Pre-publication copy.

Published in: *The Journal of Experimental Psychology: Animal Learning & Cognition*, 45, 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, upon publication.

The Partial Reinforcement Extinction Effect Depends on Learning about Non-Reinforced Trials Rather than Reinforcement Rate.

Justin A. Harris¹, Dorothy W. S. Kwok¹, and Daniel A. Gottlieb²

1. University of Sydney, Australia
2. Sweet Briar College, USA

Address for correspondence:

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

Running head: The Partial Reinforcement Extinction Effect

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

Abstract

Conditioned responding extinguishes more slowly after partial (inconsistent) reinforcement than after consistent reinforcement. This Partial Reinforcement Extinction Effect (PREE) is usually attributed to learning about nonreinforcement during the partial schedule. An alternative explanation attributes it to any difference in the rate of reinforcement, arguing that animals can detect the change to nonreinforcement more quickly after a denser schedule than a leaner schedule. Experiments 1a and 1b compared extinction of magazine responding to a conditioned stimulus (CS) reinforced with one food pellet per trial and a CS reinforced with two pellets per trial. Despite the difference in reinforcement rate, there was no reliable difference in extinction. Both experiments did demonstrate the conventional PREE comparing a partial CS (50% reinforced) with a consistent CS. Experiments 2 and 3 tested whether the PREE depends specifically on learning about nonreinforced trials during partial reinforcement. Rats were trained with two CS configurations, A and AX. One was partially reinforced, the other consistently reinforced. When AX was partial and A consistent, responding to AX extinguished more slowly than to A. When AX was consistent and A was partial, there was no difference in their extinction. Therefore, pairing X with partial reinforcement allowed rats to show a PREE to AX that did not generalise to A. Pairing A with partial reinforcement meant that rats showed a PREE to A that generalised to AX. Thus, the PREE depends on learning about nonreinforced trials during partial reinforcement and is not due to any difference in per-trial probability of reinforcement.

Keywords: Extinction; Pavlovian conditioning; Reinforcement schedules; Sequential Theory; Rate Estimation Theory.

The Partial Reinforcement Extinction Effect (PREE) is one of the most widely recognised phenomena in associative learning. It refers to a change in the rate at which a conditioned response extinguishes depending on the reinforcement schedule delivered during conditioning, and is observed for both operant (instrumental) and Pavlovian conditioning. For operant conditioning, a response that has been partially (i.e., inconsistently) reinforced is more resistant to extinction than a response that has been continuously (i.e., consistently) reinforced (e.g., Mowrer & Jones, 1945; Sheffield, 1949; Weinstock, 1954). Similarly, in Pavlovian conditioning, responding to a conditioned stimulus (CS) will extinguish more slowly if the CS had been inconsistently paired with the unconditioned stimulus (US) than if the CS had been consistently paired with the US (e.g., Fitzgerald, 1963; Fitzgerald, Vardaris, & Teyler, 1966; Humphrey, 1939, 1940).

In free operant conditioning schedules, where the animal can perform the response at any time, evidence for slower extinction after partial reinforcement than after continuous reinforcement is confounded by higher rates of responding during conditioning under partial reinforcement (Ferster & Skinner, 1957). This difference in the acquisition of responding occurs because the animal must respond multiple times to obtain the reinforcer on a partial schedule whereas it only responds once to obtain the reinforcer on a continuous schedule. If, on the partial schedule, the animal learns to perform the whole reinforced cluster of responses as a single response unit, then both the higher rates of responding in acquisition and the slower extinction of responding can be easily understood in terms of differences in what constitutes a response (Mowrer & Jones, 1945; Skinner, 1938). The effects of partial reinforcement on acquisition are different for Pavlovian conditioning and discrete-trial operant conditioning schedules, such as running in an alley to obtain food, because the animal's rate of responding does not influence that rate at which it is reinforced. In these

cases, partial reinforcement usually results in weaker acquisition of responding than that established under a continuous schedule, but responding still extinguishes more slowly after partial reinforcement (Bacon, 1962; Fitzgerald, 1963; Humphrey, 1939, 1940; Sheffield, 1949). As discussed below, at least one major theory explains the PREE obtained with discrete trial and Pavlovian conditioning in terms of the idea that animals learn about a sequence of trials as a single reinforced unit, much like the response unit that is reinforced on a partial free operant schedule.

The most popular accounts of the PREE attribute the difference in rate of extinction to differences in generalisation from conditioning to extinction (Jenkins, 1950; Mackintosh, 1974). During partial reinforcement training, the after-effects of non-reinforcement become part of the cue configuration that is (eventually) reinforced on a subsequent trial, and as a result these cues serve to promote responding during extinction. In contrast, after consistent reinforcement the after-effects of non-reinforcement are experienced for the first time during extinction, and thus these novel cues decrease generalisation and reduce responding (Hull, 1941, cited in Jenkins, 1950; Sheffield, 1949). Amsel (1958, 1962) identified these after-effects with frustration evoked by the experience of non-reinforcement. Frustration evokes behaviours that normally compete with conditioned responding, and the strengthening of these behaviours across extinction leads to the progressive decline in responding. However, during conditioning on a partial reinforcement schedule, the animal is reinforced while experiencing frustration, and as a result frustration behaviours during extinction are a cue for responding. Capaldi (1966, 1994) identified the after-effects with a less specified memory of non-reinforcement. In effect, the partial reinforcement schedule delivers reinforcement in the presence of a memory of one or even multiple non-reinforced trials, and thus the animal learns that non-reinforced trials are part of a sequence that leads up to reinforcement. The

animal persists in responding during extinction because the experience of non-reinforcement is part of the pattern of events that has been conditioned. This notion of trial sequences bears some resemblance to Skinner's (1938) idea of a response unit, which comprises multiple individual responses, that is reinforced as a whole during conditioning on a ratio schedule, and leads to a greater number of single responses being performed under extinction (Mowrer & Jones, 1945).

An alternative theoretical account of the PREE has been proposed more recently by Gallistel (2012). This is part of a general theory that identifies conditioning with learning about the rate at which a CS or response is reinforced, where rate is defined in units of time rather than discrete trials, and identifies extinction with the detection of a change in that rate of reinforcement (Balsam & Gallistel, 2009; Gallistel & Gibbon, 2000). Since detection of change depends on the size of the change, animals will take longer to detect a change from sparse reinforcement to zero reinforcement (extinction) than a change from dense reinforcement to zero reinforcement. Therefore, if partial reinforcement means sparse reinforcement, it follows that extinction will be delayed and a PREE will result (see also Gibbon, Farrell, Locurto, Duncan, & Terrace, 1980). This account is challenged by demonstrations of a PREE when the rate of reinforcement is matched between consistent and partial schedules. For example, Chan and Harris (2017, 2019) trained rats with two CSs: a relatively long CS (lasting on average 30 s) was consistently reinforced, and a shorter CS (last on average 10 s) was partially reinforced (reinforced on one in 3 trials). Even though the overall rate of reinforcement was the same for the two CSs, at one US for every 30 s of the CS, responding extinguished much more slowly to the partially reinforced CS. Similar results have been reported by Haselgrove, Aydin and Pearce (2004) and by Bouton, Woods and Todd (2014). While these demonstrations are incompatible with a time-based account of the PREE, they are compatible with a trial-

based account. Indeed, if extinction depends on detection that the likelihood of reinforcement per trial has changed, then partial reinforcement will slow extinction by a specific amount that is inversely proportion to the likelihood itself. For example, if a partial CS is reinforced on 1 in 3 trials, responding to that CS will take 3 times longer to extinguish than responding to a consistently reinforced CS. This particular prediction has received support from several recent studies (Bouton et al., 2014; Chan & Harris, 2017, 2019).

In many respects, the trial-based account of the PREE just discussed seems indistinguishable from a CS-specific version of Capaldi's Sequential Theory. In the latter, if animals learn that a sequence of X non-reinforced trials leads up to a reinforced trial, then an extinction schedule only violates this sequence every X trials, meaning that responding should take X -times as many trials to extinguish. This is equivalent to the argument that, if animals learn that a CS is reinforced at a rate of one US every X trials, then under extinction they will only detect that an expected US has been omitted every X trials. While the two accounts converge on this point, they are, nonetheless, different. The sequential account identifies the PREE specifically with learning about the presence of non-reinforced trials during partial reinforcement. Therefore, two different reinforcement schedules should only lead to different rates of extinction if at least one schedule includes non-reinforced trials. The other account links the speed of extinction to the rate of reinforcement in general, and, as such, does not treat partial reinforcement as a special case. A difference in the rate of extinction should still be observed between reinforcement schedules that do not include non-reinforced trials but nonetheless differ in reinforcement rate—for example, if one schedule signals one US per trial and the other signals multiple USs on every trial. These differences between the accounts give rise to meaningful differences in predictions about extinction. The four experiments we present here were designed to test these predictions.

Experiments 1a and 1b

The first two experiments test the prediction that the speed of extinction is sensitive to differences in reinforcement rate even when there are no non-reinforced trials. According to the reinforcement rate account, animals learn how often to expect the US during conditioning, and extinction is sensitive to this previous reinforcement rate because extinction only progresses whenever the animal detects the omission of an expected US. Thus, an animal might detect that the US has been omitted on every extinction trial following consistent reinforcement but would only detect that a US has been omitted after multiple extinction trials following partial reinforcement. However, since rate is defined here in general terms, and is not specifically identified with non-reinforced trials, the relationship between extinction and reinforcement rate should also apply when the rate of reinforcement is greater than one US per trial. For example, if animals are exposed to a schedule in which they receive two USs per trial (per CS presentation), then they should detect that two USs have been omitted per trial during extinction, and therefore responding should extinguish twice as quickly compared to a CS that had been reinforced with a single US per trial. A similar prediction would be made by Frustration Theory (Amsel, 1958, 1962). If the animal has learned to expect two USs per trial, then the omission of all USs during extinction should produce more frustration, and therefore a faster decline in responding, than during extinction of a CS that signalled only one US per trial. However, this prediction is not made by Sequential Theory because the theory does not differentiate between memories of non-reinforced trials in terms of the number of USs that were expected.

Experiments 1a and 1b tests the prediction described above by comparing the rates of extinction of two CSs between two groups of rats. For Group 1-vs-2, CS A was conditioned

with two USs delivered per trial, and CS B was conditioned with one US per trial. For Group 1-vs-0.5, CS A was reinforced with one US on half its trials (i.e., partial reinforcement) and CS B was reinforced with one US on every trial. This second group was included to provide a within-experiment demonstration of the standard PREE effect. The design is summarised in Table 1. If the speed of extinction is set by the rate of reinforcement during conditioning, then the speed of extinction will differ between the two CSs in both groups, with extinction being faster for the CS trained on a higher per-trial rate of reinforcement. The sequential account, on the other hand, predicts there will be a standard PREE in the second group, trained with one CS on a partial schedule, but it does not predict any difference in the first group because, for that group, neither CS has been conditioned with a sequence that includes non-reinforced trials.

