The Modified Law of Effect Explains the Partial Reinforcement Extinction Effect

Benjamin M. Seitz¹, Alexandra Stolyarova¹, & Aaron P. Blaisdell¹

¹University of California, Los Angeles

Submitted: September 2020
Correspondence to: bseitz1@ucla.edu

Abstract

Thorndike's Law of Effect provides a framework for understanding the selection of behaviors given specific environmental reward contingencies. Though a highly influential model, especially given its resurgence in popularity to understand habitual behaviors, it fails to predict several well-documented behavioral phenomena and incorrectly views extinction as the unlearning of a previously acquired association. Blaisdell, Stolyarova, & Stahlman (2016) proposed modifications to Thorndike's original law that address these issues and greatly increases the model's explanatory power. This modified Law of Effect (MLOE) also provides a testable account of the Partial Reinforcement Extinction Effect (PREE). The PREE is the paradoxical finding of more rapid extinction to a continuously reinforced cue than to a partially reinforced cue, and has challenged many theoretical accounts of learning. Simulations of the MLOE confirm these predictions. Two experimental paradigms, one using pigeons and the other using humans, show support for the Modified Law of Effect's explanation of the PREE.

Thorndike's *Law of Effect* (LOE) was the first formal model to address the selection of behavior over time (Thorndike, 1911). The law is computationally simple. An association between a stimulus (S) and subsequent response (R) will be strengthened if followed by a satisfying outcome (O+) such as a food reward or other forms of reinforcement. This S-R association will be weakened if followed by a dissatisfying outcome (O-) such as nonreinforcement or punishment. What serves as an O+ or O- depends on many factors, including species, current motivational state, and an individual's developmental history. By analogy to the evolutionary process of natural selection which filters out disadvantaged phenotypes in favor of those more fit to the current environment, the LOE is a simple trial and error (i.e. feedback) process that selects behaviors based on the valence of the outcome and their fit (aka belongingness) to a given stimulus, context, or situation. The LOE fell out of favor during the *cognitive revolution* but has recently seen a resurgence in attention for its role in habitual behavior that appears to be unaffected by a representation of the outcome (Balleine & Dickinson, 1998; Wood & Neal, 2007), and has been incorporated as an important learning process in neural models of behavior (Athalye, Santos, Carmena, & Costa, 2018; Balleine & O'Doherty, 2010), as well as contemporary forays into modeling of reinforcement learning (e.g., Niv & Schoenbaum, 2008; Sutton & Barto, 1998).

One tenet of the LOE that has not received much support is the weakening of the S-R association through nonreinforcement or punishment. While responding to a particular stimulus decreases during extinction, this decrease in performance does not necessarily reflect a decrease in the learned association. Rather, the last century of empirical work strongly supports the view that extinction results in the formation of a separate inhibitory association that interferes with or inhibits excitatory behavior (Bouton, 2004; Colwill, 1991; Rescorla, 1993;

Todd, Vurbic, & Bouton, 2014). Similarly, the LOE fails to account for the well-documented phenomenon whereby extinction results in the increase of behavioral variability (Antonitis, 1951; Eckerman & Lanson, 1969; Herrick & Bromberger, 1965; Neuringer, Kornell, & Olufs, 2001). The LOE has difficulty explaining this phenomenon because simply nonreinforcing a particular response should have no impact on the associative strength of other competing responses. Similar logic can be applied to the resurgence phenomenon. In resurgence, a target response is reinforced, and then extinguished as a new alternative response is reinforced. Following extinction of the alternative response, the target response increases in frequency despite being nonrewarded (Epstein & Skinner, 1980; Epstein, 1985; Winterbauer & Bouton, 2010). Again, according to the original LOE, extinction of one response should have no effect on the strength of a competing S-R association.

The Modified Law of Effect

Blaisdell, Stolyarova, and Stahlman (2016) proposed the *Modified Law of Effect* (MLOE) which expands the LOE's explanatory power (Figure 1). The first principle of the MLOE remains the same as the original model: An R followed by an O+ will strengthen the excitatory S-R association. The second principle of the MLOE is that an R followed by an O- results in the strengthening of a separate inhibitory S-R association, but does not weaken the excitatory S-R association. The excitatory and inhibitory S-R associations develop in parallel between any S and R that occur together (i.e., a response, R, made in the presence of a stimulus, S). Note that, according to the MLOE, excitatory and inhibitory S-R associations can increase in strength, but not decrease in strength. The third principle of the MLOE involves lateral inhibition (*aka* mutual inhibition), and states that all possible responses (Rs) that can occur in the presence of an S share mutual, fixed (i.e., non-plastic) inhibitory connections to each other. Although mutual inhibitory connections neither strengthen nor weaken as a result of experience, their impact on responding changes as a result of changes in net activation of the target response from

