Bioserve, Frenchtown, NJ). The entire chamber was enclosed within a sound- and light-resistant wooden shell. The remaining 8 chambers were the same as the 24 described, with the following differences. They measured 21 X 30.5 X 24 cm (height X length X depth), and the food cup was positioned left-of-centre. A well was located right of centre. The well could be backfilled with liquid but was not used in this experiment. These chambers 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. Two auditory stimuli (white noise and clicker) and one visual stimulus (steady light) was used in this experiment. 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 Associates, product ENV 135 M) was located on the end wall of the chambers, delivering four clicks per second (24 dB above background). For the first 24 chambers, steady light was provided using an incandescent bulb located on the back wall of the wooden shell (30 cd/m²). For the 8 remaining chambers, steady light was produced by an LED strip along the ceiling of the wooden outer shell. The white noise and steady light stimuli were counterbalanced across all rats, with one stimulus signalling partial reinforcement (the PRf CS) and the other signalling continuous reinforcement (the CRf CS). For all rats, the clicker was continuously

reinforced throughout the entire experiment. Reinforcement consisted of a single food pellet delivered at the offset of the stimulus.

Procedure. All procedures were approved by the Animal Ethics Committee at the University of Sydney (Project number 2016/1070). One day prior to the beginning of the experiment, 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. Across all sessions, the timing of each photobeam interruption by head entry into the magazine was recorded during each CS, and also during the 20-s pre-CS period.

Rats received a total of 35 sessions, conducted in two phases. Each of the first 20 sessions (acquisition phase) consisted of 18 presentations of PRf, and 6 presentations each of CRf and the clicker. For rats 1 to 16, the clicker was only present from session 2 onwards. CS durations varied from trial to trial according to a uniform distribution. PRf was reinforced on one third of trials (6 of 18), while CRf and clicker were reinforced on every trial. The mean duration of PRf was 10 s (2 to 18 s), and the mean duration of CRf was 30 s (2 to 58 s). The mean duration of the clicker was 17 s (2 to 32 s). Trials were intermixed randomly, with a mean ITI of 120 s. The mean total session length was 65 min.

Sessions 21 to 35 (extinction phase) consisted of 15 reinforced presentations of the clicker CS, and 10 non-reinforced presentations of each of PRf and CRf. The mean CS durations for both PRf and CRf was 17 s, varying trial to trial according to a uniform distribution (2 to 32 s). Trials were intermixed randomly, with a mean ITI of 110 s. The mean total session length was 66 min.

Results and Discussion

Data for Rats 2 and 18 were removed from analysis as they did not acquire response rates to the light and noise that were above the pre-CS rate, suggesting a failure to condition to those CSs. Data for Rat 11 was also removed due to equipment failure on Days 19, 20 and 21. For the remaining rats, elevation scores for both the acquisition phase and the extinction phase were calculated for responding to PRf and CRf by subtracting the Pre-CS rate from the response rate during the CS.

Panel A in **Figure 1** shows mean elevation scores for all 29 rats during the 20 acquisition sessions. A 2 (CS) X 20 (Session) repeated measures ANOVA with Greenhouse-Geisser correction to degrees of freedom where necessary revealed a significant main effect of Session during acquisition, $F(3.6, 100.0) = 34.423, p < .001, \eta^2_p = .551$. A significant interaction between CS and Session was found, $F(19, 153.4) = 2.251, p = .047, \eta^2_p = .074$. No main effect of CS was found, $F(1, 28) = .570, p = .562$.

The comparison of interest was the rate of responding to the PRf and CRf CSs during extinction. Elevation scores during the 15 sessions of extinction are shown in panel B of **Figure 1**. A 2(CS) X 15(Session) repeated measures ANOVA revealed a significant main effect of Session during extinction, $F(2.4, 67.7) = 53.004, p < .001, \eta^2_p = .654$, confirming that responding declined across the extinction phase. A significant main effect of CS was found, $F(1, 28) = 39.397, p < .001, \eta^2_p = .585$, and there was a significant interaction between CS and Session, $F(4.5, 127.3) = 11.474, p < .001, \eta^2_p = .291$. These results indicate that responding to the PRf and CRf CSs extinguished at different rates, consistent with the PREE.

Model fitting analysis of individual extinction curves. Extinction curves of group-averaged data may not accurately reflect each individual subject's extinction to each CS. Gallistel et al. (2000) showed that individual learning curves can differ; some subjects show an abrupt change in responding, whereas others show a gradual learning curve. We examined

individual extinction curves by fitting the inverse of the cumulative distribution of a Weibull function to the extinction data for each subject in each experiment. Cumulative Weibull functions have been used in the past to describe the acquisition of conditioned responding (e.g., Gallistel et al, 2004; Harris, 2011). Here we have used the inverse to describe the decrease in responding across extinction. The function is expressed in Equation 1.

$$R_t = A \cdot 2^{\left(-\left(\frac{t}{L}\right)^s\right)} \quad \text{Equation 1}$$

It describes a decreasing function where R represents responding at trial t. The parameters L and s determine the shape of the function, which can vary in form from a continuously decreasing exponential decay function (when $s = 1$, and L sets the rate of decay) to a sharp step function (when s is large, and L determines the position of the step along the abscissae). Thus, this function can provide a relatively objective description of extinction for an individual subject, as it can take the form of a gradual extinction curve, or a curve that decreases abruptly. In our modelling, L and s are free parameters (≥ 1) used to fit Equation 1 to the response rate data across extinction. Figure 2 shows two examples of per-trial responding during extinction for individual subjects and their respective fitted functions from Equation 1.

The parameter A in Equation 1 sets the starting height of the function, and represents the estimated response rate at the end of acquisition. In our modelling, A was not used as a free parameter when fitting Equation 1 to the extinction data, but was obtained for each rat from the response rates in the acquisition phase. It was derived by fitting a cumulative Weibull function (equal to A minus the right-hand side of Equation 1) to the rats' own elevation scores across the acquisition phase. For this operation, A, L, and s were free parameters, and the obtained value of R_t on the last trial of acquisition provided the estimate of the terminal level of responding that was used as A when fitting Equation 1 to the extinction data. The CSs in our experiment are matched for overall rates of reinforcement, so