----- Table 1 about here -----

The experiments adopted two procedures that we have used previously to compare speed of extinction between CSs (Chan & Harris, 2017, 2019; Harris & Kwok, 2018). First, the length of presentations of the CS, and thus the timing of the US (which coincided with CS offset), varied randomly from trial to trial around a specified mean duration. This ensures that response rates do not change systematically across the length of the CS (Harris & Carpenter, 2011; Harris, Gharaei, & Pincham, 2011) which would otherwise confound our attempts to compare response rates between CSs that have different durations. Second, the two CSs were matched for rate of reinforcement per unit time, with a US being delivered after an average of 20 s of accumulated CS time. The procedure should ensure that response rates to the two CSs are equated across the conditioning phase (Harris, Patterson, & Gharaei, 2015), which makes it easier to compare the subsequent extinction of responding between the two CSs. However, this goal to equate responding to CS A and CS B was thwarted for Group 1-vs-2 in

Experiment 1a, even though the two CSs were matched for their rate of reinforcement per unit time. One possible reason for the difference in responding was that we gave the rats in Group 1-vs-2 twice as many trials per session with B than with A (12 versus 6) in order to match the two CSs on total number of reinforcements. The difference in number of trials may have contributed to the difference in acquisition of responding. Therefore, Experiment 1b was conducted as a near replication of Experiment 1a except that we matched CS A and CS B for the number of trials each received during the conditioning phase across Sessions 1 to 20, rather than matching the CSs for total number of USs (see Table 1). For both Experiments, we predicted faster extinction of B than A in Group 1-vs-0.5, and faster extinction of A than B in Group 1-vs-2, if extinction is sensitive to the per-trial rate of reinforcement during conditioning. However, if extinction is only sensitive to the presence of non-reinforced trials during conditioning, then we should only see a difference in the speed of extinction between A and B in Group 1-vs-0.5, and not in Group 1-vs-2.

Another notable feature of these experiments, and the subsequent experiments, is the use of a within-subject design—rats were concurrently trained with both a partially reinforced CS and a continuously reinforced CS, and then both CSs were concurrently extinguished. While we have reliably observed a PREE under these conditions (Chan & Harris, 2017, 2019; Harris & Kwok, 2018), such findings pose a serious challenge for many of the theoretical accounts of the PREE described earlier. According to these accounts, the intermixed schedule should allow the after-effects of non-reinforcement to be conditioned, and so promote generalisation of responding to both CSs equally under subsequent extinction. The significance of this will be discussed in more detail later in the General Discussion.

Methods

Subjects

In each experiment, 32 experimentally-naive female albino Sprague Dawley rats (8 to 10 weeks of age at the start of the experiment) were obtained from Animal Resources Centre, Perth, Western Australia. The rats were housed in groups of four in split-level ventilated plastic tubs (Techniplast™), measuring 40 x 46 x 40cm (length x width x height), located in the School of Psychology animal colony. They had unrestricted access to water in the home boxes. One week after transportation, and 3 days prior to commencement of the experiment, they were placed on a restricted food schedule. Each day, half an hour after the end of the daily training session, each tub of rats received a ration of their regular dry chow (3.4 kcal/g) equal to 5% of the total weight of all rats in the tub. This amount is approximately equal to their required daily energy intake (Rogers, 1979), and took at least 2 h to be eaten (but was usually finished within 3 h).

Apparatus

Rats were trained and tested in 32 Med Associates™ conditioning chambers. Twenty-four chambers measured 28.5 x 30 x 25 cm (height x length x depth); the other 8 measured 21 x 30.5 x 24 cm. The rats from both groups were assigned in equal numbers to the different-size chambers. The end walls of each chamber were aluminum; the sidewalls and ceiling were Plexiglas™. The floor of the chamber consisted of stainless-steel rods, 0.5 cm in diameter, spaced 1.5 cm apart. Each chamber had a recessed food magazine in the center of one end wall, with an infra-red LED and sensor located just inside the magazine to record entries by the rat. A small metal cup measuring 3.5 cm in diameter and 0.5 cm deep was fixed on the floor of each food magazine. The cup was in the center of the magazine in 24 of the chambers,

and to the left of center in the other 8 chambers. Attached to the food magazine was a dispenser delivering 45 mg food pellets (purified rodent pellets; Bioserve, Frenchtown, NJ) as the US. Each chamber was enclosed in a sound- and light-resistant wooden shell. Throughout all sessions, fans located in the rear wall provided ventilation; the operation of these created a background noise level measuring between 61 and 66 dB (depending on the chamber). Experimental events were controlled and recorded automatically by computers and relays located in the same room. Three stimuli served as the three CSs, A, B, and C. White noise (19 dB above background) was presented from a speaker mounted on the wall of each conditioning chamber above and to the left of the food magazine. A steady light was produced either by an incandescent bulb (30cd/m²) located on the back wall of the wooden shell (24 chambers), or by an LED strip along the ceiling of the wooden outer shell (8 chambers). A clicker stimulus (4 clicks per second, each click 24 dB above background) was emitted from a module (Med Associates, product ENV 135M) located on the wall of the conditioning chamber above and to the right of the magazine. The white noise and light were allocated in counterbalanced fashion as CSs A and B; the clicker fulfilled the role of CS C for all rats.

Procedure

All procedures were approved by the Animal Ethics Committee at the University of Sydney (Project number 2013/6026). Prior to the start of conditioning, rats received a 60-min magazine training session during which 20 food pellets were presented on a variable time (VT) 3-min schedule, with no stimulus presentations. Rats that ate fewer than half of the pellets were given a second session of magazine training the following day. After all rats completed magazine training, they commenced daily conditioning sessions 5 days per week for 20 sessions. Just prior to this, the rats were split into two equal groups ($n=16$ per group).

Experiment 1a. For Group 1-vs-2, there were 30 conditioning trials per session, which comprised 6 trials with CS A and 12 trials with each of CS B and CS C. The 6 presentations of CS A had a mean duration of 40 s (range 4 to 76 s), the 12 presentations of B had a mean duration of 20 s (2-38 s), and the 12 presentations of C had a mean duration of 10 s (2-18 s). Each presentation of A was reinforced with delivery of two food pellets, one at a randomly determined time during the CS and the second to coincide with offset of the CS. Each presentation of CSs B and C was reinforced with delivery of a single food pellet timed to coincide with CS offset. For Group 1-vs-0.5, each conditioning session contained 48 trials, which comprised 24 trials of A (mean duration 10 s, range 2-18 s), 12 trials of B (mean 20 s, 2-38 s), and 12 trials of C (mean 10 s, 2-18 s). Twelve of the 24 trials of A, and all 12 of the trials of B and C, were reinforced with delivery of a single food pellet at CS offset. Therefore, for both groups, CS A and CS B were reinforced with food after on average 20 s of cumulative exposure to the CS. For both groups, CS durations varied randomly from trial to trial, sampled from a uniform distribution of durations. In Group 1-vs-0.5, the 12 reinforced trials had a mean duration of 10 s, as did the 12 non-reinforced trials. Trials of each type were randomly intermixed through each session. The inter-trial interval varied randomly with a mean of 170 s (range 40 s to 300 s) for Group 1-vs-2 and a mean of 110 s (40 s to 180) for Group 1-vs-0.5. For both groups, each session lasted approximately 95 min.

Responding to A and B was extinguished across Sessions 21 to 30. The sessions were identical for both groups. There were 40 trials per session, comprising 8 non-reinforced presentations of each of CSA and CS B (mean duration 14 s, range 5-23 s), and 24 reinforced trials of CS C (mean 10 s, 2-18 s). CS C was included in the experiment so that the rats could continue to receive food pellets during the extinction sessions for CSs A and B, thereby ensuring that any decrease in responding during A and B was due to extinction of those

specific CSs, and was not due to extinction of the magazine itself which might otherwise obscure differences in the extinction of CSs A and B. The 14-s durations were adopted for CSA and CSB because this is the geometric mean of the 10-s and 20-s durations used during conditioning in Group 1-vs-0.5. However, this duration is obviously closer to the 20-s conditioning duration of CSB than the 40-s conditioning duration of CSA in Group 1-vs-2. This could, in principle, mean that responding to CSA suffers a greater generalisation decrement than does responding to CS B during extinction, which would inflate evidence for the predicted result that responding to CSA will extinguish faster than to CSB. However, to anticipate the results, we did not find reliable evidence for this effect in either this experiment or the subsequent one.

For both groups, the inter-trial interval varied randomly with a mean of 110 s (40 s to 180), each session lasted approximately 80 min. Throughout the experiment, both the start and end of all photo-beam interruptions by entry into the magazine were recorded during each CS and each 20-s pre-CS period. Only the initiation of each beam break was used when analysing the number of responses made per second.

Experiment 1b. The procedure of Experiment 1b replicated that of Experiment 1a in every detail except as follows. During Sessions 1 to 20, rats in Group 1-vs-2 received 8 trials of each of the three CS types (A, B, and C), and rats in Group 1-vs-0.5 received 12 trials of each of CS A and CS B and 6 trials of CS C. Therefore, each session lasted approximately 75 min for Group 1-vs-2 and 90 min for Group 1-vs-0.5. Sessions 21 to 30 were identical to those in Experiment 1a.

Results

Experiment 1a

One rat's data were lost due to equipment failure on Days 27 and 28. Because these data were critical to the comparison of response extinction between CS A and CS B, we omitted all data from this rat (in Group 1-vs-0.5) for the entire experiment. The top half of Figure 1 shows the response rates for both groups during presentations of CS A, CS B, and during the pre-CS period, for each of the 30 sessions in Experiment 1a. Both groups acquired much higher levels of responding during the CSs than during the pre-CS period across the conditioning phase (Sessions 1-20), and these responses gradually decreased across the extinction phase (Sessions 21-30). As expected, response rates to A and B across Sessions 1-20 were very similar for rats in Group 1-vs-0.5, but they did differ unexpectedly for rats in Group 1-vs-2. Response rates to A and B were also different across the extinction sessions for both groups.

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

All data were analysed using repeated measures ANOVAs run with JASP (JASP Team, 2018), with Greenhouse-Geisser correction to the degrees of freedom when assumptions of sphericity were violated. A 2 x (20) mixed model ANOVA comparing the two groups on their pre-CS response rates across the acquisition phase (Sessions 1-20) showed that there was no significant difference overall between the groups, $F(1,29) = 0.30$, $p = .590$. There was a significant main effect of Session, $F(5,149) = 2.39$, $p = .041$, $\eta^2_p = 0.07$, that did not interact with Group, $F(5,149) = 0.82$, $p = .535$. A 2 x (10) ANOVA on pre-CS response rates across extinction (Sessions 21-30) also found no significant difference overall between the two

groups, $F(1,29) = 0.30$, $p = .591$. There was no main effect of Sessions, $F(1.8,53) = 0.76$, $p = .463$, nor interaction between Sessions and Group, $F(1.8,53) = 0.99$, $p = .372$.