changes in the S-R associations. For example, rewarding (with O+) S-R1 increases activation of R1, which in turn increases R1's mutual inhibitory output to all other competing responses (R_i). This reduces the likelihood of the occurrence of competing responses. Alternatively, if S-R1 is followed by no reward or punishment (O-), the S-R1 inhibitory association increases, thereby reducing activation of R1 and reducing R1's inhibitory output to R_i —with the result that all other Rs increase in likelihood of occurrence (up to their baseline or post-training level). It is the inclusion of the mechanism of lateral inhibition that allows the MLOE to explain the relationship between reinforcement and response variability (discussed in more detail in Blaisdell et al.). The response rule is that the probability of any given R is a direct reflection of the net strengths of its excitatory and inhibitory S-R associations, and the net lateral inhibition received from all other Rs. We note that our operationalization of a response is intentionally vague because no two responses are likely to ever be truly identical. Nevertheless, for the sake of making specific empirical predictions, for our experimental setup we operationally define the target response as the pecking or pressing of a visual key presented on the computer display (see below), and the set of all non-target responses as any response in the operant chamber other than key presses (or pecks). Thus, in the following studies, we strictly measure pecking (or pressing) a key as the target response of interest. The MLOE is theoretically logical and is based on well-established principles of associative learning and behavior modification. Furthermore, it reflects a modern understanding of extinction, including how response variability increases following nonreinforcement and with reduced probability of reward (Stahlman & Blaisdell, 2011; Stahlman, Roberts, & Blaisdell, 2010; Stahlman, Young, & Blaisdell, 2010). Nevertheless, specific unique predictions of the MLOE have yet to be put to rigorous empirical testing.

The Partial Reinforcement Extinction Effect (PREE)

One such test involves the MLOE's explanation of the PREE. The PREE is the finding of more rapid extinction of an instrumental or Pavlovian response that had previously been

reinforced on a continuous schedule compared to a response that had only been partially, or intermittently reinforced or rewarded (Humphreys, 1939). The PREE has been a challenge for many learning theories (e.g., Rescorla & Wagner, 1972) that assume the strength of conditioned responding to be a product of the associative strength between stimulus and response, or cue and outcome.

The two most popular accounts of the PREE have been Amsel's (1967) frustration theory and Capaldi's (1966) sequential hypothesis. According to Amsel's frustration theory, frustration is elicited by nonreinforcement on nonreinforced trials involving a partially reinforced cue or response. This frustration remains to some degree on the following trial. Sometimes a rewarded trial follows a nonrewarded trial, and thus an association develops between the internal state of frustration and reward. Thus, frustration itself becomes an S+ that motivates the conditioned response. Frustration never develops into an S+ for responding in subjects trained only with a continuously reinforced cue or response. Thus, when extinction treatment begins, and every trial evokes frustration due to nonreward, responding continues to be motivated to a partially reinforced cue or response, but not to a continuously reinforced cue or response. This extended responding to the partially reinforced cue but not the continuously reinforced cue results in the PREE.

Capaldi's sequential hypothesis is quite similar, except with memory for nonreward becoming an S+ to motivate conditioned responding. For example, if during a partial-reinforcement schedule, a nonrewarded trial is followed by a rewarded trial, and the subject has a memory of nonreward (N) from the previous trial when the reward is delivered on the subsequent trial, N can acquire properties of an S+. Since all continuously reinforced trials are rewarded, the subject never experiences N, nor develops any association between N and reward. Thus, during extinction, N continues to motivate responding in the partially reinforced condition, but not in the continuously reinforced condition, leading to the PREE.

While both theories have garnered much theoretical support, an experiment by Rescorla (1999) provides evidence that PREE might not be explained by either account. Rescorla used a completely within-subject procedure in which trials with separate partially reinforced and continuously reinforced cues were interleaved within a session. During extinction, the partially reinforced cue extinguished more slowly than did the continuously reinforced cue, thereby demonstrating the PREE. This finding challenges both theoretical accounts of the PREE discussed above because both frustration and memory for nonreward should motivate responding on all trials during extinction, both the partially reinforced cue or response AND the continuously reinforced cue or response. This is because each extinction trial, whether for the continuously reinforced cue or partially reinforced cue, ends in nonreinforcement. Thus, frustration (Amsel) or memory for nonreward (Capaldi) should be elicited on each extinction trial, which in turn should contribute motivation to respond equally to continuously reinforced and partially reinforced cues.