the response rates at the end of acquisition for both the partial and continuous CSs (parameter A in equation 1) should be comparable. This is confirmed in panel C of Figure 1, which plots A for the Partial CS against A for the Continuous CS for each rat. The best-fitting regression line to these values ($y=0.99x + 0.05$) is very close to the identity line ($y=x$), showing that the values of A were comparable for the two CSs. Therefore, Weibull functions describing the rates of extinction for the partially reinforced and continuously reinforced CSs in each experiment should be directly comparable because they start at the same part of the response scale. To compare the rates of extinction between CSs in Experiment 1, and between groups in Experiments 2 to 4, we measured the mean number of trials taken for each rats' fitted extinction function to reach a particular level, defined as a percentage of the initial response rate. For example, to calculate the number of trials taken for responding to have extinguished by 10 %, we identified the first trial on which the fitted function was lower than 90% of its starting height (A). Panel D of Figure 1 shows the mean number of trials taken to reach 10%, 20%, 30%, up to 90% extinction for both CSs in Experiment 1. As the figure shows, it took more trials for responding to reach each extinction criterion for the Partial (33%) CS than for the Continuous (100%), consistent with the PREE. This was confirmed using a repeated measures ANOVA, revealing a significant main effect of CS ($F(1,28) = 80.152, p < .001, \eta^2_p = .741$) and percentage of extinction ($F(1.03, 28.8) = 150.55, p < .001, \eta^2_p = .843$). A significant interaction was found between CS and extinction, $F(1,29.7) = 21.88, p < .001, \eta^2_p = .439$. Panel D of Figure 1 also plots the ratio of trials to each criterion between the Partial CS and the Continuous CS. This shows that it took 2.9 times as many trials for responding to the Partial CS to reach the first extinction criterion (when the Weibull function had decreased by 10% from the starting value) compared with the number of trials for responding to the Continuous CS to reach this criterion. This value is very close to our prediction of a 3-fold difference in rate of extinction. However, it is also clear that this ratio gets smaller as the

extinction criterion increases, and the ratio is less than 2:1 for extinction to each CS to be 90% complete.

The within-subjects design of Experiment 1 meant that we could directly compare how quickly each rat extinguished responding to the PRf CS and the CRf CS. The scatterplot in panel E of Figure 1 shows the relationship between the number of trials each rat took to reach 10% of extinction for the PRf and CRf CSs. A linear regression showed there to be a strong relationship between the two CSs, $R^2 = .6096$, $F(1, 27) = 42.16$, $p < .001$. However, from inspection of the scatterplot it is clear that a single data point was responsible for most of the apparent correlation: Rat 31 took many more trials than any other rat to start extinguishing responding to both CSs. After excluding this single data point, there was no longer any correlation between CRf and PRf extinction, $R^2 = .009$, $F(1, 26) = .2308$, $p = .635$. Therefore, unlike the level of responding acquired to each CS over Phase 1, for which there was a strong correlation across rats, there appeared to be little or no systematic relationship in the rate at which responding to the two CSs extinguished.

The values shown in panel D of Figure 1 represent the mean number of trials to reach each criterion. The ratio of these values for the 10% criterion was very close to the 3:1 ratio predicted based on the reinforcement schedules of the two CSs. However, this result, based on the means, may not represent the relationship across the full distribution of the data, particularly if the data are skewed. A more complete presentation of the relationship is shown in the Q-Q plot in panel F of Figure 1. For this plot, the trials to the 10% criterion for 28 rats (excluding the outlier Rat 31) were rank ordered and divided into 7 equal quantiles (the first quantile being the average of the 4 smallest values, and the last quantile being the average of the highest 4 values). This was done separately for the two CSs, and then each quantile of the Partial CS is plotted against the corresponding quantile of the Continuous CS. The 7 quantiles lie close to the regression line $y = 2.82x$, $R^2 = 0.974$. The slope of this line (2.82) is close to

the ratio of the means (2.9) and again consistent with the 3:1 ratio predicted. The fact that the data are well explained by a line with zero intercept shows that the relationship between the values for the two CSs is very consistent across the full distribution of the data.

Non-Weibull analysis of individual extinction data. The model-fitting analysis described above makes some assumptions about the form of the extinction curve based on the shape of the Weibull function. As such, our choice of this particular model may have influenced our measure of trials to extinction. Therefore, we have performed an additional analysis of the extinction data that does not involve fitting a curve to the response rate. In this analysis, the number of trials to commence extinction is visually read off a cumulative record of the elevation scores of each rat. We used the cumulative record because it shows overall trends in the data that are often obscured by the variance in the per-trial response rate. First, for each rat, an elevation score was calculated for every trial as the mean response rate per second in that trial minus the rats' average pre-CS response rate for that session. (The pre-CS rate was taken as the average over the session to reduce trial by trial variability.) Cumulative records were then plotted across trials for the PRf CS and CRf CS. Two examples of these are shown in the lower plots of Figure 2. Because there were three times as many PRf trials as CRf trials in each conditioning session, we averaged the elevation scores across bins of three consecutive PRf conditioning trials before plotting the cumulative response records. On each plot, we added a regression line fitted to the cumulative record from the last two conditioning sessions (12 trials for the CRf CS, 12 3-trial bins for the PRf CS)¹ and projected this forward into the extinction sessions. This projected regression line shows the expected cumulative record if responding were to be maintained at same the level as at the end of training (i.e., if there were no extinction of responding). From this, we were able to identify the trial on the

¹ In a small number of cases, the cumulative record over these trials was not well fitted by a straight line, indicating some irregularities in response rates. In these cases, we included a larger number of trials to obtain a more reliable characterization of the terminal response rate.

cumulative record when the response rate dropped below that at the end of training. This judgement did not necessarily select the trial when the cumulative record first fell below the projected regression line because, in some cases, the response rate increased at the start of the extinction session (e.g., in Figure 2D). Rather, we selected the first trial when the slope of the cumulative record was shallower than the regression line, because it is the slope, rather than the height of the cumulative record, that indicates rate of responding.

This method makes no assumptions about the shape of the extinction curve and does not rely on arbitrary criteria to define the onset of extinction². A disadvantage of this method is that it calls on a human observer to make a subjective judgement from the cumulative response record. Therefore, to improve the reliability of the values obtained with this method, each cumulative record was rated independently by two observers (the two authors). To remove any opportunity for systematic bias in the judgements, these ratings were performed blind to the identity of the subject and CS of each record. There was good agreement between the two observers (Pearson $r = 0.82$). The ratings were within 5 trials for 44 response records. In these cases we took the average of the two ratings as the final value for number of trials to commence extinction. For the other 14 cumulative records, the two raters jointly re-inspected each record, still blind to the rat and CS that the record belonged to, and resolved on a single value. With this method, across all records the mean number of trials to extinction was 34.0 for the PRf CS and 10.2 for the CRf CS (plotted as N in Figure 1D). The ratio of these is 3.3:1. A Q-Q plot across 7 quantiles of the data (similar to that shown in Figure 1F) was well described by a regression line with slope of 3.1 ($R^2 = 0.94$), which is consistent with the 3.3:1 ratio of the means of the two CSs. In sum, this analysis has confirmed the results of the model-fitting analysis using the Weibull function in showing that the number of trials to

² As compared with methods that, for example, use the trial by trial record to identify the first trial on which responding falls below some predefined value and remains below that value for at least X of the next Y trials.

commence extinction was close to 3 times larger for responding to the PRf CS than for responding to the CRf CS.