The acquisition of responding to CSs A and B across Sessions 1 to 20 were analysed by a $2 \times (2 \times 20)$ ANOVA. There was no overall difference in responding between the two groups, $F(1,29) = 0.93$, $p = .340$. There was no overall difference in responding between the two CSs $F(1,29) = 1.55$, $p = .222$, nor was there an interaction between CS and Group, $F(1,29) = 0.77$, $p = .387$. There was a main effect of Session, $F(4.1,122) = 25.00$, $p < .001$, $\eta^2_p = 0.45$, confirming that response rates during the CSs increased over the course of conditioning. This effect of Session did not interact significantly with either Group, $F(4.1,122) = 1.05$, $p = .384$, or CS, $F(8.8,263) = 1.78$, $p = .074$. However, there was a 3-way interaction, $F(8.8,263) = 3.47$, $p < .001$, $\eta^2_p = 0.10$.

To unpack the interaction, we conducted separate 2×20 repeated measures ANOVAs for each group. For Group 1-vs-2, there was not a significant main effect of CS, $F(1,15) = 3.49$, $p = .082$, but there was a significant interaction between CS and Session, $F(19,285) = 4.16$, $p < .001$, $\eta^2_p = 0.22$. The significant interaction reflects the fact that response rates did not differ between A and B across the first 10 sessions, but did differ over the last 10 sessions. A *post-hoc* comparison of response rates, averaged over Sessions 11 to 20, confirmed the difference between responding to A and responding to B, $t(15) = 4.39$, $p < .001$, $d = 1.10$. For Group 1-vs-0.5, there was no main effect of CS, $F(1,14) = 0.05$, $p = .826$, nor interaction between CS and Session, $F(19,266) = 1.46$, $p = .101$. To test the strength of evidence that responding to the two CSs did not differ in Group 1 vs 0.5, we conducted a Bayesian repeated measures ANOVA using JASP. For the overall difference between the CSs, the Bayes Factor in favour of the null hypothesis (BF_{01}) was 10.6, meaning that the observed data are 10.6 times more likely

to have occurred if the response rates to A and B were the same than if they were different. There was also strong evidence that the comparison between A and B did not interact with Session ($BF_{01} = 567$). Therefore, we had strong evidence that the rats in this group acquired responding at the same rate to the two CSs.

The extinction of responding to CSs A and B across Sessions 1 to 20 were analysed by a $2 \times (2 \times 20)$ ANOVA. There was a significant overall difference between the two groups, $F(1,29) = 8.89$, $p = .006$, $\eta^2_p = 0.24$. The overall difference between the two CSs was not significant, $F(1,29) = 3.09$, $p = .089$, but it did interact significantly with Group, $F(1,29) = 25.71$, $p < .001$, $\eta^2_p = 0.47$. There was a significant main effect of Session, $F(2.4,71) = 48.43$, $p < .001$, $\eta^2_p = 0.63$, that interacted with both Group, $F(9,71) = 5.95$, $p = .002$, $\eta^2_p = 0.17$, and CS, $F(9,140) = 4.80$, $p < .001$, $\eta^2_p = 0.14$. There was a significant 3-way interaction, $F(9,140) = 5.04$, $p < .001$, $\eta^2_p = 0.15$.

To unpack the interactions, we ran separate 2×10 repeated measures ANOVAs on the extinction data of each group. For Group 1-vs-2, there was a significant main effect of Session, $F(2.3,34) = 16.95$, $p < .001$, $\eta^2_p = 0.53$, confirming that response rates declined over the extinction phase. There was also a significant difference between overall responding to A and overall responding to B, $F(1,15) = 12.25$, $p = .004$, $\eta^2_p = 0.43$, which interacted significantly with Session, $F(4.1,61) = 8.53$, $p < .001$, $\eta^2_p = 0.36$. The overall difference between A and B, and interaction with Sessions, is consistent with the prediction that responding to A would extinguish more quickly than responding to B. However, interpretation of the difference in extinction between A and B is confounded by the fact that responding to A was already significantly lower than responding to B at the end of the conditioning phase. One approach to solving this problem is an analysis that compares the difference between A and B across

extinction against the difference between these CSs at the end of conditioning. We conducted a Bayesian 2 x 2 repeated measures ANOVA in which Factor 1 compared the average response rate across Sessions 11-20 with the average response rate across Sessions 21-30, and Factor 2 compared response rates to CS A with response rates to CS B. This showed there was moderate evidence that the difference between A and B was the same across the two blocks of sessions ($BF_{01} = 3.0$, favouring the null for the interaction between CS and Session). In other words, extinction did not increase the existing difference in responding to A versus B. However, the analysis just describe introduces its own problems because it assumes a uniform (linear) response scale. Therefore, a better approach is to use procedures that properly match response rates between A and B at the end of conditioning.

The analysis of the extinction data from Group 1-vs-0.5 revealed there was a significant main effect of session, $F(2.4,33) = 33.69$, $p < .001$, $\eta^2_p = 0.71$, confirming decrease in responding to the two CSs. There was a significant difference in responding to CS A and CS B overall, $F(1,14) = 14.74$, $p = .002$, $\eta^2_p = 0.51$, but the interaction between the CSs and Session was not significant, $F(3.8,54) = 2.39$, $p = .065$.

Experiment 1b

The data from one rat in Group 1-vs-0.5 of Experiment 1b was excluded from all analyses because its response rates during each of the CSs did not differ from that during the pre-CS period across the entire experiment. Response rates to CS A and CS B, as well as during the pre-CS period, are shown for both groups in the bottom half of Figure 1. Responding to A and B reached very similar levels across conditioning (Sessions 1-20) for Group 1-vs-2 but appeared to differ for Group 1-vs-0.5. Response rates to A and B differed markedly across

extinction (Sessions 21-30) for Group 1-vs-0.5 but were very similar for Group 1-vs-2. These data from each group were subject to repeated measures ANOVAs, as for Experiment 1.

A 2 x (20) mixed model ANOVA comparing the two groups on their pre-CS response rates across Sessions 1-20 showed that there was no significant difference overall between the groups, $F(1,29) < 0.01$, $p = .981$. There was a significant main effect of Session, $F(7.6,220) = 10.90$, $p < .001$, $\eta^2_p = 0.27$, that did not interact with Group, $F(7.6,220) = 1.68$, $p = .108$. A 2 x (10) ANOVA on pre-CS response rates across Sessions 21-30 also found no significant difference overall between the two groups, $F(1,29) = 0.12$, $p = .731$. There was a significant main effect of Sessions, $F(4.4,261) = 2.47$, $p = .043$, but no interaction between Sessions and Group, $F(4.4,261) = 1.16$, $p = .330$.

Responding to CSs A and B across Sessions 1-20 were analysed by a 2 x (2 x 20) ANOVA. There was no overall difference between the groups, $F(1,29) = 0.08$, $p = .785$. There was no overall difference in responding between the two CSs $F(1,29) = 2.54$, $p = .122$, nor was there an interaction between CS and Group, $F(1,29) = 0.21$, $p = .653$. There was a main effect of Session, $F(4.2,122) = 14.82$, $p < .001$, $\eta^2_p = 0.34$, confirming that response rates during the CSs increased over the course of conditioning. This effect of Session did not interact significantly with Group, $F(4.2,122) = 0.45$, $p = .781$, or CS, $F(8.5,247) = 0.51$, $p = .856$, but there was a 3-way interaction, $F(8.5,247) = 3.43$, $p < .001$, $\eta^2_p = 0.11$.

The interaction was followed up by separate 2x20 repeated measures ANOVAs on the acquisition data of each group. The ANOVA performed on response rates during conditioning for Group 1-vs-2 showed there was a significant main effect of Session, $F(3.4,50) = 8.17$, $p < .001$, $\eta^2_p = 0.35$, but no significant difference between responding to A vs B overall, $F(1,15) = 0.91$, $p = .354$. There was, however, a significant interaction between CS and Session,

$F(6.7,100) = 2.34, p = .032, \eta^2_p = 0.14$. This interaction reflects the fact that responding to A is higher than B over the first 10 sessions, but there is little difference in responding between the CSs across the last 10 sessions. A post-hoc comparison of response rates to A and B, averaged over Sessions 11-20, showed that there was no significant difference, $t(15) = 0.37, p = .718$. A repeated measures Bayesian ANOVA tested the strength of evidence for no difference between responding to A and B over the last 10 sessions. This showed there was moderate evidence that responding to the two CSs did not differ ($BF_{01} = 6.4$, meaning that the observed data were 6.4 times more likely to occur if the CSs produced the same response rates than if they produced different levels of responding). There was strong evidence that responding to the CSs did not interact with Session ($BF_{01} = 102$).

A 2x20 ANOVA of response rates in Group 1-vs-0.5 across Sessions 1-20 showed a significant main effect of Session, $F(19,266) = 7.14, p < .001, \eta^2_p = 0.34$, but no significant difference between responding to A and B overall, $F(1,14) = 1.58, p = .229$. There was a marginal interaction between these factors, $F(19,266) = 1.62, p = .050$. this marginal interaction reflects what appears to be a difference between responding to A versus B over the last half of the conditioning phase. To follow this up, and in keeping with the analysis for Group 1-vs-2, we conducted a repeated measures Bayesian ANOVA on the data from Sessions 11-20. This revealed strong evidence for a difference between responding to A versus B overall ($BF_{10} = 2890$), but no interaction between CS and Session ($BF_{01} = 95$).

A 2 x (2 x 10) mixed-model ANOVA was conducted on the extinction data from both groups. There was not a significant overall difference between the two groups, $F(1,29) = 2.81, p = .105$. There was a significant difference overall between the two CSs, $F(1,29) = 11.78, p = .002, \eta^2_p = 0.29$, that interacted significantly with Group, $F(1,29) = 8.23, p = .008, \eta^2_p = 0.22$.

There was a significant main effect of Session, $F(3.0,86) = 54.30, p < .001, \eta^2_p = 0.65$, that did not interact significantly with Group, $F(3.0,86) = 2.30, p = .084$, or with CS, $F(3.0,87) = 1.71, p = .171$. There was not a significant 3-way interaction, $F(3.0,87) = 1.49, p = .222$.

Separate 2 x 10 repeated measures ANOVAs were run on the extinction data for each group. For Group 1-vs-2, there was a significant main effect of Session, $F(2.6,38) = 27.87, p < .001, \eta^2_p = 0.65$, but no difference in responding to A and B overall, $F(1,15) = 0.20, p = .665$, and no interaction between CS and Session, $F(2.7,40) = 0.28, p = .816$. A Bayesian ANOVA was conducted to test the strength of evidence for no difference in responding to A versus B in Group 1-vs-2. This revealed moderate evidence that overall levels of responding to A and B were the same ($BF_{01} = 6.1$), and strong evidence that there was no interaction with Session ($BF_{01} = 106$). Thus, there was moderate evidence that responding to A and B did not differ overall, and strong evidence that any difference did not interact with session.