The reviewed theoretical explanations of the PREE rely on trial-based accounts of learning. Gallistel (2012) provided a time-based (rather than trial-based) account of the PREE. In this model (see Gallistel & Gibbon, 2000) the subject learns about the rate at which a CS or discriminative stimulus is reinforced. Thus, the subject computes a lower rate of reinforcement on a partially reinforced than continuously reinforced schedule. Extinction involves a change in the rate of reinforcement from higher to much lower. The PREE occurs, according to this theory, because it takes longer to detect the change from sparse reinforcement on the partial-reinforcement schedule to no reinforcement during extinction compared to the change from dense reinforcement on the continuous reinforcement schedule to no reinforcement during extinction. Recent attempts to test this account by equating the rate of reinforcement for continuously reinforced and partially reinforced cues still result in the PREE, thereby challenging the rate estimation theory account (Bouton, Woods, & Todd, 2014; Chan & Harris, 2017).

The MLOE provides an alternative trial-based account of the PREE that does not suffer the same weaknesses as frustration or memory sequence accounts, because it predicts the PREE for both between and within-subject procedures. According to the MLOE, a continuously reinforced cue should acquire a strong excitatory S-R association but no inhibitory S-R association. A partially reinforced cue, however, should acquire excitatory and inhibitory S-R associations. The excitatory association strengthens when responses are followed by reward and the inhibitory association strengthens when the response is followed by nonreward. The change in associative strength has been shown to follow an error-reduction process. When a stimulus is followed by reward, the increment in excitatory associative strength is largest on the first rewarded trial, and progressively decreases as the association approaches asymptotic level supported by the reward. Likewise, the increment in inhibitory strength is largest on the first nonrewarded trial, and progressively decreases as the inhibitory association approaches asymptotic level supported by nonreward. These acquisition functions for excitatory and inhibitory associations have been best documented by Rescorla (2000; 2001). This understanding allows the MLOE to account for the PREE. During the extinction phase, both the continuously reinforced and partially reinforced cues are under an extinction schedule in which on every trial they are nonrewarded. The continuously reinforced cue has never previously been nonrewarded and thus its inhibitory S-R association starts at zero strength and should increase rapidly. The partially reinforced cue, however, has previously received a history of nonreward during acquisition, and thus should have a relatively well-established inhibitory S-R association. As a result, on each nonrewarded trial during extinction, its inhibitory S-R association should undergo much smaller increments. Thus, while according to the MLOE both continuously reinforced and partially reinforced cues undergo increments in inhibitory S-R strength during extinction, these negatively accelerating functions start at different levels of preexisting S-R inhibition, and thus change at different rates during extinction. Assuming adequate response

acquisition to continuously reinforced and partially reinforced cues, the more rapid increase in inhibitory S-R learning to the continuously reinforced than partially reinforced cue during extinction predicts more rapid extinction to the continuously reinforced than partially reinforced cue—that is, the PREE.

If the basis for the PREE is the different levels of S-R inhibition at the start of extinction treatment, then blocking the development of inhibitory S-R associations during acquisition of the partially reinforced cue should prevent the PREE. That is, blocking the development of an inhibitory S-R association to the partially reinforced cue should result in both the partially reinforced and continuously reinforced cues starting extinction with the same, zero strength inhibitory S-R association. As a result, they should both undergo similar rates of acquisition of the inhibitory S-R association, and thus the response should extinguish at the same rate to both cues. This prediction can be tested by preventing responses from being made to a partially reinforced cue during nonreinforced trials in the acquisition phase of training. This can be achieved by omitting presentation of the response key on nonreinforced trials, which according to the MLOE should prevent the formation of inhibitory S-R associations. Amsel and Capaldi's theories still predict the PREE, even when subjects are prevented from making the instrumental response on nonreinforced trials, because the partially reinforced cue is still partially reinforced in a Pavlovian manner. That is, on nonreinforced trials, the partially reinforced cue is present and nonreinforced. According to Amsel's theory, nonreinforcement would elicit frustration. Some nonreinforced trials are followed by reinforced trials, and thus the frustration produced by the preceding trial now has the opportunity to become associated with, or serve as, a discriminative cue for reinforcement on the subsequent reinforced trial. As a result, during extinction, frustration should slow down extinction of both partially reinforced and continuously reinforced cues using a within-subject procedure. An analogous prediction is made by Capaldi's sequential hypothesis in terms of memories of nonreward slowing down the rate of extinction on all trial

types during extinction. It is only the MLOE that predicts more rapid extinction specifically to the partially reinforced cue for which the instrumental key peck response was prevented during acquisition (i.e. no PREE), and not to the partially reinforced cue for which instrumental key pecking was allowed during acquisition, because no S-R inhibitory associations should form during acquisition to the former partially reinforced cue. We first present a simulation of the MLOE to formally verify its predictions regarding PREE and the prevention of PREE by removing the response key during acquisition treatment with a partially reinforced cue. Following the simulation, we present two experiments in which we tested the MLOE predictions using a within-subject procedure similar to that described by Rescorla (1999) in pigeons (Experiment 1) and a mixed design in humans (Experiment 2).