Comparison of the number of trials taken to extinguish responding. The ratios plotted in Panel D of Figure 1 speak directly to the hypothesis that partial reinforcement slows extinction by an amount that is equal to the inverse of the reinforcement rate, which, in this case, is 3:1. The obtained ratio was close to this value if calculated using the earliest point on the extinction curve, when responding had fallen by just 10% of its starting value. However, as already noted, the ratios decreased below this value for later points on the extinction curve. The results suggest that the point at which extinction commenced did correspond to that predicted based on the reinforcement schedule, however this relationship was not maintained as extinction progressed.

We conducted a further analysis to compare how extinction for each CS progressed once it had commenced. Using the fitted Weibull functions, we measured the number of trials taken for extinction to progress from 20% to 80%. If the only difference between the two CSs was in the point at which extinction commenced, then there should be no differences between them when comparing how many trials it took for extinction to progress from 20% to 80%. However, the analysis showed that the two groups did differ. The partially reinforced CS took on average 74.0 trials to progress from 20% extinction to 80% extinction, whereas the continuously reinforced CS took just 45.8 trials. This difference was statistically significant; $t(28) = 4.921, p < .001$, Cohen's $d = .681$. Therefore, the progress of extinction was slower for the inconsistently reinforced CS throughout the course of extinction. Nonetheless, the ratio of these values is 1.6:1, which is well below the 3:1 ratio predicted by the difference in reinforcement schedules.

Experiments 2, 3 and 4

We found a PREE in Experiment 1, despite matched overall rates of reinforcement. This was evidence that extinction depends not on the overall rate of reinforcement, but on the per-trial probability of reinforcement. We also found that the latency to commence extinction matched the difference in the per-trial probabilities of reinforcement learned during acquisition. That is, a 3:1 difference in the per-trial probability of reinforcement during conditioning led to a similar difference in the number of trials taken to begin extinction. Experiments 2, 3, and 4 were conducted to replicate this result using a between-subjects design, and extend it by testing how extinction is affected by partial reinforcement schedules other than the 3:1 difference tested in Experiment 1. These experiments used between-subjects designs to avoid the possibility that generalisation between the Partial and Continuous CS might have contributed to our evidence for equivalent rates of responding across acquisition, or might have affected the values we obtained for rates of extinction.

Experiment 2 made the same comparison as Experiment 1, between a CS reinforced on 33% of trials and a CS reinforced on 100% of trials, except that the two schedules were given to different groups of rats. Experiment 3 compared extinction in a Group reinforced on 20% of trials with a group reinforced on 100% of trials. If the rate of extinction depends on the per-trial probability of reinforcement, then the ratio of the trials taken to begin extinction should be close to 3:1 for Experiment 2 as in Experiment 1. For Experiment 3, this ratio should be close to 5:1. Finally, instead of comparing a partially reinforced CS against a continuously reinforced CS, Experiment 4 compared two partially reinforced CSs, one reinforced on 33% of trials and one reinforced on 66% of trials. Assuming that the PREE is sensitive to the density of non-reinforcement, rather than just resulting from exposure to non-reinforcement per se, we predicted that we would still see a difference in the rate of extinction between the two groups in Experiment 4 (see also Rescorla, 1999). Moreover, this

difference should be comparable to the ratio of the per-trial probabilities of reinforcement during conditioning, in this case a 2:1 ratio.

Method

Subjects and apparatus. Each of the three experiments was conducted with 32 subjects. Experimentally naive female Sprague Dawley rats (*Rattus norvegicus*) were obtained and housed in the same manner as in Experiment 1.

The same 32 Med Associates conditioning chambers used in Experiment 1 were used to train and test the rats. The steady light and white noise stimuli, described for Experiment 1, were used. Reinforcement consisted of a single food pellet as in Experiment 1.

Procedure. One day prior to the beginning of each experiment, rats received a single magazine training session as described for Experiment 1. They then received daily conditioning sessions 5 days per week. Across all sessions, the timing of each photo-beam interruption by head entry into the magazine was recorded during each CS, and also during the 20-s pre-CS period.

After magazine training, all rats were divided into two groups. For Experiments 2 and 3, the two groups received either continuous or partial reinforcement of the noise CS, and continuous reinforcement of the light CS. For Experiment 4, both groups received partial reinforcement with the noise, but with different per-trial probabilities of reinforcement, and continuous reinforcement of the light. Thus, in all three experiments, the light served the same function as the clicker in Experiment 1, to provide a CS that could continue to receive reinforcement while responding to the target CS (the noise) was being extinguished.

Experiment 2. For Group Continuous, each of the first 20 sessions consisted of 12 presentations of the light, and 6 presentations of the noise. The duration of each CS presentation varied from trial to trial according to a uniform distribution. Both CSs were reinforced on every trial. The mean duration for the light was 17 s (2 to 32 s). The mean

duration of the noise was 30 s (2 to 58 s). Trials were intermixed randomly, with a mean inter-trial interval (ITI) of 200 s.

For Group Partial 33%, presentations of the light were identical to those for Group Continuous. The noise was presented 18 times, with a mean duration of 10 s, varying according to a uniform distribution (2 to 18 s). The noise was reinforced on one-third of trials (6 of 18). Trials were intermixed randomly, with a mean ITI of 120 s. The total session lengths for both Group Continuous and Group Partial were approximately 64 min.

Sessions 21 to 35 were identical for both groups. Presentations of the light stimulus were unchanged, reinforced on each trial with a mean duration of 17 s. The Noise was presented 10 times per session with a mean duration of 17 s (2 to 32 s). No reinforcers were provided for the noise. Light and noise trials were intermixed randomly, with a mean ITI of 170 s. The total session length for an extinction session was 65 minutes.

Experiment 3. The procedure for Experiment 3 was identical to that of Experiment 2, except for differences in the number of extinction sessions, average CS durations, and proportion of trials reinforced during acquisition. After magazine training, Group Continuous received 20 sessions of acquisition consisting of 12 presentations of the light (mean duration = 22 s, chosen as the geometric mean of 10 s and 50 s), and 6 presentations of the noise (mean duration = 50 s). All noise presentations were reinforced with a single food pellet. Group Partial 20% received identical presentations of the light over 20 sessions of acquisition, but for this group the noise was presented 30 times, with a mean duration of 10 s. Only 6 of these 30 noise presentations were reinforced. Trials of each type were randomly intermixed, with a mean ITI of 100 s. Total session length for acquisition was approximately 56 min for both groups.