The 2x10 ANOVA on response rates across the extinction phase for Group 1-vs-0.5 revealed there was a significant main effect of Session, $F(2.7,37) = 28.57, p < .001, \eta^2_p = 0.67$, and a significant difference in responding to A and B overall, $F(1,14) = 16.50, p = .001, \eta^2_p = 0.54$, but not a significant interaction between these factors, $F(2.3,32) = 2.71, p = .076$. Even though the difference in responding between A and B during the conditioning phase was not significant, there was a difference in mean response rates in the last half of that phase that was in the same direction as the difference observed across the extinction phase. Therefore, as we did for Group 1-vs-2 in Experiment 1, we performed a 2x2 ANOVA on the response rates averaged over Sessions 11-20 and over Sessions 21-30, in order to test whether the difference between responding to A and to B during the extinction phase was greater than the difference in the last half of the conditioning phase. This analysis showed there was a significant

difference overall between the last half of the conditioning phase and the extinction phase, $F(1,14) = 62.34, p < .001, \eta^2_p = 0.82$, and between overall responding to A versus B, $F(1,14) = 9.45, p = .008, \eta^2_p = 0.40$, and there was also a significant interaction between these factors, $F(1,14) = 6.80, p = .021, \eta^2_p = 0.33$. This interaction establishes that the difference between responding to A and to B during extinction was significantly greater than any difference during conditioning. But once again, caution must be heeded when interpreting this result because the analysis just describes makes a strong, and questionable, assumption that the response scale is linear (such that a difference at one part of the scale can be meaningfully compared with a difference at another part of the scale).

Discussion

In both experiments, responding took longer to extinguish following partial (50%) reinforcement than following consistent (100%) reinforcement among rats in Group 1-vs-0.5. In Experiment 1b, this difference in extinction of responding to the two CSs is confounded by an unexpected difference between the CSs at the end of acquisition. The difference during acquisition was probably due to the fact that the partially-reinforced CS received half as many reinforced trials as did the consistently reinforced CS (but see Harris et al., 2015). Nonetheless, the difference in responding between the CSs was significantly larger across the extinction phase than at the end of acquisition, suggesting that greater responding to the partially-reinforced CS across extinction was not due entirely to an existing difference during acquisition. Moreover, the difference between the CSs during extinction was also observed in Experiment 1a, even though there was no difference between them in response rates during acquisition. These findings confirm that partial reinforcement produces responding that is more resistant to extinction than does consistent reinforcement.

The results concerning extinction of responding were less clear for Group 1-vs-2. We had predicted that responding to CS A, which had been reinforced with two pellets per trial, would extinguish faster than responding to B, which had been reinforced with one pellet per trial. Consistent with this, across extinction in Experiment 1a, responding to A was significantly lower than to B. However, the reason for this difference is confounded by the fact that responding to the two CSs was already different by the end of the conditioning sessions. Thus, it is unclear whether the difference between A and B during the extinction phase reflects a genuine difference in the rate of extinction between the two CSs, or is simply carried over from the difference during conditioning. Experiment 1b addressed this problem by using a conditioning schedule that did match the acquisition of responding to A and B in Group 1-vs-2. This experiment failed to provide any evidence for accelerated extinction of responding after a conditioning schedule that delivered two reinforcers per trial. Indeed, the data from Experiment 1b provided considerably stronger support for the hypothesis that there would be no difference in extinction for CSs conditioned with one versus two reinforcements per trial than for the hypothesis that responding would extinguish more quickly after conditioning with two reinforcements per trial.

The combined results of Experiments 1a and 1b support the following conclusions. Responding extinguished more slowly to a CS that had been reinforced on 50% of trials than a CS that had been reinforced on 100% of trials, regardless of whether response rates to the CSs were well matched or differed at the end of conditioning. Responding to a CS that had been reinforced with 2 pellets per trial extinguished at the same rate as responding to a CS reinforced with one pellet per trial, as long as response levels to the CSs were matched at the end of acquisition. (When responding to the CSs differed at the end of acquisition, that difference persisted across extinction.) Therefore, these experiments show that a

reinforcement schedule less than 100% can slow extinction, but a reinforcement schedule greater than 100% does not appear to accelerate extinction. This is consistent with the sequential account discussed in the Introduction, inasmuch as they show that an effect of reinforcement rate on extinction depends on the presence of non-reinforced trials (when the reinforcement schedule is less than 100%). The findings are inconsistent with the alternative account that links the speed of extinction with the trial-based rate of reinforcement in more general terms, since this would predict accelerated extinction when the conditioning schedule delivers more than one US per trial.

These conclusions depend on specific assumptions about what defines a trial. Here, we have treated each presentation of the CS as a trial. Presentation of the CS without the US counts as a non-reinforced trial; presentation of the CS that includes two USs counts as a double-reinforced trial. But the second of these assumptions would not hold up if the occurrence of the US also defines a trial, for example if the animal parses its experience such that any change in stimulus conditions delineates an episode. In that case, a long CS presentation that includes two separate USs would be treated by the rat as two independent trials. Thus, our manipulation of reinforcement rate in Group 1-vs-2 would have been ineffective because the reinforcement rate would be the same for both CSs (each trial would include one and only one US), which would explain why we saw no effect of reinforcement rate on extinction in this group. This post-hoc explanation of the present results does lead to several testable predictions. If the occurrence of a US part way through a CS divides that event into two trials, then a schedule that delivers one US part way through each presentation of the CS will constitute 50% reinforcement because the second half of each CS presentation will be a new non-reinforced trial. This should have two consequences. It should slow down acquisition of responding and reduce the terminal level of responding when compared with

a CS trained that has the same CS-US interval but without the “over-hanging” CS segment that remains after the US delivery. The second consequence is that it should produce a PREE.

We do not know of any empirical findings that speak to the second of the predictions just proposed, but there are findings relevant to the first. Recently, Austen and Sanderson (2019) conducted experiments in which mice were presented with CSs that had fixed CS-US intervals (a 10-s interval between the onset of the CS and delivery of a food pellet) but differed in how long the CS remained after US delivery on each trial. They found that acquisition of responding was unaffected by the over-hanging CS segment. For example, mice were no slower to acquire responding to a CS that remained for 30 s after US delivery than to a CS that terminated at the moment of US delivery. We have observed a similar result using variable CS-US intervals with rats. In an unpublished experiment using procedures very similar to those described here, we trained rats with four CSs, one of which was presented for a fixed duration of 60 s every trial and a food pellet was delivered at a variable time with a mean of 30 s (range 2-58 s) after CS onset. Acquisition of responding to this CS was indistinguishable from a second CS that had a variable duration of 30 s and was always reinforced with a food pellet at CS offset, $F(1,14) = 0.10$, $p = .755$ (see Figure 2). Therefore, extending the first CS for, on average, 30 s beyond the delivery of the US did not appear to impact on the acquisition of responding. Responding to this first CS was, nonetheless, significantly greater than to a third CS that had a mean duration of 30 s and was reinforced at CS offset on 50% of trials, $F(1,14) = 8.60$, $p = .011$, $\eta^2_p = 0.38$, as well as being significantly greater than to a fourth CS that had a mean duration of 60 s and was reinforced at CS offset on every trial, $F(1,14) = 16.34$, $p = .001$, $\eta^2_p = 0.53$. Therefore, presenting the CS for an additional 30 s either before the US or as a separate non-reinforced trial did affect the acquisition of responding. These results are inconsistent

with the suggestion that delivery of the US part way through a CS presentation creates two conditioning trials.

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

Experiment 2

Experiments 1a and 1b failed to find evidence linking the speed of extinction with the rate of reinforcement beyond the effect of partial reinforcement in slowing subsequent extinction. This supports the assumption made by the sequential account that the PREE depends specifically on learning about non-reinforced trials. Experiment 2 took a different approach to test this assumption more directly. It adapted a design that we previously used to tease apart what animals learn about a CS conditioned on an omission schedule versus one conditioned with a yoked schedule (Harris, Andrew, & Kwok, 2013). The current design involved two types of trials: on some trials a single stimulus, A, was presented alone; on other trials it was presented in compound with a second stimulus, as AX. One of these stimulus configurations was consistently reinforced (after an average length of 30 s) and the other configuration was partially reinforced (on 1 out of 3 trials, after an average length of 10 s). Responding to both configurations was then extinguished. Because the overall rates of reinforcement per unit time were matched, we expected rats to acquire responding to A and AX at the same rate. The critical question was whether we would see evidence for a PREE given that this would depend on rats being able to discriminate between trials with A and trials with AX, rather than between the two very different stimuli (A versus B) used in Experiments 1a and 1b. In particular, we were interested in whether evidence for the PREE would depend on which configuration had been consistently reinforced and which had been

partially reinforced. To this end, the experiment tested two groups of rats. For one group, termed “AX- partial”, A was conditioned with the consistent reinforcement schedule, and AX was conditioned with the partial schedule. For the other group, termed “AX-consistent”, AX was given consistent reinforcement and A was given partial reinforcement. A summary of the design is shown in Table 2.

----- Table 2 about here -----

In both groups, evidence for a PREE depends on the rats learning that the two configurations had different reinforcement schedules. More specifically, if the PREE depends on learning about non-reinforcement in the partial schedule, then we should only see a difference in rate of extinction if the rats have learned which of the two configurations was associated with partial reinforcement. Critically, learning this discrimination should not be equally easy for the two groups due to the asymmetrical nature of generalisation between A and AX. Studies of excitatory conditioning and extinction have shown that generalisation from A to AX is greater than from AX to A (e.g., Bouton, Doyle-Burr, & Vurbic, 2012; Brandon, Vogel, & Wagner, 2000). This explains the common observation that animals master a feature positive discrimination (AX+ vs A-) faster than they master a feature negative discrimination (A+ vs AX-) (Holland, 1991; Reberg & Leclerc, 1977; Sainsbury & Jenkins, 1967). In more general terms, this type of discrimination is easier when the added feature, X, is paired with the presence of the target outcome (that which supports a difference in responding or, in the present case, a difference in extinction) than when X is paired with the absence of that outcome. If, as assumed by the sequential account, evidence for a PREE depends on rats learning that non-reinforced trials are part of the reinforcement schedule of one configuration but not the other configuration, then this difference should be easier to learn

with A+ vs AX+/- training than with AX+ vs A+/- training because non-reinforced trials are associated with X in the former case but not the latter case. That is, the sequential account predicts a stronger PREE for the AX- partial group than for the AX-consistent group.

The same prediction hold for Frustration Theory (Amsel, 1958, 1962): if X is associated with frustration, this will slow extinction of AX but not A, whereas if A is associated with frustration it will slow extinction of A and AX. The alternative account that we have entertained here does not make this prediction. According to that account, there is nothing special about partial reinforcement. Extinction is lawfully proportional to the rate of reinforcement during conditioning, and this applies equally for partial and consistent reinforcement. In this sense, the PREE is a misnomer because one could equally describe evidence for the PREE as “faster extinction after continuous reinforcement” or “slower extinction after partial reinforcement”. This symmetry in the way the reinforcement schedules affect extinction means that the asymmetry in generalization between AX and A will not be of any consequence. Any difference in extinction should be equally large (or small) because it won’t matter whether learning from the consistent schedule generalizes to the inconsistent schedule or vice versa. It argues that any difference in the speed of extinction between two CSs depends on the mutually-defined difference in their reinforcement schedules. Therefore, the difference in reinforcement rate should be learned equally well regardless of whether X is associated with the consistent or partial schedules.