Simulation of the MLOE Account of PREE

The task implemented in the simulation was the same as the task used for pigeons as described for Experiment 1 below, with the same numbers of sessions and trials, and the same reinforcement contingencies. In short, the model simulated response frequencies on a task for pigeons that involved key responses during the separate presentations of four different discriminative stimuli, A, B, X, and Y, which were colored rectangles presented on a computer monitor. Each trial involved the presentation of one of the discriminative stimuli, with an equal number of each trial type per session. During acquisition, key pecks during A and X were continuously reinforced, while key pecks during B and Y were reinforced on 25% of trials. During extinction, A continued to be reinforced at 100%, whereas X, B, and Y were nonreinforced on every trial (see *Table 1*). A critical manipulation is that there was no response key available on nonreinforced trials with Cue Y during acquisition, though the response key was available on Cue Y trials during extinction and on reinforced trials during acquisition (see *Figure 2*). This manipulation prevents the formation of an inhibitory S-R association to Cue Y during acquisition, but allows it to form and strengthen during extinction.

We simulated a total of 1000 agents performing the task. Each agent formed a set of excitatory and inhibitory S-R associations for the four stimuli used in the study. One of the responses contingent on each stimulus was the target action (meant to represent pecking the key). The other response represented all alternative actions; these actions were modeled as a single response because reinforcement contingencies for all behaviors other than pecking the key are the same (i.e., no reward delivery). At the beginning of the simulation each response had the same excitatory strength, $Q(R_{i,ex}|S_j)$, and the same inhibitory strength $Q(R_{i,in}|S_j)$. Here R_i represents response i and S_j represents stimulus j . All responses also had mutual inhibitory connections that were proportional to the net strength of the S-R associations:

$$\omega_{m_in} \cdot (Q(R_{i,ex}|S_j) - Q(R_{i,in}|S_j)), \text{ where } \omega_{m_in} \text{ is the strength of mutual inhibition.}$$

Therefore, the total strength of response i when presented with stimulus j , $Q(R_i|S_j)$, is calculated as:

$$Q(R_i|S_j) = Q(R_{i,ex}|S_j) - Q(R_{i,in}|S_j) - \omega_{m_in} \cdot (Q(R_{\sim i,ex}|S_j) - Q(R_{\sim i,in}|S_j)), .$$

Because each trial lasted 20 seconds for pigeons thereby allowing for multiple actions, for the agents each trial was also broken down into several bins (set to 4) and a response was selected in each bin. The probability of the target response (key peck) in each bin was equal to $Q(R_i|S_j)$: a single uniformly distributed random number in the interval (0,1) was drawn in each bin and if $Q(R_i|S_j)$ exceeded that number the response was made.

The strength of excitatory and inhibitory associations was updated after the response was made and an outcome was observed. A response followed by a reward strengthens the excitatory S-R association. In contrast, a response followed by no reward leads to strengthening of a separate inhibitory S-R association. The rules for updating the excitatory and inhibitory associations were as follows:

1. When a response is made and a reward is delivered, the strength of the excitatory association for response i when presented with stimulus j , $Q(R_{i,ex}|S_j)$, is increased according to:

$$Q(R_{i,ex}|S_j) \leftarrow Q(R_{i,ex}|S_j) + \alpha_r \cdot (1 - Q(R_{i,ex}|S_j)),$$

where α_r is the update rate for excitatory associations following reinforcement.

The maximum possible value of $Q(R_{i,ex}|S_j)$ is 1 (because the rewarding event is represented as 1) and the change in the strength of the excitatory association is greater early during conditioning. As conditioning progresses $Q(R_{i,ex}|S_j)$ approaches its maximum value and the difference, $1 - Q(R_{i,ex}|S_j)$, that drives learning grows progressively smaller.

2. When a response is made and a reward is *not* delivered, the strength of the inhibitory association of response i when presented with stimulus j , $Q(R_{i,in}|S_j)$, is increased according to:

$$Q(R_{i,in}|S_j) \leftarrow Q(R_{i,in}|S_j) + \alpha_{nr} \cdot (1 - Q(R_{i,in}|S_j)).$$

Here, α_{nr} is the update rate for inhibitory associations following nonreinforcement.