Extinction lasted 20 sessions and was identical between the two groups. Each session consisted of 12 non-reinforced presentations of the noise, with a mean duration of 22 s. The

light continued to be reinforced with a mean duration of 22 s. Total session length for extinction was 59 min for both groups.

Experiment 4. After magazine training, rats were divided into two groups, Partial 33% and Partial 66%, before commencing 20 sessions of conditioning with the light and noise. For both groups, the light was presented 6 times, with a mean duration of 28 s (geometric mean of 20 s and 40 s), and was reinforced on every trial. For Group Partial 33%, the noise was presented 18 times (mean duration 20 s) and reinforced on 6 of these occasions. For Group Partial 66%, the noise was present 9 times (mean duration 40 s) and reinforced on 6 of these occasions.

After the conditioning phase, both groups underwent 15 sessions of extinction, consisting of 12 non-reinforced presentations of the noise, with a mean duration of 28 s. Presentations of the light remained unchanged.

Results and Discussion

Acquisition. As in Experiment 1, elevation scores were calculated for responding to the noise by subtracting the pre-CS rate from the response rate during the noise. Elevation scores during the acquisition phase for all three experiments are shown in panels A, B, and C of Figure 3. A 2 (Group) X 20 (Session) mixed-model ANOVA (with Greenhouse-Geisser correction) was used to analyse responding during acquisition for each Experiment. For Experiment 2 (Figure 3A), the analysis revealed a significant main effect of Session, $F(3.8, 113.3) = 15.178, p < .001, \eta^2_p = .336$, but no main effect of Group, $F(1, 28) = .093, p = .762$, nor interaction between Group and Session, $F(3.8, 113.3) = .867, p = .481$. For Experiment 3 (Figure 3B), there was a significant main effect of Session, $F(3.3, 99.2) = 13.625, p < .001, \eta^2_p = .312$, but no main effect of Group, $F(1, 30) = .136, p = .715$, nor interaction between Group and Session, $F(3.3, 99.2) = .619, p = .620$. The ANOVA conducted

for Experiment 4 (Figure 3C) revealed a significant main effect of Session, $F(3.2, 88.2) = 8.246, p < .001, \eta^2_p = .227$, but no main effect of Group, $F(1, 28) = .074, p = .788$, nor interaction between Group and Session, $F(3.2, 88.2) = .770, p = .520$. Therefore, in each experiment, responding to the noise was acquired in a very similar manner for both groups, despite the difference in the reinforcement schedules. This shows, for the first time using a between-subjects design, that the acquisition of responding is equivalent despite differences in per-trial probability of reinforcement, as long as the total reinforcement rate is equated per unit time.

Extinction. Elevation scores for responding during the extinction phase for Experiments 2, 3, and 4 are shown in panels D, E, and F of Figure 3. A 2 (Group) X 15 (Session) mixed-model ANOVA (with Greenhouse-Geisser correction) was used to analyse responding during extinction for Experiments 2 and 4. Analysis for Experiment 3 was conducted using a 2 (Group) X 20 (Session) mixed ANOVA. The analysis for Experiment 2 revealed a significant main effect of Session, $F(4.4, 132.4) = 23.170, p < .001, \eta^2_p = .436$. There was also a significant main effect of Group, $F(1, 30) = 8.049, p = .008, \eta^2_p = .212$, and a significant interaction between Group and Session, $F(4.4, 132.4) = 2.512, p = .039, \eta^2_p = .077$. For Experiment 3, there was a significant main effect of Session, $F(3.0, 89.8) = 16.579, p < .001, \eta^2_p = .356$, as well as a significant main effect of Group, $F(1, 30) = 7.257, p = .011, \eta^2_p = .195$. However, the interaction between Group and Session was not significant, $F(3.0, 89.8) = 1.958, p = .126$. For Experiment 4, there was a significant main effect of Session, $F(3.4, 95.0) = 23.743, p < .001, \eta^2_p = .459$, but no significant main effect of Group, $F(1, 28) = 3.243, p = .082$, nor was there a significant interaction between Group and Session, $F(3.4, 95.0) = 1.586, p = .192$. These analyses show that extinction of responding to the noise was significantly slower for Group Partial than Group Continuous in Experiments 2 and 3,

but that the difference in responding during extinction between Group Partial 33% and Group Partial 66% was not significant.

Model-fitting analysis of individual extinction curves. Following the procedure developed for Experiment 1, we obtained estimates of the rate of extinction for each rat by fitting an inverse cumulative Weibull distribution to the trial by trial data across the extinction sessions for both groups in all three experiments. The left panels of Figure 4 show the mean number of trials taken to reach each extinction criterion for each group in Experiments 2, 3, and 4. These results support the results of Experiment 1. In each experiment, the group that received a higher probability of reinforcement per trial during acquisition took fewer trials to reach the same extinction criteria compared to the group that received a lower probability of reinforcement per trial. The ratio of the number of trials to reach each extinction criterion between the two groups are shown in the same plots. In each experiment, the ratio was highest for the earliest point of extinction (10%) and decreased thereafter. Moreover, at the earliest extinction criterion, the ratios were close to the values predicted based on the ratios of the per-trial reinforcement rates: in Experiment 2, the 10% extinction ratio was 2.9:1, in Experiment 3 it was 4.6:1, and in Experiment 4 it was 2.0:1.

A 2 (Group) x 9 (Percentage of Extinction) mixed ANOVA was used to analyse whether the difference between trials taken to extinguish was significantly different for each group in each experiment. Analysis for Experiment 2 revealed there were significant main effects of Group ($F(1,30) = 98.368, p < .001, \eta^2_p = .766$), and percentage of extinction ($F(1, 30.3) = 50.640, p < .001, \eta^2_p = .628$), as well as a significant interaction between Group and extinction, $F(1,30.3) = 7.387, p = .011, \eta^2_p = .136$. For Experiment 3, there were significant main effects of Group ($F(1,30) = 55.283, p < .001, \eta^2_p = .648$), and percentage of extinction ($F(1, 30.2) = 36.9, p < .001, \eta^2_p = .552$), as well as a significant interaction between Group and extinction, $F(1,30.2) = 4.71, p = .038, \eta^2_p = .136$. Likewise, for Experiment 4, there were

significant main effects of Group, $F(1,28) = 158.389, p < .001, \eta^2_p = .850$, and percentage of extinction, $F(1, 28.7) = 120.194, p < .001, \eta^2_p = .811$, as well as a significant interaction between Group and extinction, $F(1,28.7) = 7.468, p < .001, \eta^2_p = .211$. In sum, in all three experiments, extinction was significantly slower (i.e., it took a significantly larger number of trials to reach criterion) for the group conditioned with the leaner partial-reinforcement schedule, consistent with a PREE.