Methods

Subjects and Apparatus

Thirty-two experimentally-naïve female albino Sprague Dawley rats (8 to 10 weeks) were obtained from the same source and were housed in the same manner as described for

Experiments 1a and 1b. The rats were trained in the same conditioning chambers, with the same food pellets, and using the same stimuli, as described previously.

Procedure

Rats were first given magazine training, as described for Experiments 1a and 1b. They were then split into equal two groups ($n=16$) before commencing 20 sessions of training containing intermixed trials with A, AX, and C. The light and white noise stimuli served as A and X in counterbalanced fashion, whereas the clicker served as stimulus C for all rats. For both groups, stimulus C was presented and reinforced 14 times per session. The food pellet was delivered at the offset of each presentation, which varied from trial to trial with a mean duration of 17 s (range 2-32 s). Trials with C were intermixed randomly among trials with A and AX, and these differed between the two groups.

Group AX-partial received 8 trials of A and 24 trials of AX each session; Group AX-consistent received 24 trials of A and 8 trials of AX. Seven trials of each type were reinforced with delivery of a single food pellet. Thus, strictly speaking, both configurations had inconsistent reinforcement schedules. Nonetheless, one schedule was much more consistently reinforced (7 out of 8 trials) than the other (7 out of 24 trials), and this difference does reliably produce a PREE (Harris & Kwok, 2018). The advantage of these schedules is that, because they include non-reinforced trials for both trial types, they allow us to record post-trial response rates (immediately after CS offset) that are not contaminated by the presence of the food pellet itself. Using these exact schedules, we recently showed that post-CS responding is sensitive to the likelihood of reinforcement per trial, even when responding during the same CSs is not (Harris & Kwok, 2018). Using these schedules here allowed us to test for differences in post-CS responding after non-reinforced trials of A or AX.

Another procedural difference from the previous two experiments was that, on reinforced trials of A or AX, the food pellet was delivered during the CS presentation rather than at the end. On these trials, the CS was presented for a fixed duration and the food pellet was delivered at a random time during the CS. The more consistently reinforced configuration (i.e., A for Group AX-partial, and AX for Group AX-consistent) was presented for a fixed duration of 60 s, and the food pellet was delivered between 2 s and 58 s into the presentation. The one non-reinforced presentation of this configuration lasted for a variable duration with a mean of 30 s. The other configuration (AX for Group AX-partial, A for Group AX-consistent) was presented for a fixed duration of 20 s, and the food pellet was delivered between 2 and 18 s. The 17 non-reinforced presentations of this configuration lasted for a variable duration with mean of 10 s. The food pellet was delivered during, rather than at the end of, the CS presentation in this experiment to prevent rats from learning that stimulus offset is the most reliable predictor of food delivery, which might otherwise drive post-CS responding. On reinforced trials of A or AX, magazine activity was only recorded for the period prior to delivery of the food pellet so that response rates were not contaminated by activity induced by the pellet itself. On non-reinforced trials of A and AX, we recorded post-trial activity for a period of 15 s after CS offset. The mean inter-trial interval was 120 s (range 40 to 200 s), and the sessions lasted approximately 2 h.

Both groups received identical schedules across Sessions 21-35. In each session, A and AX were presented 12 times each for a mean duration of 17 s (range 2 to 32 s) and were never reinforced with food. Intermixed among these extinction trials were 24 reinforced trials of the clicker, C. The mean duration of these presentations was 17 s (2 to 32 s) and all presentations terminated with delivery of a food pellet. The mean inter-trial interval was again 120 s (range 40 to 200 s), and the session length was approximately 2 h.

Results

One rat from Group AX-consistent was removed from the experiment due to ill health. The acquisition of responding to A and AX across Sessions 1 to 20, and the subsequent extinction of that responding (Sessions 21-35) is shown for both groups in Figure 3. Both groups acquired responding to both configurations in very similar fashion. Responding to A and AX extinguished equally quickly for Group AX-consistent. Extinction was slower for AX than for A in Group AX-partial, providing evidence for a PREE in that group.

----- Figure 3 about here -----

A 2 x (20) mixed model ANOVA comparing the two groups on their pre-CS response rates across Sessions 1-20 showed that there was no significant difference overall between the groups, $F(1,29) = 0.91$, $p = .348$. There was a significant main effect of Session, $F(5.9,177) = 4.06$, $p = .001$, $\eta^2_p = 0.12$, that did not interact with Group, $F(5.9,177) = 1.22$, $p = .297$. A similar 2 x (15) mixed model ANOVA comparing pre-CS response rates across Sessions 21-35 also showed there was no significant difference overall between the groups, $F(1,29) = 0.70$, $p = .408$. The main effect of Session was not significant, $F(4.4,128) = 1.91$, $p = .105$, and did not interact with Group, $F(4.4,128) = 1.37$, $p = .244$.

The response rates across acquisition (Sessions 1-20) were first analysed using a 2 x (2 x 20) mixed model ANOVA. There was a significant main effect of Session, $F(5.1,148) = 19.84$, $p < .001$, $\eta^2_p = 0.41$, but no main effect of Group, $F(1,29) = .112$, $p = .740$, or Trial-type, $F(1,29) = 2.51$, $p = .124$. There was a significant interaction between Group and Trial-type, $F(1,29) = 4.58$, $p = .041$, $\eta^2_p = 0.14$, but Session did not interact with either Group, $F(5.1,148) = 0.35$, $p = .883$, or Trial-type, $F(7.4,214) = 1.59$, $p = .136$, and there was no 3-way interaction, $F(7.4,214) = 1.51$, $p = .163$.

To test for the source of the significant Group by Trial-Type interaction, for each group a simple pairwise comparison was made between the response rates to A and AX, averaged over the 20 acquisition sessions. In Group AX-partial, there was no difference in responding to A and AX, $t(15) = 0.43$, $p = .675$. A Bayesian paired samples t-test showed there was modest evidence that responding to A and to AX was the same ($BF_{01} = 3.6$). In Group AX-consistent, there was a significant difference in responding to A versus AX, $t(14) = 2.44$, $p = .029$, $d = 0.63$. It is clear from Figure 3 that the difference in responding to A and AX had all but disappeared by the last few sessions of conditioning. To test this, a Bayesian paired sample t-test was conducted on response rates average over Sessions 18-20. This revealed there was moderate evidence that response rates to A and AX were equal ($BF_{01} = 3.7$).

The data from the extinction sessions (21-35) were first analysed using a 2 x (2 x 15) mixed model ANOVA. There was no overall difference between the two groups, $F(1,29) = .60$, $p = .447$, but there was a significant overall difference between responding to A and AX, $F(1,29) = 10.52$, $p = .003$, $\eta^2_p = 0.27$, which interacted significantly with the factor of Group, $F(1,29) = 16.29$, $p < .001$, $\eta^2_p = 0.36$. The significant main effect of Session, $F(2.3,68) = 67.14$, $p < .001$, $\eta^2_p = 0.70$, did not interact significantly with Trial-type, $F(5.3,155) = 1.64$, $p = .148$, or Group, $F(2.3,68) = 0.78$, $p = .479$, but there was a significant 3-way interaction between Session, Trial-type, and Group, $F(5.3,155) = 2.67$, $p = .021$, $\eta^2_p = 0.08$. Both the Group by Trial-type interaction, and the 3-way interaction, demonstrate that the difference in extinction between A and AX was significantly greater in Group AX-partial than in Group AX-consistent.

The mixed model ANOVA was followed up with two 2 x 15 repeated measures ANOVAs conducted on the extinction data from each group individually. Responding across extinction in Group AX-partial showed a main effect of Session, $F(1.9,28) = 35.36$, $p < .001$, $\eta^2_p = 0.70$,

and a significant difference in overall responding between A and AX, $F(1,15) = 18.78$, $p = .001$, $\eta^2_p = 0.56$, as well as a significant interaction between these factors, $F(3.3,49) = 3.01$, $p < .035$, $\eta^2_p = 0.17$. Analysis of responding during the extinction phase in Group AX-consistent showed a significant main effect of Session, $F(2.4,34) = 32.46$, $p < .001$, $\eta^2_p = 0.70$, but no difference in responding overall to A and AX, $F(1,14) = 0.60$, $p = .453$, nor an interaction between this difference and Sessions, $F(5.2,73) = 0.96$, $p = .448$.

The graph on the right-hand side of Figure 3 plots the response rates to A and AX for both groups together. This makes it clear that the rate of extinction of responding to A and AX in Group AX-consistent is very similar to the rate of extinction of responding to AX in Group AX-partial, whereas responding to A in the latter group extinguished more quickly. To test this, we have conducted four between-group pairwise comparisons using response rates averaged over the 15 extinction sessions. These confirmed that responding to A in Group AX-partial was significantly lower than responding to both A and AX in Group AX-consistent: for A vs A, $t(29) = 3.01$, $p = .005$, Cohen's $d = 1.08$ (95% Confidence Intervals on $d = 0.35$ and 1.82); for A vs AX, $t(29) = 2.14$, $p = .041$, $d = 0.77$ (0.03 and 1.50). Responding to AX in Group AX-partial was not significantly difference from either A or AX in Group AX-consistent: for AX vs A, $t(29) = 0.46$, $p = .643$; for AX vs AX, $t(29) = 0.67$, $p = .507$.

The reinforcement schedules used in this experiment, which included one non-reinforced trial per session of the consistently reinforced configuration, provided us with the opportunity to record post-CS response rates. We have recently shown post-CS responding to be sensitive to the trial-based likelihood of reinforcement (Harris & Kwok, 2018). However, in this experiment we did not observe any such difference in responding post presentations of A and AX in either group. In Group AX-partial, there was no difference overall in post-trial

responding between trials with A and trials with AX, $F(1,15) = 2.41$, $p = .141$, nor was there a change in post-trial responding over sessions, $F(6.5,97) = 1.06$, $p = .396$, nor was there an interaction between Trial-type and Session, $F(7.1,106) = 1.71$, $p < .114$. Similarly, in Group AX-consistent, there was no difference between A and AX trials, $F(1,14) = 3.73$, $p = .088$, or across sessions, $F(5.6,78) = 1.67$, $p = .144$, nor was there an interaction between Trial-type and Session, $F(6.8,95) = 0.82$, $p = .568$.

Discussion

In this experiment, rats showed slower extinction of responding to a compound, AX, that was partially reinforced (7 out of 24 trials) than to the single stimulus, A, that was more consistently reinforced (7 out of 8 trials). However, a second group of rats did not show any difference in extinction of responding to the same configurations when A had been partially reinforced and AX had been more consistently reinforced. Therefore, for the two reinforcement schedules to have a differential impact on extinction, the rats must be able to learn that the unique feature, X, is associated with the occurrence of non-reinforced trials. This confirms the assumption of the sequential account that a PREE depends on animals learning about the presence of non-reinforced trials in the partial reinforcement schedule.