The response data from this simulation were summarized in the same way as the pigeon behavioral data and the results are presented in Figure 2, separately for acquisition and extinction treatments. The following parameter values were used in the simulation presented in Figure 2: $\omega_{m_in} = 0.2$, $\alpha_r = 0.02$, $\alpha_{nr} = 0.0015$. The strengths of excitatory and inhibitory associations at the beginning of the simulation for each agent were set to 0. The higher update rate for the strength of the excitatory, compared to inhibitory, S-R associations (α_r vs α_{nr}) is needed to compensate for the lower frequency of rewarded, compared to non-rewarded, trials. If the update rates are set to the same value, inhibition quickly overpowers excitation with a 75%

non-reinforced schedule (because there are 5 bins in each trial, non-reinforcement can occur 15 times more frequently than reinforcement). The difference between these parameter selections rates of excitatory versus inhibitory learning are also consistent with the huge empirical literature showing that excitatory responding develops faster than inhibitory responding (both for Pavlovian conditioned inhibition and for discriminative operant responding), and for the ability to establish excitatory responding on lean schedules of reinforcement, such as 25% (or lower).

The simulation was implemented in MATLAB (MathWorks R2017a). Figure 3 shows the results of the simulation for both acquisition treatment (upper panel) and extinction treatment (lower panel). Continuously reinforced Cues A and X approach asymptotic levels of responding sooner (around Blocks 2 or 3) than partially reinforced Cue Y (around Blocks 5 or 6), for which the target response was not available on nonreinforced trials. Partially reinforced Cue Y did eventually reach similar asymptotic level of responding as did the continuously reinforced cues. The slower rate of acquisition to Cue Y is due to having fewer reinforced trials per session. Partially reinforced Cue B also exhibited a slower rate of acquisition (around Blocks 5 or 6), but unlike Cue Y, Cue B also reached a lower asymptotic level of responding. The lower asymptote is due to the acquisition of an inhibitory S-R association that competes with the excitatory S-R association to produce a net level of responding lower than observed to the continuously reinforced cues for which no inhibitory S-R association was acquired.

The lower panel of Figure 3 depicts the proportion of simulated key pecks during extinction for the presentation of each cue as a function of the number of pecks made on the first day of extinction. This mirrors the normalization of extinction data in the experiments below. Cue A continues to be reinforced during extinction, and thus maintains 100% of responding as observed on Day 1 of extinction treatment. Continuously reinforced Cue X shows a more rapid rate of extinction than does partially reinforced Cue B. Thus, the MLOE accurately simulates the PREE. More importantly, responding to partially reinforced Cue Y extinguishes as rapidly as that

to continuously reinforced Cue X, and faster than to partially reinforced Cue B. This provides a demonstration of the MLOE's prediction that preventing key responses on nonreinforced trials of Cue Y during acquisition, which should prevent the formation of an inhibitory S-R association for Cue Y, would prevent the PREE during extinction. This is due to both Cues X and Y entering into the extinction treatment with a zero strength inhibitory S-R association for the target response, and thus, the inhibitory S-R association should undergo rapid and large increases at equal rates for each cue.

Now that we've formally validated the predictions of the MLOE for the PREE and for the prevention of PREE through the omission of the target response on nonreinforced trials with Cue Y, we turn to the empirical investigation of these predictions.

Experiment 1

As in the simulated experiment above, in Experiment 1, pigeons were trained to peck a response key during the separate presentations of four different discriminative stimuli, A, B, X, and Y, which were colored rectangles presented on a computer monitor. Key pecks during A and X were continuously reinforced, while key pecks during B and Y were reinforced on 25% of trials (see *Table 1*). The MLOE's account of the PREE was tested by omitting the response key on nonreinforced trials with Y. Omitting the response key should prevent the formation of inhibitory S-R associations to Cue Y. As a result, compared to extinction of continuously reinforced Cue X, we predicted similar, rapid extinction to Y, but slower extinction to B.

Participants: Two male and two female homing pigeons (*Columba livia*) from Double T Farms served as subjects[‡]. Subjects were individually housed in steel home cages with metal wire mesh floors in a vivarium and they were maintained at 80% of their free-feeding weight. Testing occurred at approximately the midpoint of the light portion of the 12-hour light-dark cycle.

Materials: Testing was conducted in a flat-black Plexiglas chamber (38 x 36 x 38 cm). Stimuli were generated on a monitor (NEC MultiSync LCD1550M). Pecks to the monitor were detected

by an infrared touchscreen (Carroll Touch, Elotouch Systems, Fremont, CA) mounted on the monitor. A customized food hopper was located in the center of the front panel, its access hole flush with the chamber floor, and could deliver mixed grain. All experimental procedures were programmed using PsychoPy Software (Peirce et al., 2019). Four colored (red, blue, green, yellow) rectangles (20 x 15 cm) served as discriminative stimuli (DS) and are subsequently referred to as A, B, X, and Y. Color assignments for A, B, X, and Y were counterbalanced across birds. Below the DS, a 2.5-cm diameter white circle centered on the lower fifth of the screen served as the response key. See Figure 3a for an example screenshot of a trial display.

Procedure

A summary of the full reinforcement schedule for each cue is provided in *Table 1*.