The right panels of Figure 4 show Q-Q plots comparing the number of trials taken to reach 10% extinction for each Group, divided into 4 equal quantiles, in Experiments 2, 3, and 4. In each case, the data were well-fitted by a line with a zero intercept: In Experiment 2, the slope was 2.83 ($R^2 = 0.983$); in Experiment 3, the slope was 5.23 ($R^2 = 0.938$); and in Experiment 4, the slope was 2.0 ($R^2 = 0.991$). These slopes are similar to the ratios of the means reported above.

Non-Weibull analysis of individual extinction data. As we did for Experiment 1, we conducted an additional analysis using the cumulative record of elevation scores across trials to identify the number of trials to commence extinction for each group. In Experiment 2, the mean number of trials to extinction was 26.3 for the Partial (33%) group and 8.3 for the Continuous group (shown above point N in Figure 4A). The ratio of the values is 3.2:1, and is consistent with the slope (3.2) of the regression line fitted ($R^2 = 0.85$) to the Q-Q plot of the quartiles of these data (similar to that shown in Figure 4D). In Experiment 3, the mean number of trials to extinction was 32.9 for the Partial (20%) group and 6.7 for the Continuous group (Figure 4B). The ratio of these values (4.9:1) is similar to the slope (5.1) of the regression line fitted ($R^2 = 0.94$) to the Q-Q plot of the quartiles of each group. Finally, in Experiment 4, the mean number of trials to extinction was 26.6 for the Partial 33% group and 12.0 for the Partial 66% group (Figure 4C). The ratio of these values is 2.2:1, which is very close to the slope (2.3) of the regression fitted ($R^2 = 0.98$) to the Q-Q plot of the quartiles of

each group. In sum, these analyses have confirmed the results obtained from the Weibull fitting analyses. In each experiment, the number of trials to extinction differed between the two groups by a factor that was very close to the ratio of the per-trial reinforcement rates.

Comparison of the number of trials taken to extinguish responding. As we had observed in Experiment 1, the ratio of the number of trials taken to reach 10% extinction was close to that predicted by the reinforcement schedules, but these ratios steadily decreased below that predicted for later extinction criteria. Therefore, as we did for Experiment 1, we used the Weibull functions fitted to the extinction data to estimate the number of trials taken for extinction to progress from 20% to 80% and compared this between the two groups in each of Experiments 2, 3, and 4. In Experiment 2, the Partial group took 75.1 trials and the Continuous group took 33.0 (ratio = 2.3:1), and this difference was statistically significant, $t(30) = 2.863$, $p = .008$, $d = 1.762$. In Experiment 3, The Partial group took 113.7 trials and the Continuous group took 53.3 trials (ratio = 2.1:1), which again was a significant difference, $t(30) = 2.233$, $p = .033$, $d = .595$. Finally, in Experiment 4, the Partial 33% group took than 98.2 trials, compared with 57.4 trials for the Partial 66% group (ratio = 1.7:1), and this difference too was significant, $t(28) = 2.850$, $p = .008$, $d = .870$. These results show that, even after responding had started to extinguish, the progress of extinction remained slower for the group given the leaner reinforcement schedule in training. Nonetheless, they also show that the difference in the rate that extinction progressed did not match that predicted by the difference in reinforcement schedules.

General Discussion

Our results for Experiments 1, 2, and 3 showed that, overall, responding during extinction was significantly greater for the partially reinforced CS than for the continuously reinforced CS, consistent with a PREE. In Experiment 4, responding over 20 extinction

sessions was higher for a group trained with a 33% partial reinforcement schedule than for a group trained with a 66% partial schedule, although the difference between the groups was not statistically significant. We used two independent methods to quantify the rate of extinction. One method fitted an inverted Weibull function to each individual rat's response data across all extinction trials. From this Weibull function, we could extract estimates of the number of trials that each rat took for their responding to extinguish by a given amount (from 10% to 90%). In the second method, the cumulative response record of each rat was visually inspected to identify the onset of extinction. We found that, in all four experiments, responding to the CS that signalled a lower per-trial probability of reinforcement took significantly more trials to extinguish to any given criterion than did responding to the CS that signalled a higher probability of reinforcement per trial. That is, responding extinguished more quickly following continuous reinforcement than partial reinforcement in Experiments 1, 2, and 3, and responding extinguished more quickly after 66% reinforcement than after 33% reinforcement in Experiment 4.

In each of the present experiments, we compared responding to CSs that differed in per-trial probability of reinforcement but were matched for their rate of reinforcement per unit of time. As we have observed previously (Harris et al., 2015; Harris et al., 2017), these conditions produce virtually identical levels of responding across acquisition. Experiment 1 replicated our previous results using a within-subjects design, while Experiments 2, 3, and 4 showed this holds true for between-group comparisons as well. This finding in itself suggests that rats only learn about the total rate of reinforcement, and do not learn about the probability of reinforcement per trial. However, the fact that responding extinguished at different rates following the two reinforcement schedules demonstrates that rats do indeed learn something about the per-trial reinforcement rate (see also Harris & Kwok, 2018). This confirms other demonstrations of a PREE despite matching the overall reinforcement rates

(Bouton et al., 2014; Harris and Chan, 2017; Haselgrove et al., 2004), and directly contradicts the account of the PREE offered by timing models like RET (Gallistel & Gibbon, 2000; Gallistel, 2012). According to RET, the PREE only exists if the partial schedule gives a lower overall rate of reinforcement than the continuous schedule, making it harder for animals to detect a change to zero reinforcement under extinction. If, as argued by RET, the PREE does depend on differences in how quickly animals can identify the change to zero reinforcement, then the animals in the current experiments must have encoded the per-trial probability of reinforcement rather than the overall rate of reinforcement during conditioning.