The absence of a difference in extinction between A and AX in Group AX-consistent is, on face value, ambiguous. It could mean that extinction was fast for both configurations (i.e., there no PREE for A or AX) or that extinction was slow for both (i.e., there was a PREE for both A and AX). To test between these alternatives, we plotted responding to A and AX for both groups on the same graph on the right side of Figure 3. This reveals that responding to A and AX in Group AX-consistent extinguished at the same slow rate as responding to AX in Group AX-partial. In other words, it appears that Group AX-consistent showed a PREE to both A and

AX. This is what one would expect if the partial reinforcement schedule of stimulus A meant that the rats learned A was associated with both reinforced and non-reinforced trials, and thus responding to A, whether presented alone or in compound with X, would extinguish slowly.

Experiment 3

Experiment 2 showed that rats must learn that the added feature, X, is specifically associated with non-reinforced trials in the partial schedule in order to show differential extinction to A vs AX. If, instead, X is associated with consistent reinforcement, it exerts no differential control over extinction. However, there is a potential problem with the evidence for this difference in the effect of X on extinction. In Experiment 2, the influence of X on extinction was tested against the influence of A on extinction. This is problematic because the two groups in Experiment 2 differed in what they learned about A as well as what they learned about X, as reflected in the fact that the two groups differed in how quickly responding to A itself extinguished. An alternative way to assess any difference between the groups in the influence of X over extinction would be to test X in compound with a third stimulus that had been given identical treatment for the two groups. Experiment 3 was designed for this purpose.

Experiment 3 trained two groups of rats for 20 sessions on an A vs AX discrimination, with one configuration being consistently reinforced (which, in this experiment, meant a 100% reinforcement schedule) and the other being partially reinforced (33%). Group AX-partial was trained with A+ vs AX+/-, and Group AX-consistent was trained with A+/- vs AX+. Both groups were concurrently trained with another stimulus, B, which was consistently reinforced. (As for the previous experiments, they also received reinforced trials with a clicker

CS, which continued to be reinforced throughout the entire experiment to prevent extinction of all magazine activity.) For both groups, responding to the compound BX was extinguished to compare the effect of X on extinction between the two groups. These extinction sessions also included occasional probe trials (one per session) in which B was presented alone, to test for any differences in extinction of B itself.

The experiment also provided an opportunity to test whether X exerted a direct influence over responding, independently of its possible effect on extinction, and whether this differed between the groups. In particular, we were interested in the possibility that X, when associated with partial reinforcement in Group AX-partial, might acquire some inhibitory control. There was no evidence for such an effect in Experiment 2, when comparing responding between A and AX during training. However, as already noted, this comparison confounds any putative difference between groups in the control exerted by X with a difference in the control exerted by A. The design of Experiment 3 provided a means to avoid this confound by assessing the control exerted by X when tested against B which received the same training across the groups. Therefore, before we extinguished responding to the compound BX, we tested the influence of X on responding to B during a further 8 training sessions that included non-reinforced probe trials (one per session) with BX. A summary of the design is shown in Table 3.

----- Table 3 about here -----

Methods

Subjects and Apparatus

Thirty-two experimentally-naïve female albino Sprague Dawley rats (8 to 10 weeks) were obtained from the same source and were housed in the same manner as described for Experiment 1. The rats were trained in the same conditioning chambers, with the same food pellets, and using the same stimuli, as described previously. This experiment required an additional auditory stimulus that was not used in the previous experiments. It was a high-pitch tone (2.9 kHz, 17 dB above background), produced from a piezo buzzer positioned on the floor of the sound-attenuating shell behind each conditioning chamber. This tone, and the white noise stimulus, were assigned as Stimulus A and Stimulus B in counterbalanced fashion. The steady light served as X for all rats. This arrangement was used because the design required that X be combined in compound with both A and B, which ideally meant that X should be in a different modality from both A and B for all rats. This could only be achieved with these particular stimuli if the light always served as X.

Procedure

After magazine training, the rats began 20 sessions of training, with 48 trials per session. In each session for Group AX-partial, there were 8 reinforced trials of A, each with a fixed duration of 60 s and a single food pellet delivered at a random time between 2 and 58 s. There were 24 trials with AX, 8 of which had a fixed duration of 20 s and were reinforced with a food pellet at a random time between 2 and 18 s, and the other 16 varied in duration between 2 and 18 s (mean 10 s) and were not reinforced with food. For Group AX-consistent, there were 8 trials of AX (60-s duration, all reinforced with food at random time between 2 and 58 s) and 24 trials of A (8 trials of 20-s duration reinforced with food at a random time between 2 and

18 s, and 16 non-reinforced trials of variable duration between 2 and 18 s). For both groups, there were also 8 reinforced trials of B and of the clicker. These trials had a fixed duration of 20 s, and food was delivered at a random time between 2 and 18 s during the stimulus. The mean inter-trial interval was 120 s (25 to 215 s), and each session lasted approximately 2 hr 20 min.

The training schedule was extended for a further 8 sessions (Sessions 21-28). These sessions were identical to the first 20 sessions except that they included a single probe trial with the compound BX. This trial had a fixed duration of 20 s and was not reinforced. It occurred mid-way through the session (as trial 25 out of 49).

Responding to BX was extinguished across Sessions 29 to 40. For both groups, the compound BX was presented 12 times, for a mean duration of 10 s (varying between 2 and 18 s), and was never reinforced with food. These were intermixed among 24 reinforced trials with the clicker (each had a fixed 20-s duration, with food delivered at a random time during the stimulus). Each session also included a single probe trial with B, occurring on Trial 19 (out of 37). On this trial, B was presented for a fixed duration of 10 s and not reinforced. The inter-trial interval was 140 s (40 to 240 s), and each session lasted approximately 2.5 h.

Results

A 2 x (20) mixed model ANOVA comparing the two groups on their pre-CS response rates across Sessions 1-20 showed that there was no significant difference overall between the groups, $F(1,29) = 0.16$, $p = .690$. There was a significant main effect of Session, $F(3.9,117) = 6.64$, $p < .001$, $\eta^2_p = 0.18$, that did not interact with Group, $F(3.9,117) = .30$, $p = .872$. A 2 x

(15) mixed model ANOVA comparing pre-CS response rates across Sessions 21-35 also showed there was no significant difference overall between the groups, $F(1,29) = 0.19, p = .667$. The main effect of Session was significant, $F(4.6,134) = 3.58, p = .006, \eta^2_p = 0.11$, but did not interact with Group, $F(4.6,134) = 0.38, p = .849$.

Both groups of rats acquired responding to A, AX, and B across Sessions 1 to 20. A 2 x (3 x 20) mixed model ANOVA on the data from Sessions 1-20 showed there was no overall difference between the groups, $F(1,30) = 0.19, p = .666$. There was a significant main effect of CS type, $F(1.4,41) = 23.16, p < .001, \eta^2_p = 0.44$, that did not interact with group, $F(1.4,41) = 0.48, p = .552$. There was a significant main effect of Session, $F(3.6,108) = 19.56, p < .001, \eta^2_p = 0.40$, that did not interact with Group, $F(3.6,108) = 0.40, p = .789$, but did interact with trial type $F(8.5,256) = 2.69, p = .006, \eta^2_p = .08$. The 3-way interaction with groups was not significant $F(8.5,256) = 1.07, p = .388$. Two orthogonal contrasts were used to unpack the main effect of CS type. One contrast compared responding to B with responding to A and AX, the other compared responding to A versus AX. The first contrast showed that responding to B was significantly higher than responding to A and AX, $F(1,15) = 42.24, p < .001$. The second contrast found there was no difference in responding between A and AX, $F(1,15) = 0.18, p = .676$. Responding to B was higher than to A and AX in both groups because B was reinforced at a higher rate (see Harris & Carpenter, 2011).

Responding during B, and during probe trials with BX, averaged across Sessions 21-28, are shown for both groups in the left part of Figure 4. A 2 x (2) mixed model ANOVA showed that there was no difference overall between the two groups, $F(1,29) = 0.18, p = .672$. There was a significant difference in responding to B and responding during BX, averaged across groups, $F(1,29) = 7.41, p = .011, \eta^2_p = 0.20$, but there was no interaction between these

factors, $F(1,29) = 0.03$, $p = .863$. Therefore, while X produced a small decrease in responding to B, this effect did not differ between the two groups, indicating that any effect of X was not based on learning about its association with a specific reinforcement schedule.

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

Responding to BX, and during probe trials with B, across Sessions 28-40 are shown in the right graph in Figure 4. Responding to BX extinguished across Sessions 29-40 in both groups and appeared to extinguish more slowly in Group AX-partial. A $2 \times (2 \times 12)$ mixed model ANOVA showed that there was no difference overall between the two groups, $F(1,29) = 0.00$, $p = .961$, but there was a main effect of Session, $F(1,29) = 18.76$, $p < .001$, $\eta^2_p = 0.39$, and a main effect for Trial-type, $F(1,29) = 26.54$, $p < .001$, $\eta^2_p = 0.48$. This latter effect shows that responding to BX was less than to B. The effect of Trial-type interacted significantly with Session, $F(6.4,187) = 3.38$, $p = .003$, $\eta^2_p = 0.10$, but not with Group, $F(1,29) = 3.71$, $p = .064$. There was not a significant interaction between Group and Session, $F(4.7,137) = 0.52$, $p = .750$, nor was there a 3-way interaction, $F(6.4,187) = .661$, $p = .692$. Some caution is required when interpreting the outcome of these statistical comparisons between BX and B because response rates to BX were sampled across 12 trials per session whereas response rates to B were sampled from only one probe trial per session.

The comparison of key interest is between the two groups in their levels of responding to BX. A $2 \times (12)$ mixed model ANOVA showed there was a significant main effect of Sessions, $F(2.5,71) = 7.41$, $p = .011$, $\eta^2_p = 0.20$, and a significant main effect of Group, $F(1,29) = 6.77$, $p = .014$, $\eta^2_p = 0.19$, but no interaction between Group and Sessions, $F(2.5,71) = 0.47$, $p = .664$. The significant overall difference between the groups establishes that responding to BX was

lower among rats in Group AX-consistent than among rats in Group AX-partial. The absence of a significant interaction with session might imply equivalent rates of extinction between the two groups. However, since there was no such difference between the groups prior to extinction (as seen on Session 28 in Figure 4), the significant overall difference between the groups can only mean that responding to BX decreased more in Group AX-consistent than in Group AX-partial.

The equivalent ANOVA was conducted on the data from probe trials with B. It also showed there was a significant main effect of Sessions, $F(6.6,193) = 7.98, p < .001, \eta^2_p = 0.22$, confirming that responding to B extinguished over sessions. However, there was no main effect of Group, $F(1,29) = 0.60, p = .447$, nor was there a Group by Sessions interaction, $F(6.6,193) = 0.60, p = .747$.

Discussion

Experiment 3 has confirmed that when rats are trained on an A vs AX discrimination in which one configuration is reinforced consistently and the other partially, the cue X comes to exert a different influence over extinction depending on whether the configuration AX had been paired with consistent or partial reinforcement. In this experiment, X was extinguished in compound with a third CS, B, which had been consistently reinforced on its own. Therefore, any difference in the extinction of responding to BX was due to differences in the influence of X. We observed that responding to BX extinguished more slowly if AX had been partially reinforced (and A consistently reinforced) than if AX had been consistently reinforced (and A partially reinforced). The difference between the groups in extinction of their responding to BX shows that they differed in what they learned about X.