Acquisition

A mixed autoshaping/operant shaping procedure to the key began by displaying the key below one of four colored rectangles (DS). Trials lasted 20s and multiple rewards could be obtained on each rewarded trial. Each 60-min session consisted of the presentation of 25 trials of each stimulus (A, B, X, and Y). The first presentation of each stimulus was always rewarded. For the remaining 24 trials, Cues A and X were reinforced on each trial (continuously reinforced), while Cues B and Y were only reinforced on 25% of trials (partially reinforced). All reinforced trials ended with a non-contingent delivery of the reward (i.e. autoshaping). Stimulus presentations were followed by a variable (ITI) with a mean time of 10s, where the screen was grey and the key was present. On rewarded trials, each time the key was pecked, the subject received a reward, which consisted of a 3-s delivery of grain from the hopper. While the key was present on all trials with Cues A, X, B, and during the ITI, the key was not present on nonreinforced trials with Cue Y (see Figure 3b). Nevertheless, the key was present for the 25% of Y trials that were reinforced, thus allowing excitatory but not inhibitory S-R associations to form. We also recorded pecks on nonreinforced Y trials from the screen location where the

response key was otherwise present on other trials, which allowed us to continue to document the subject's behavior during these trials. Sessions were conducted five days per week and acquisition lasted a total of 12 daily sessions.

Extinction

After twelve acquisition sessions, peck rates to A, X, B, and Y had reached asymptotic levels (visually assessed, see Figure 4a), and birds were placed on an extinction schedule, in which key pecks to Cues X, B, and Y were never reinforced. Each cue was presented 24 times per session during extinction. Key pecks during presentations of Cue A continued to be reinforced on 100% of trials as was what had occurred during acquisition. The key was present on all trials during extinction, including trials with Cue Y. The study concluded after 6 days of extinction training as peck rates had reached asymptotically low levels (visually assessed, see Figure 4b).

Results and Discussion

The mean number of pecks made to the key on rewarded trials per acquisition session is averaged across all 4 birds and presented in 2-session blocks (Figure 4a, see figure S1a for responding on nonreinforced trials). A 4 (Cue) by 6 (Block) Repeated Measures ANOVA did not reveal a significant main effect of block, $F(5,15) = 2.45$, $p = 0.082$, $\eta^2 = 0.450$, cue, $F(3, 9) = 1.27$, $p = 0.34$, $\eta^2 = 0.298$, nor interaction between cue and block, $F(15,45) = 1.08$, $p = 0.41$, $\eta^2 = 0.264$. While there was not a statistically significant effect of block on responding, visually (see Figure 4a) it is clear the birds increased their responding with continued training—from a mean of 0.5 pecks per 20-s trial in the beginning of training to a mean of 3 pecks trial at the end of training. Thus, birds appeared to learn to peck the key during the presentation of each cue at similar rates, though there was great variability in the number of pecks made during each cue by each bird, though this conclusion is not statistically supported. Also, peck rates to the key were highest during nonreinforced B trials, (see S1a) which should result in the acquisition of a strong

inhibitory S-R association.

The extinction phase lasted 6 days and the extinction data are presented in Figure 4b. On one hand, on the first day of extinction, responding to partially reinforced Cue B remained similar to responding at the end of the acquisition phase—which is typical in the transition from acquisition to extinction. On the other hand, continuously reinforced Cue X and partially reinforced Cue Y elicited massive extinction bursts on the first day of extinction (note continuously reinforced Cue A remained reinforced to motivate responding during the extinction sessions). Extinction bursts are not uncommon when transitioning from continuous reinforcement to an extinction schedule (Katz, 2018), but one must ask why an extinction burst was observed to partially reinforced Cue Y. A frustration account is unsatisfying, because frustration putatively would have been experienced during the 75% of Y trials that were nonreinforced during training. By contrast, the MLOE readily predicts that the lack of inhibitory S-R associations for Cues X and Y should result in an extinction burst observed to each cue at the beginning of extinction treatment. Because the extinction burst to Cues X and Y result in much higher absolute levels of responding than observed to Cues A and B, we next analyzed the proportion of pecks made to each cue during sessions 2-6 of extinction relative to the number of pecks made to each cue on the first day of extinction (Figure 4c). This approach mirrors that of Rescorla (1999, Exp 3) who examined performance during extinction as a percentage of the final asymptotic level for the same reason—that differences in absolute levels of response complicate interpretation of extinction results. A 4 (Cue) by 5 (Session) Repeated Measures ANOVA was used to analyze the normalized data, eliminating the first day due to no variability between cues. This revealed a main effect of Cue, $F(3,9) = 21.44$, $p < 0.001$, $\eta^2 = 0.877$, and Session, $F(4, 12) = 12.07$, $p < 0.001$, $\eta^2 = 0.801$, and Cue by Session interaction, $F(12,36) = 3.18$, $p = 0.004$, $\eta^2 = 0.515$. A planned comparison between partially reinforced Cue B and continuously reinforced Cue X revealed a significant difference, $t(3) = 20.42$, $p < 0.001$.