The evidence we present here for a PREE is also relevant for Behavioural Momentum Theory (BMT; Nevin & Grace, 2000) which, like RET, predicts there should be no PREE when rates of responding during acquisition are equated. According to BMT, matching reinforcement rates should produce equal response rates during conditioning, as we observed, but should also result in equivalent resistance to extinction. Some evidence in support of that prediction, rather than the PREE we observed here, has been provided by Nevin and Grace (2005) who compared resistance to extinction for three schedules in a free-operant paradigm. They trained pigeons to respond to a multiple-component schedule consisting of either a variable interval (VI) 20-s schedule (providing 3 reinforcers/min), or one of two VI 60-s schedules (each providing 1 reinforcer/min). One of the VI 60-s components (3F) was presented three times as often as the VI 20-s component, while the other VI 60-s component (3D) was presented for three times as long as the VI 20-s component—ensuring that all three components provided an equal number of reinforcers overall. Nevin and Grace (2005) found that the VI 20-s schedule produced greatest resistance to extinction, with no difference in resistance to extinction for the other two components. Of most relevance to our study is the latter comparison between components 3D and 3F. These two schedules were matched for overall rate of reinforcement but differed in the number of trials reinforced: 3D was

reinforced on 50% of trials whereas 3F was reinforced on only 17% of trials. The difference in per-trial reinforcement rate should have led slower extinction of the 3F component than the 3D component, according to the results we present here. The fact that the rate of extinction was not significantly different appears to contradict our results.

There is a key difference between Nevin and Grace's (2005) experiment and our experiments in the way extinction was conducted. In our experiments, the presentation schedules of the two CSs were completely matched during extinction, they were presented the same number of times and for the same duration. Nevin and Grace, on the other hand, maintained the different presentation schedules that they had used in training for the two components. That is, the 3F component was presented 3 times as often as the 3D component, and presentations of the 3D component were 3 times as long as the 3F component. The fact that responding extinguished at the same rate per session under these conditions is consistent with a recent experiment from our lab (Chan & Harris, 2017). In that experiment, evidence for a PREE was eliminated when the PRf CS (reinforced on 33% of trial) was given 3 times as many extinction trials as the CRf CS (a similar finding was reported by Bouton et al., 2014). That result confirms a number of studies showing that the rate that responding extinguishes is related to the number of extinction trials rather than the duration of those trials (Drew, Yang, Ohyama & Balsam, 2004; Golkar, Bellander & Öhman, 2013; Harris & Andrew, 2017; Haselgrove & Pearce, 2003). In sum, the absence of evidence for a PREE in the session-by-session rates of extinction in Nevin and Grace's experiment can be attributed to the fact that the 3F component was presented 3 times more often than the 3D component.

The current set of experiments were primarily concerned with quantifying differences in the rates of extinction as a function of the reinforcement schedule. It was for this reason that we used a protocol that equated the rate of reinforcement per unit time during training, while varying the per-trial probability of reinforcement, so that we could compare extinction

functions that had the same starting point (i.e., were matched for rates of responding during acquisition). Following previous experiments by Bouton et al (2014) and Chan & Harris (2017), we suggested that rats keep track of the expected number of trials per reinforcer during acquisition and express this learning only in the face of extinction, such that responding to a partial CS would take X times as many trials to extinguish if it had been reinforced once every X number of trials. Our experiments bore out this suggestion, inasmuch as the point at which responding to each CS began to extinguish was directly related to the probability of reinforcement learned during the conditioning phase. For Experiments 1 and 2, responding to a CS reinforced on 1 in 3 trials took close to 3 times as many trials to reach 10% extinction compared with a CS reinforced on 100% of trials. In Experiment 3, responding to a partial CS reinforced on 1 in 5 trials took almost 5 times longer to extinguish by 10% than a continuously reinforced CS. In Experiment 4, responding to a CS reinforced on 1 in 3 trials took twice as long to extinguish by 10% as did responding to a CS reinforced on 2 in 3 trials. It is remarkable that the ratio of the probability of reinforcement during acquisition is so consistently related to the ratio of trials to extinction, across the 4 experiments. The finding is consistent with Sequential Theory (Capaldi, 1966, 1994) in showing that persistence of responding under extinction is directly related to the length of runs of non-reinforced trials prior to reinforcement.

Sequential Theory emphasises that the PREE depends on the animal remembering sequences of non-reinforced trials and associating those sequences with eventual reinforcement. An alternative view is that the PREE is not a special case of extinction that applies only to animals that have been conditioned to respond to non-reinforced trials during acquisition. Rather, during conditioning animals might learn specifically how often each presentation of the CS is reinforced, and use this information to keep track of each time that an expected reinforcer is missed during extinction. When the number of missed reinforcers

exceeds some threshold, then responding decreases (see Gallistel, 2012, and Gibbon et al., 1980, for a similar argument, but based on time to expected reinforcement rather than number of trials). The PREE occurs because it takes many more trials for the animal to accumulate evidence that expected reinforcers have been missed after partial reinforcement, but the extinction process itself is no different from that following continuous reinforcement. There are two caveats to this argument, which suggest that the PREE may be controlled by other factors as well. First, early experiments using a runway paradigm showed that that a PREE could occur when partial reinforcement was followed by continuous reinforcement prior to extinction (Jenkins, 1962). One might expect the intervening exposure to continuous reinforcement would increase how often the animal expected reinforcement and thereby eliminate the PREE. It remains to be determined whether, in such a case, the latency to onset of extinction would still be directly related to the proportion of non-reinforcement trials during conditioning, or whether the intervening exposure to continuous reinforcement would break that relationship. It may be useful to revisit Jenkins's procedure using a paradigm similar to the current set of experiments that match the overall rates of reinforcement in order to fully tease apart the different factors that influence extinction learning. The second caveat concerns the observation, discussed below, that the rate at which extinction progressed after onset was slower following partial reinforcement, but it was not as slow as would be predicted if extinction continued to depend on the accumulation of missed opportunities for reinforcement. Thus, different processes may control when extinction begins and how it progresses.

While the data reported here for the onset of extinction fit remarkably well with the hypothesis that the rate of extinction is directly proportional to the previous rate of reinforcement per trial, this pattern was not maintained so clearly across the remainder of the extinction process. In each experiment, the ratio of the number of trials to reach successive

extinction criteria steadily decreased below the predicted ratio. For example, in Experiment 1, responding to the partial (33%) CS took only 2.1 times longer to extinguish by 50% than responding to the continuous CS (compared with a ratio of 2.9 for responding to extinguish by 10%). Evidently, the per-trial probability of reinforcement did not correspond to the rate of extinction throughout the entirety of extinction. To examine this point in more detail, we conducted an additional analysis on the rate that extinction progressed once it had been initiated (measured as the number of trials taken for responding to decrease from 80% down to 20% of the pre-extinction level). This showed that responding to the CS reinforced on the thinner schedule continued to decrease more slowly throughout the extinction process, but the difference in rate of extinction no longer matched the difference in reinforcement schedules during conditioning. It might be that reinforcement rate only directly predicts the onset of extinction, and extinction of responding beyond that point depends on a different set of processes that do not track the prior ratio of non-reinforced to reinforced trials. However, this conclusion relies entirely on estimates of the rate of extinction derived from Weibull functions fitted to the extinction data, and therefore the specific results may reflect constraints imposed by the shape of the function itself rather than the pattern of the empirical data.