The difference in extinction of responding to BX contrasts with our earlier observation that responding to AX extinguished in near identical fashion for the two groups in Experiment 3. That result can be attributed to differences in what the two groups learned about A alone, since A alone was consistently reinforced in Group AX-partial but was partially reinforced in Group AX-consistent. Thus, in Experiment 3, extinction of responding to AX was equally slow for both groups but for different reasons: in Group AX-partial, AX extinguished slowly because X has been associated with inconsistent reinforcement, producing a PREE; in Group AX-consistent, AX extinguished slowly because A was associated with partial reinforcement, which produced a PREE. It is only when X was combined with a third CS, B, which had been given consistent reinforcement for both groups, that we were able to see the difference in X's influence over extinction between the groups.

While the effect of X on extinction was different between the two groups in this experiment, there was no difference in its effect on responding to B prior to extinction. X produced a small but reliable decrease in responding to B in both groups. The fact that the effect was the same in both groups shows that it was not a product of learning about partial or consistent reinforcement schedules, and was not related to the subsequent influence of X over extinction. Perhaps the most parsimonious explanation is that the slight decrease in responding reflected some external inhibition caused by the novelty of the BX configuration. It is worth noting that in numerous previous experiments we have reliably observed summation of responding when two CSs are presented in compound, even when that compound is novel and is presented as occasional probe trials such as those used here. Therefore, the fact that responding to BX was not above that to B suggests that X did not acquire its own excitatory associative strength that could summate with B's. This is plausible given that, in the A vs AX design, X did not signal any change in overall reinforcement rate,

thus rendering X redundant as a predictor of reinforcement. Yet, at the same time, the present experiments show that rats did learn something about X when it signalled partial reinforcement. All we can say about this learning is that it neither promoted nor suppressed responding per se, but it did selectively influence subsequent extinction of responding.

General Discussion

Four experiments explored how different schedules of reinforcement give rise to subsequent differences in how quickly responding extinguishes. Each experiment compared two reinforcement schedules that were matched for their total rate of reinforcement (i.e., the cumulative duration of exposure to the CS per reinforcement) to equate the levels of responding across conditioning. This meant that subsequent differences in extinction were not confounded by differences in initial levels of responding. The experiments were designed to test between two explanations of the PREE. According to one explanation, extinction depends on the animal detecting that an expected US has been omitted (Gallistel, 2012; Gallistel & Gibbon, 2000; Gibbon et al., 1980). Therefore, more consistent reinforcement is followed by more rapid extinction because it increases how often the US is expected, and thus increases how often the absence of the US is detected. By this account, the difference in extinction after consistent versus partial reinforcement is mutually determined by the difference in the rate of reinforcement per trial. This contrasts with most other explanations of the PREE that specifically identify the difference in extinction with perseveration of responding after partial reinforcement due to the presence of non-reinforced trials during that schedule (Amsel, 1958, 1962; Capaldi, 1966, 1994; Sheffield, 1949). In other words, partial reinforcement includes an additional feature—non-reinforced trials—that is responsible for the difference in extinction following consistent and partial reinforcement.

Using within-subjects designs, Experiments 1a and 1b confirmed that extinction proceeds more slowly after a 50% reinforcement schedule, which includes reinforced and non-reinforced trials, than after a consistent reinforcement schedule in which all trials are reinforced. However, while removing reinforcement from 50% of trials slows subsequent extinction, these experiments found no evidence for the complementary effect, that adding extra USs to each trial accelerated subsequent extinction. That is, there was no difference in subsequent extinction between a conditioning schedule that delivered one US per trial and a schedule that delivered two USs per trial. A difference in extinction would be predicted under these circumstances if extinction were tied to each omission of an expected reinforcer. According to that account, reducing the likelihood of reinforcement to 50% per trial should reduce how often animals expect reinforcement and therefore reduce the frequency that they detect the omission of an expected US during extinction. It follows, from this account, that doubling the number of reinforcers per trial should increase how often animals expect reinforcement per trial and so increase the frequency with which they detect the omission of an expected US. The fact that we did not observe any such effect of doubling the number of reinforcers per trial indicates that differences in extinction are not a consequence of any difference in the likelihood of reinforcement per trial. Instead, differences in extinction are only observed when the likelihood of reinforcement is reduced below one US per trial, which supports the argument that changes in extinction depend on animals learning about the presence of non-reinforced trials during conditioning.

As noted earlier, these conclusions depend on an assumption that each presentation of the CS constitutes a single trial, and any presentation that includes two USs counts as one double-reinforced trial. If, on the other hand, the occurrence of the US itself also defines a trial, then a long CS presentation that includes two USs might in fact represent two

independent trials. This would mean that, for Group 1-vs-2, both CSs were reinforced at the same rate, thus explaining the absence of any difference in extinction between the CSs. Moving beyond this as a *post-hoc* explanation of the present finding, this definition of a trial predicts that a long CS with a single food pellet delivered part way through the presentation will be treated as a 50% partially-reinforced CS because the segment before the US will constitute one reinforced trial and the segment after the US will constitute a non-reinforced trial. Responding to such a CS should be acquired more slowly, and also extinguish more slowly, than responding to a CS that is reinforced after the same CS-US interval but which terminates at US delivery. However, recent findings by Austen and Sanderson (2019), and some unpublished findings from our laboratory (presented in Figure 2), indicate that this prediction concerning the acquisition of responding is not correct. A non-reinforced CS segment that extends after delivery of the US does not impact negatively on the acquisition of responding. However, we are not aware of any evidence that speaks to the other prediction, that extending the CS beyond the US will retard subsequent extinction.

Experiment 2 used another approach to identify what it is about consistent versus partial reinforcement that results in differences in extinction. As described above, most theories argue that the PREE is due to learning about the presence of non-reinforced trials in the partial schedule. That is, the difference in extinction following consistent and partial reinforcement is attribute to a feature that is added to the partial schedule. To test this account, rats were trained concurrently on two different schedules, one consistent and one partial, with two CS configurations—a single CS, A, and a compound CS, AX. For one group of rats (Group AX-partial) A was consistently reinforced and AX was partially reinforced; for the second group (Group AX-consistent) this assignment was reversed. Evidence for a PREE will depend on animals learning that one of the configurations is associated with non-reinforced

trials. However, if that learning generalises to the other configuration, this will reduce the observed difference in rate of extinction of A versus AX. Crucially, generalisation between A and AX is asymmetrical: generalisation is greater from A to AX than from AX to A (Bouton et al., 2012; Brandon et al., 2000). Therefore, we expected that if A is associated with non-reinforcement, we should see a PREE to A that generalises strongly to AX, whereas if AX was associated with non-reinforcement, we should see a PREE to AX that generalises much less to A. This prediction was confirmed in Experiment 2 when responding to both A and AX was extinguished in the two groups. Extinction was slower for AX than for A in Group AX-partial, whereas A and AX extinguished at the same rate in Group AX-consistent. In the latter group, the rate at which A and AX extinguished matched the slower rate of extinction shown to AX in the former group. The results are incompatible with the argument that the PREE results from a mutually defined difference in reinforcement rates because this account predicts that a PREE should be equally evident following A versus AX training regardless of which configuration was paired with consistent reinforcement and which with partial reinforcement.

The two groups in Experiment 2 differed in their extinction of AX relative to A, but this appeared to arise from a difference between the groups in how quickly responding to A extinguished—extinction of A was slower in Group AX-consistent than Group AX-partial, whereas extinction of AX was very similar across the groups. A reasonable explanation for this is that the rats in Group AX-consistent showed a PREE to A because it had been partially reinforced during conditioning, whereas rats in Group AX-partial showed rapid extinction of A because it, on its own, had been consistently reinforced during conditioning. However, the difference in extinction of A complicates interpretation of the finding that X slowed extinction of responding in Group AX-partial but not in Group AX-consistent. To demonstrate that X itself

had different effects on extinction between the groups, it was necessary to extinguish responding to X when in compound with a common stimulus that had been given identical conditioning in both groups. This was addressed in Experiment 3. Two groups were given the same training with A and AX as those in Experiment 2, but were also trained with a third CS, B, which was given the same consistent reinforcement schedule in both groups. When responding to the compound BX was extinguished, that extinction was significantly slower in Group AX-partial than in Group AX-consistent, confirming that X differentially affected extinction depending on whether it had been associated with consistent or partial reinforcement.

What do these findings mean for theories of the PREE? First, as already mentioned, they establish the importance of non-reinforcement for the effect. A difference in per-trial reinforcement rate does not in and of itself create a difference in susceptibility to extinction. One of the CSs must have been conditioned on a partial reinforcement schedule that included non-reinforced trials intermixed among reinforced ones. Further, the difference in extinction after partial versus consistent reinforcement depends on learning that is specific to the partial schedule—a difference in extinction is only observed if animals do not generalise what they learned about the partial CS to the consistent CS. These conclusions are broadly consistent with the theoretical account of the PREE that was developed through the 1940s, 50s, and 60s, (Amsel, 1958, 1962; Capaldi, 1966, 1994; Hull, 1941, cited in Jenkins, 1950; Sheffield, 1949). According to these theories, a partial reinforcement schedule creates conditions under which the after-effects of non-reinforcement, variously defined, can become cues that are associated with subsequent reinforcement. As a result, the experience of non-reinforcement will, in itself, promote responding under extinction in a way that does not occur following

consistent reinforcement because that schedule does not establish any association between non-reinforcement and subsequent reinforcement

The present findings confirm this emphasis on learning about non-reinforcement, but at the same time the evidence they present poses a significant challenge for those theories. Each demonstration of the PREE was made using a within-subjects design in which rats were concurrently trained with both a consistently reinforced and a partially reinforced CS on intermixed trials. If partial reinforcement produces resistance to extinction because it allows the after-effects of non-reinforcement to be associated with later reinforcement, then this resistance to extinction should be manifest for both the consistently reinforced CS as well as the partially reinforced CS in the within-subjects design. In other words, it should not be possible to observe a difference in extinction between a consistent and a partial CS if both have been conditioned in intermixed fashion in the same subjects.

The reliable evidence we have obtained here, and has been reported previously (Chan & Harris, 2017, 2019; Crawford, Steirn, & Pavlik, 1985; Harris & Kwok, 2018; Rescorla, 1999), for differential extinction after partial and consistent reinforcement within the same subjects means that the PREE is not due to learning about the after-effect of non-reinforcement *per se*, since this would promote generalisation of responding to both CSs during extinction. Rather, the experience of non-reinforcement during extinction must selectively promote generalisation of responding to the partial CS and not to the consistent CS. One way this might be achieved is if each experience of non-reinforcement encodes information about the CS that was not reinforced on that trial. For example, when a CS is presented without the usual reinforcement, the animal is not simply frustrated about the absence of the US, but is frustrated with the CS; or their memory of that event is not simply that the US was absent,

but that the particular CS was not reinforced. Because this reaction or memory stipulates which CS was not reinforced, it is only retrieved on subsequent trials with that CS. If this is part of a partial reinforcement schedule, then this memory can become associated with the US on trials when the CS is reinforced. During subsequent extinction, this CS-specific memory will only be retrieved on trials when the partial CS is presented, and therefore will only promote responding to that CS. On trials when the continuously reinforced CS is presented during extinction, the animal will learn that this CS is not reinforced (or learn to be frustrated with this CS) but this particular memory has never been associated with reinforcement and therefore helps to discriminate extinction from conditioning rather than promoting generalisation.