This difference in extinction rates to a partially reinforced cue and a continuously reinforced cue demonstrates the PREE. In line with MLOEs account of the PREE, a planned comparison between continuously reinforced Cue X and partially reinforced Cue Y revealed no difference, $t(3) = 1.65$, $p = 0.20$. Furthermore, a planned comparison revealed more responding to partially reinforced Cue B than to partially reinforced Cue Y, $t(3) = 4.90$, $p = 0.016$. Thus, partially reinforced Cue Y did not demonstrate the resistance to extinction that was observed for partially reinforced Cue B and that is typically observed in the PREE. This suggests the acquiring an inhibitory S-R association is what mediates the PREE, and that preventing a key response on nonreinforced trials prevented the formation of this association.

Is it possible that on nonreinforced Y trials during acquisition, pigeons simply learned not to respond to Cue Y, thus making this an easy pattern to continue during extinction? We find this alternative explanation unsatisfactory. First, on reinforced trials when the key was present, responding to Cue Y was equally as robust as responding to the key while other cues were presented. Also, by the end of acquisition, some birds made pecks on nonreinforced Y trials to the location where the key would have been on reinforced trials (see Y_null in S1a), which suggests they did not learn a response rule of not making responses during the nonreinforced Y trials. With that said, demonstrating a similar pattern of responding in a similar procedure but with a different animal model would further strengthen the MLOE's explanation of the PREE.

Experiment 2

The results from Experiment 1 are congruent with the MLOE's account of the PREE. By preventing the pigeon from responding to the key on nonreinforced trials for a partially reinforced cue, we were able to attenuate resistance to extinction—the defining feature of the PREE. Experiment 2 tests the generality of the MLOE account of PREE using a similar but adapted procedure in humans.

Participants: Forty UCLA undergraduates (29 female) participated in this study for course credit. Twenty participants per condition afforded us over 80% power to observe a medium-sized ($f = 0.25$) within-subject PREE based on our pilot testing.

Materials: The same touchscreen monitor as was used in Experiment 1 was also utilized in Experiment 2, but the plexiglass chamber was removed and the apparatus was moved to a small human testing room. Points were utilized as a rewarding outcome, and sufficient points earned participants their choice of candy from an assorted bag. Trials lasted 7 s and points could be awarded with every key press on rewarded trials. Additionally, to facilitate participant engagement, there was a 3-s delay between presentation of the cue and presentation of the key. Thus, participants had only 4 s to make responses on any given trial.[‡]

Procedure: Table 2 summarizes the full reinforcement schedule of each cue. Participants received acquisition and extinction training in a single session. To facilitate this, participants were randomly assigned to one of two conditions ($n = 20$). Participants in *Condition AB* were trained on continuously reinforced Cue A and partially reinforced Cue B (25% rewarded), as well as Cue C, which was not reinforced during the acquisition phase but was continuously reinforced during the extinction phase. This served to make the transition from acquisition to extinction more salient and to still allow participants to obtain some reward during the extinction of Cues A and B. Participants in *Condition XY* were trained on continuously reinforced Cue X and partially reinforced Cue Y (25% rewarded), and Cue C. As with Experiment 1, the key was not available during nonreinforced trials of acquisition for Cue Y. Cue C was similarly nonrewarded during acquisition and then continuously rewarded during extinction. The noncontingent reward at the end of the trial (i.e. autoshaping) that was used in Experiment 1 was removed and points could only be obtained by pressing the button. We found through pilot testing, that this omission helped facilitate the PREE in a single session in humans^{‡‡}. Thus, we hypothesized that with one group of 20 participants on our established procedure we would

observe the PREE but for the other group who were incapable of making a response on nonreinforced trials to the partially reinforced cue—they would not show the effect. That eliminating the response key would eliminate the effect provides evidence that the response key—and thus learning the inhibitory association—is necessary for the PREE.

Participants were told that the aim of the study was to learn when to correctly press a button to obtain points. To incentivize this behavior, they were told that if they obtained more points than the average participant, they would receive a piece of candy of their choosing. On rewarded trials, pressing the response key immediately resulted in being awarded a point, and participant's point total was displayed and live updated on the top of the screen. Above the participant's respective point total was text that showed the average participant score of 1500 points. In total, participants were presented with 17 trials of each cue during acquisition, and 24 presentations of each cue during extinction. The first presentation of Cues A, B, X, and Y was rewarded for all participants during acquisition and the remaining 16 were always reinforced or reinforced 25% of the time (i.e. on 4 out of the 16 trials).