One issue that we have not yet addressed concerns the demonstration in Experiment 1 of a PREE using a within-subjects design in which rats were trained with both a continuously reinforced CS and a partially reinforced CS. This adds to a growing number of experiments to make within-subjects demonstrations of a PREE (Chan & Harris, 2017; Harris & Kwok, 2018; Rescorla, 1999), despite claims that the effect cannot be observed within-subjects (Amsel, 1966; Pearce, Redhead & Aydin, 1997). A common factor for the experiments that have shown a PREE within-subjects is the presence of an alternative source of reinforcement during extinction. For Rescorla's experiments, alternative reinforcement came in the form of a separate response key that scheduled a different rate of reinforcement throughout the

experiment; in our experiments, it was provided in the presence of a third CS (a clicker) that continued to signal delivery of food. This continued reinforcement would have prevented extinction of responding to the magazine itself, and thus allowed us to observe extinction of responding that was specific to the CSs of interest. This argument concerning the rate of extinction of the magazine may also explain how changes in context can differentially affect extinction after partial versus continuous reinforcement (Pearce et al., 1997). All one needs to assume is that extinction of the magazine is slower in a context associated with partial reinforcement than in a context associated with continuous reinforcement (i.e., a PREE of the magazine). For example, Pearce et al.'s evidence for a within-subjects PREE when the PRf and CRf CSs were conditioned in different contexts might have been the result of slower extinction of the magazine in the context where partial reinforcement had occurred. Bouton and Sunsay (2001) found very rapid extinction of responding to both a CRf CS and a PRf CS if they were tested in a context that had only been associated with continuous reinforcement, consistent with the argument that a magazine associated with continuous reinforcement extinguishes more quickly than one associated with partial reinforcement.

Demonstrations of a within-subjects PREE are important because they pose a challenge for Sequential Theory, as well as similar explanations of the PREE such as Amsel's (1958) Frustration Theory and the Hull-Sheffield hypothesis (Sheffield, 1949). According to these theories, when animals are conditioned on a partial reinforcement schedule, reinforcement is frequently delivered in the presence of the after-effects of non-reinforced trials. As a result, non-reinforced trials continue to elicit responding during extinction because the after-effects of non-reinforcement are part of what has been conditioned. But this type of account is hard pressed to explain evidence for the PREE obtained using a within-subjects design because any after-effects produced by non-reinforcement should impact on responding to both CSs equally. Therefore, these generalisation accounts of the PREE need to

be revised if they are to accommodate evidence for the PREE within-subjects. One would need to argue that the PREE is not based on learning about the after-effects of non-reinforcement per se during conditioning, but that it involves learning specifically about the non-reinforcement of a particular CS. That is, the animal's experience of non-reinforcement must include information about which CS was not reinforced. For example, the animal could learn that a sequence of non-reinforced trials of one CS, X, is followed by reinforcement of X, but learn nothing about non-reinforcement of a second CS, Y, that had been consistently reinforced throughout training. During subsequent extinction, as the animal experiences both CSs without reinforcement, only non-reinforcement of CS X is associated with reinforcement, and so X alone will benefit from greater generalisation of conditioning. We see this as the most cogent explanation of our findings.

References

- Amsel, A. (1958). The role of frustrative nonreward in noncontinuous reward situations. *Psychological Bulletin*, 55(2), 102–119. <http://doi.org/10.1037/h0043125>
- Amsel, A. (1962). Frustrative nonreward in partial reinforcement and discrimination learning: Some recent history and a theoretical extension. *Psychological Review*, 69(4), 306–328. <http://doi.org/10.1037/h0046200>
- Amsel, A., Rashotte, M. E., & Mackinnon, J. R. (1966). Partial reinforcement effects within subject and between subjects. *Psychological Monographs: General and Applied*, 80(20). Retrieved from <http://www.psycontent.com/index/JJ00H168631J7005.pdf>
- Balsam, P. D., Fairhurst, S., & Gallistel, C. R. (2006). Pavlovian contingencies and temporal information. *Journal of Experimental Psychology. Animal Behavior Processes*, 32(3), 284–294. <http://doi.org/10.1037/0097-7403.32.3.284>
- Balsam, P., Drew, M., & Gallistel, C. (2010). Time and associative learning. *Comparative Cognition & Behavior Reviews*, 5, 1–22. <http://doi.org/10.3819/ccbr.2010.50001>
- Bouton, M. E., & Sunsay, C. (2001). Contextual control of appetitive conditioning: influence of a contextual stimulus generated by a partial reinforcement procedure. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, 54(2), 109–125. <http://doi.org/10.1080/02724990042000083>
- Bouton, M. E., & Sunsay, C. (2003). Importance of trials versus accumulating time across trials in partially reinforced appetitive conditioning. *Journal of Experimental Psychology. Animal Behavior Processes*, 29(1), 62–77. <http://doi.org/10.1037/0097-7403.29.1.62>
- 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. <http://doi.org/10.1016/j.beproc.2013.08.006>
- Capaldi, E. J. (1966). Partial reinforcement: a hypothesis of sequential effects. *Psychological Review*, 73(5), 459–477. <http://doi.org/10.1037/h0023684>
- Capaldi, E. J. (1971). Memory and learning: a sequential viewpoint. In W. K. Honig & P. H. R. James (Eds.), *Animal memory* (pp. 111–154). New York: Academic Press. Retrieved from <http://www.annualreviews.org/doi/abs/10.1146/annurev.ps.41.020190.001125>
- 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.
- Capaldi, E. J., & Deutsch, E. A. (1967). Effects of severely limited acquisition training and pretraining on the partial reinforcement effect. *Psychonomic Science*, 9(3), 171–172. <http://doi.org/10.3758/BF03330814>