References

- Amsel, A. (1958). The role of frustrative nonreward in noncontinuous reward situations. *Psychological Bulletin*, 55, 102-119.
- Amsel, A. (1962). Frustrative non-reward in partial reinforcement and discrimination learning: Some recent history and a theoretical extension. *Psychological Review*, 69, 306-328.
- Austen, J. M., & Sanderson, D. J. (2019). Delay of reinforcement versus rate of reinforcement in Pavlovian conditioning. *Journal of Experimental Psychology: Animal Learning and Cognition*.
- Bacon, W. E. (1962). Partial-reinforcement extinction effect following different amounts of training. *Journal of Comparative and Physiological Psychology*, 55, 998-1003.
- Balsam, P. D., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neurosciences*, 32, 73-78.
- Bouton, M. E., Doyle-Burr, C., & Vurbic, D. (2012). Asymmetrical generalization of conditioning and extinction from compound to element and element to compound. *Journal of Experimental Psychology: Animal Behavior Processes*, 38, 381-393.
- 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.
- Brandon, S. E., Vogel, E. H., & Wagner, A. R. (2000). A componential view of configural cues in generalization and discrimination in Pavlovian conditioning. *Behavioural Brain Research*, 110(1-2), 67-72.
- 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. (2017). Extinction of Pavlovian conditioning: The influence of trial number and reinforcement history. *Behavioural Processes*, 141, 19-25.

- Chan, C. K. J., & Harris, J. A. (2019). The Partial Reinforcement Extinction Effect: The proportion of trials that are reinforced during conditioning predicts rate of extinction. *Journal of Experimental Psychology: Animal Learning and Cognition*, 45, 43-58.
- Crawford, L. L., Steirn, J. N., & Pavlik, W. B. (1985). Within- and between-subjects partial reinforcement effects with an autoshaped response using Japanese quail. *Animal Learning and Behavior*, 13, 85-92.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.
- Fitzgerald, R. D. (1963). Effects of partial reinforcement with acid on the classically conditioned salivary response in dogs. *Journal of Comparative and Physiological Psychology*, 56, 1056-1060.
- Fitzgerald, R. D., Vardaris, R. M., & Teyler, T. J. (1966). Effects of partial reinforcement followed by continuous reinforcement on classically conditioned heart rate in the dog. *Journal of Comparative and Physiological Psychology*, 62, 483-486.
- 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.
- Gibbon, J., Farrell, L., Locurto, C. M., Duncan, H. J., & Terrace, H. S. (1980). Partial reinforcement in autoshaping with pigeons. *Animal Learning and Behavior*, 8, 45-59.
- Harris, J. A., Andrew, B. J., & Kwok, D. W. S. (2013). Magazine approach during a signal for food depends on Pavlovian, not instrumental, conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 39, 107-116.
- Harris, J. A., & Carpenter, J. S. (2011). Response rate and reinforcement rate in Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 37, 375-384.
- Harris, J. A., Gharaei, S., & Pincham, H. L. (2011). Response rates track the history of reinforcement times. *Journal of Experimental Psychology: Animal Behavior Processes*, 37, 277-286.
- Harris, J. A., & Kwok, D. W. S. (2018). The probability of reinforcement per trial affects posttrial responding and subsequent extinction but not within-trial responding. *Journal of Experimental Psychology: Animal Learning and Cognition*, 44, 23-35.

- Harris, J. A., 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., Aydin, A., & Pearce, J. M. (2004). A Partial Reinforcement Extinction Effect despite equal rates of reinforcement during Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 30, 240-250.
- Holland, P. C. (1991). Transfer of control in ambiguous discriminations. *Journal of Experimental Psychology: Animal Behavior Processes*, 17(3), 231-248.
- Humphrey, L. G. (1939). The effect of random alternation of reinforcement on the acquisition and extinction of conditioned eyelid reactions. *Journal of Experimental Psychology*, 25, 141-158.
- Humphrey, L. G. (1940). Extinction of conditioned psychogalvanic responses following two conditions of reinforcement. *Journal of Experimental Psychology*, 27, 71-75.
- JASP_Team. (2018). JASP (Version 0.9, for Windows).
- Jenkins, W. O. (1950). Partial reinforcement: A review and critique. *Psychological Bulletin*, 47, 193-234.
- Mackintosh, N. J. (1974). *The psychology of animal learning*. London, UK: Academic Press.
- Mowrer, O. H., & Jones, H. (1945). Habit strength as a function of the pattern of reinforcement. *Journal of Experimental Psychology*, 35, 293-311.
- Reberg, D., & Leclerc, R. (1977). A feature positive effect in conditioned suppression. *Animal Learning and Behavior*, 5, 143-147.
- Rescorla, R. A. (1999). Within-subject partial reinforcement extinction effect in autoshaping. *Quarterly Journal of Experimental Psychology*, 52B, 75-87.
- Rogers, A. E. (1979). Nutrition. In H. J. Baker, J. R. Lindsey & S. H. Weisbroth (Eds.), *The laboratory rat* (Vol. 1, pp. 123-152). New York: Academic Press.
- Sainsbury, R. S., & Jenkins, H. M. (1967). Feature-positive effect in discrimination learning. *Proceedings of the annual convention of the American Psychological Association*, 2, 17-18.

Sheffield, V. F. (1949). Extinction as a function of partial reinforcement and distribution of practice.

Journal of Experimental Psychology, 39, 511-526.

Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century.

Weinstock, S. (1954). Resistance to extinction of a running response following partial reinforcement under widely spaced trials. *Journal of Comparative and Physiological Psychology, 47*, 318-322.

Table 1. Summary of designs for Experiments 1a and 1b.

Conditioning	Extinction
Experiment 1a	
Sessions 1-20	Sessions 21-30
Group 1-vs-2	
A 40 s, 100%, 2 pel (6)	A 14 s, 0%, 0 pel (8)
B 20 s, 100%, 1 pel (12)	B 14 s, 0%, 0 pel (8)
C 10 s, 100%, 1 pel (12)	C 10 s, 100%, 1 pel (24)
Group 1-vs-0.5	
A 10 s, 50%, 1 pel (24)	A 14 s, 0%, 0 pel (8)
B 20 s, 100%, 1 pel (12)	B 14 s, 0%, 0 pel (8)
C 10 s, 100%, 1 pel (12)	C 10 s, 100%, 1 pel (24)
Experiment 1b	
Sessions 1-20	Sessions 21-30
Group 1-vs-2	
A 40 s, 100%, 2 pel (8)	A 14 s, 0%, 0 pel (8)
B 20 s, 100%, 1 pel (8)	B 14 s, 0%, 0 pel (8)
C 10 s, 100%, 1 pel (8)	C 10 s, 100%, 1 pel (24)
Group 1-vs-0.5	
A 10 s, 50%, 1 pel (12)	A 14 s, 0%, 0 pel (8)
B 20 s, 100%, 1 pel (12)	B 14 s, 0%, 0 pel (8)
C 10 s, 100%, 1 pel (6)	C 10 s, 100%, 1 pel (24)

Note: The mean duration (in seconds) of each CS, and the percentage of trials reinforced with 1 or 2 food pellets, are specified. The number of trials for each CS per session is in parentheses.

Table 2. Summary of design for Experiment 2.

Conditioning	Extinction
Sessions 1-20	Sessions 21-35
Group AX-partial	
A 30 s, 88% (7+ & 1-)	A 17 s, 0% (12-)
AX 10 s, 29% (7+ & 17-)	AX 17 s, 0% (12-)
C 17 s, 100% (14+)	C 17 s, 100% (24+)
Group AX-consistent	
A 10 s, 29% (7+ & 17-)	A 17 s, 0% (12-)
AX 30 s, 88% (7+ & 7-)	AX 17 s, 0% (12-)
C 17 s, 100% (14+)	C 17 s, 100% (24+)

Note: The mean duration (in seconds) of each configuration, and the percentage of trials reinforced with 1 food pellet, are specified. The numbers of reinforced (+) and non-reinforced (–) trials per session for each configuration are shown in parentheses.

Table 3. Summary of design for Experiment 3

Conditioning	Test for inhibition by X	Extinction
Sessions 1-20	Sessions 21-28	Sessions 29-40
Group AX-partial		
A 30 s, 100% (8+)	A 30 s, 100% (8+)	
AX 10 s, 33% (8+ & 16-)	AX 10 s, 33% (8+ & 16-)	
B 10 s, 100% (8+)	B 10 s, 100% (8+)	B 10 s, 0% (1-)
	BX 20 s, 0% (1-)	BX 10 s, 0% (12-)
C 10 s, 100% (8+)	C 10 s, 100% (8+)	C 10 s, 100% (24+)
Group AX-consistent		
A 10 s, 33% (8+ & 16-)	A 10 s, 33% (8+ & 16-)	
AX 30 s, 100% (8+)	AX 30 s, 100% (8+)	
B 10 s, 100% (8+)	B 10 s, 100% (8+)	B 10 s, 0% (1-)
	BX 20 s, 0% (1-)	BX 10 s, 0% (12-)
C 10 s, 100% (8+)	C 10 s, 100% (8+)	C 10 s, 100% (24+)

Note: The mean duration (in seconds) of each configuration, and the percentage of trials reinforced with 1 food pellet, are specified. The numbers of reinforced (+) and non-reinforced (–) trials per

Figure 1. Mean response rates to CSs A and B, and during the pre-CS interval, during the conditioning phase (Sessions 1 to 20) and the extinction phased (Sessions 21-30) for both groups in Experiment 1a (top) and Experiment 1b (bottom).

Figure 3. Mean response rates to CSs A and B, and during the pre-CS interval, during the conditioning phase (Sessions 1 to 20) and the extinction phased (Sessions 21-30) for both groups in Experiment 2. The data from both groups are plotted together on the right, but for clarity are shown separately in the left and middle plots.

Figure 2. Results of an unpublished experiment. Rats were trained with 4 CSs. One CS was presented for a fixed duration of 60 s and was reinforced at a variable time (VT) with a mean of 30 s after CS onset. The other 3 CSs all had variable durations of either 30 s (VD30) or 60 s (VD60) and were reinforced at CS offset. One of the VD30 CSs (PRf) was reinforced on 50% of trials.

Figure 4. The primary results of Experiment 3. **Left:** Mean response rates during probe trials with CS B and with the compound BX, averaged across Sessions 21-28, for both groups. **Right:** Mean response rates to B and BX during the extinction phase (Sessions 29 to 40) and on the final condition day (Session 28) for both groups.