- Chan, C. K. J., & Harris, J. A. (2017). Extinction of Pavlovian conditioning: The influence of trial number and reinforcement history. *Behavioural Processes*, 141(Part 1), 19–25. <http://doi.org/10.1016/j.beproc.2017.04.017>
- Church, R. M., Meck, W. H., & Gibbon, J. (1994). Application of scalar timing theory to individual trials. *Journal of Experimental Psychology: Animal Behavior Processes*, 20(2), 135–155. <http://doi.org/10.1037/0097-7403.20.2.135>
- Crawford, L. L., Steirn, W. B., & Pavlik, W. B. (1985). Within- and between-subjects partial reinforcement effects with an autoshaped response using Japanese quail (*Coturnix coturnix japonica*). *Animal Learning & Behavior*, 13, 85–92. <http://dx.doi.org/10.3758/BF03213369>
- 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.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107(2), 289–344. <http://doi.org/10.1037/0033-295X.107.2.289>
- Gallistel, C. R. (2012). Extinction from a rationalist perspective. *Behavioural Processes*, 90(1), 66–80. <http://doi.org/10.1016/j.beproc.2012.02.008>
- Gallistel, C. R., Fairhurst, S., & Balsam, P. (2004). The learning curve: implications of a quantitative analysis. *Proceedings of the National Academy of Sciences of the United States of America*, 101(36), 13124–31. <http://doi.org/10.1073/pnas.0404965101>
- 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.
- Golkar, A., Bellander, M., & Öhman, A. (2013). Temporal properties of fear extinction - Does time matter? *Behavioral Neuroscience*, 127, 59–69.
- Harris, J. A. (2011). The acquisition of conditioned responding. *Journal of Experimental Psychology: Animal Behavior Processes*, 37(2), 151–164. <http://doi.org/10.1037/a0021883>
- 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(1), 23–35.
- Harris, J. A., Patterson, A. E., & Gharaei, S. (2015). Pavlovian conditioning and cumulative reinforcement rate. *Journal of Experimental Psychology: Animal Learning and Cognition*, 41(2), 137–151.
- Harris, J. A., & Andrew, B. J. (2017). Time, trials, and extinction. *Journal of Experimental Psychology: Animal Learning and Cognition*. <http://doi.org/dx.doi.org/10.1037/str0000014>

- Harris, J. A., & Carpenter, J. S. (2011). Response rate and reinforcement rate in Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 37(4), 375–384. <http://doi.org/10.1037/a0024554>
- 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(3), 240–250. <http://doi.org/10.1037/0097-7403.30.3.240>
- 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.
- Jenkins, H. M. (1962). Resistance to extinction when partial reinforcement is followed by regular reinforcement. *Journal of Experimental Psychology*, 64(5), 441–450. <http://doi.org/10.1037/h0048700>
- Jennings, D. J., Alonso, E., Mondragón, E., Franssen, M., & Bonardi, C. (2013). The effect of stimulus distribution form on the acquisition and rate of conditioned responding: implications for theory. *Journal of Experimental Psychology: Animal Behavior Processes*, 39(3), 233–48. <http://doi.org/10.1037/a0032151>
- Mackintosh, N. J. (1974). *The psychology of animal learning*. London, United Kingdom: Academic Press.
- Mellgren, R. L., & Dyck, D. G. (1972). Partial reinforcement effect, reverse partial reinforcement effect, and generalized partial reinforcement effect within subjects. *Journal of Experimental Psychology*, 92, 339–346.
- Nevin, J. a., & Grace, R. C. (2000). Behavioral momentum and the law of effect. *The Behavioral and Brain Sciences*, 23(1), 73–90. <http://doi.org/10.1017/S0140525X00002405>
- Nevin, J. A., & Grace, R. C. (2005). Resistance to extinction in the steady state and in transition. *Journal of Experimental Psychology: Animal Behavior Processes*, 31(2), 199–212. <http://doi.org/10.1037/0097-7403.31.2.199>
- Nevin, J. A. (1988). Behavioral momentum and the partial reinforcement effect. *Psychological Bulletin*, 103(1), 44–56. <http://doi.org/10.1037/0033-2909.103.1.44>
- Pearce, J. M., Redhead, E. S., & Aydin, a. (1997). Partial reinforcement in appetitive Pavlovian conditioning with rats. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, 50(4), 273–94. <http://doi.org/10.1080/713932660>
- Rescorla, R. A. (1999). Within-subject partial reinforcement extinction effect in autoshaping. *The Quarterly Journal of Experimental Psychology*, 52b(1), 75–87.

- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*(2), 225–237. <http://doi.org/10.3758/PBR.16.2.225>
- Sheffield, V. F. (1949). Extinction as a function of partial reinforcement and distribution of practice. *Journal of Experimental Psychology*, *39*, 511-526.
- Ward, R. D., Gallistel, C. R., & Balsam, P. D. (2013). It's the information! *Behavioural Processes*, *95*, 3–7. <http://doi.org/10.1016/j.beproc.2013.01.005>

Figures

Figure 1. Panel A. Mean session-by-session responding to PRf (dotted line) and CRf (solid line) during the 20 acquisition sessions for Experiment 1. Panel B. Mean session-by-session responding to PRf and CRf during the 15 extinction sessions for Experiment 1. Panel C. Terminal response rates for CRf and PRf responding during acquisition plotted against one another. Panel D. Number of trials taken to extinguish to the CRf (solid line) and the PRf (dashed line) to different points of extinction, and the ratio of the difference between the two (orange (light-grey) line). The points at N denote the mean extinction trial number on which the cumulative record deviated from the regression line used in the non-Weibull analysis for CRf (filled black circle), PRf (unfilled circle), and the ratio of the two (orange [grey] circle). Panel E. Number of trials taken to reach 10% of extinction of PRf and CRf responding plotted against one another. The solid line represents the regression line including the outlier (square), while the dotted line represents the regression line not including the outlier. Panel F. Q-Q plot comparing the number of trials taken to reach 10% extinction of PRf and CRf responding, ranked into quantiles. Error brackets indicate SEM of the difference between the two conditions.

Figure 2. The top two figures show best fitting Weibull curves (solid lines) applied to the trial-by-trial elevation scores (grey lines) for two rats in Experiment 2. The bottom two figures show the cumulative records of the elevation of two rats in the same experiment. The arrow indicates the point at which responding deviated from the regression line continuing from the acquisition phase.

Figure 3. The panels on the left side represent the mean session-by-session responding to PRf (dotted line) and CRf (solid line) during the 20 acquisition sessions for Experiment 2 (Panel A), Experiment 3 (Panel B), and Experiment 4 (Panel C). The right-hand panels represent the mean session-by-session responding to PRf and CRf during the extinction phase for Experiment 2 (Panel D), Experiment 3 (Panel E), and Experiment 4 (Panel F). Error brackets indicate SEM.

Figure 4. The 3 plots on the left side show the number of trials taken for responding to extinguish to different criteria (from 10% to 90% extinction) for the two groups in Experiment 2 (Panel A), Experiment 3 (Panel B), and Experiment 4 (Panel C). The plots also show the ratio of the number of trials to extinction for the two groups. The points at N denote the mean extinction trial number on which the cumulative record deviated from the regression line used in the non-Weibull analysis. On the right, Q-Q plots for Experiment 2 (Panel D), Experiment 3 (Panel E), and Experiment 4 (Panel F) compare the number of trials to reach 10% extinction between the two groups split into quartiles. Error brackets indicate SEM.