Results and Discussion: Responses to Cue C were not relevant and thus were eliminated from analysis (but see Figure S2 for those data). The acquisition data are grouped into blocks of 4 consecutive trials, except for Block 1 which only includes responding for the first presentation of each cue. For example, Block 2 includes the mean value of responses made on trials 2-5 for each cue. These data for each condition are displayed in Figures 5a and 5b respectively and again reflect the mean number of key presses made on rewarded trials to each cue (for responding on nonreinforced trials, see Figure S1b/c). A 5 (Block) by 2 (Cue) repeated measures ANOVA was conducted for the acquisition phase for each condition. For Condition AB, there was a main effect of Block, $F(4,76) = 26.51$, $p < 0.001$, $\eta^2 = 0.374$, but not of Cue, $F(1,19) = 3.53$, $p = 0.076$, $\eta^2 = 0.028$, nor an interaction between Block and Cue, $F(4,76) < 1.0$. For Condition XY, there was a main effect of Block, $F(4,76) = 38.38$, $p < 0.001$, $\eta^2 = 0.533$, but

not of Cue, $F(1,19) = 2.78$, $p = 0.11$, $\eta^2 = 0.006$, nor an interaction between Block and Cue, $F(4,76) < 1.0$. Thus, for both conditions, participants learned to respond to rewarded trials of both cues at similar rates.

The 24 extinction trials were grouped into 4 trial blocks and are presented in Figures 5c and 5d. One notable difference between these extinction data and those from the pigeons in Experiment 1 is the lack of an extinction burst during the first block of trials during extinction. In fact, there was a large attenuation in overall responding at the outset of extinction compared to the end of acquisition. A number of factors might explain this. First, in Experiment 1, each bird received a total of 300 acquisition trials per cue whereas in Experiment 2 each participant only received 17 acquisition trials of each cue. Further, pigeons were trained on a total of 4 different cues while humans were only trained on 3. Last, the initial attenuated responding in extinction in humans compared to the extinction burst in pigeons may reflect species differences in inhibiting responses that lead to nonreward—though this may also be a function of the extended training in pigeons. Thus, the lack of an extinction burst in these extinction data, and similar rates of responding across all extinction cues on the first four trials (block 1) of extinction obviates the need to normalize the extinction data. Separate 6 (Block) by 2 (Cue) repeated measures ANOVA's were used to analyze the extinction data. For Condition AB, there was a main effect of Block, $F(5,95) = 3.00$, $p = 0.015$, $\eta^2 = 0.108$. Surprisingly, there was a no effect of Cue, $F(1,19) = 3.43$, $p = 0.079$, $\eta^2 = 0.010$, however there was a significant interaction between Block and Cue, $F(5,95) = 3.37$, $p = 0.008$, $\eta^2 = 0.021$. Indeed, there was a simple main effect of Cue on Block 3, $F(1) = 5.43$, $p = 0.031$, Block 5, $F(1) = 8.01$, $p = 0.011$, and Block 6 $F(1) = 5.66$, $p = 0.028$, but not on Block 1, $F(1) = 2.05$, $p = 0.168$, Block 2, $F(1) = 1.26$, $p = 0.168$, or Block 4, $F(1) = 1.79$, $p = 0.197$. Thus, we appear to have observed a small PREE for Condition AB. Meanwhile, for Condition XY there was a main effect of Block, $F(5,95) = 15.76$, $p < 0.001$, $\eta^2 = 0.395$, but no effect of Cue $F(1,19) = 1.564$, $p = 0.23$, $\eta^2 = 0.002$. There was also a significant

interaction between Block and Cue, $F(5,95) = 4.44$, $p = 0.001$, $\eta^2 = 0.020$, but there was a only a significant simple main effect of Cue on Block 1, $F(1) = 9.018$, $p = 0.007$, but not in Blocks 2-6 (all $F_s < 1.0$). Thus we appear to have observed a PREE for Condition AB but not Condition XY. Thus, preventing responses on nonrewarded trials appears to attenuate the PREE in humans as it did in pigeons. These data obtained from experiments with two animal models and the model simulations, and strong theoretical basis, are suggestive of the MLOE's account of the PREE.

General Discussion

The Partial Reinforcement Extinction Effect is one of the most robust and well-documented phenomena in associative learning. Nevertheless, the PREE appears paradoxical to most accounts of learning that presume responding to reflect the acquired associative strength between stimulus and response (or outcome). A number of theoretical accounts have been proposed to explain this effect, most notably Amsel's (1967) frustration account and Capaldi's (1966) sequential hypothesis. Nevertheless, neither account can explain how the PREE can be observed using a within-subjects paradigm (Rescorla, 1999). The MLOE, proposed by Blaisdell et al. (2016), accounts for the PREE in a more parsimonious fashion than that proposed by Amsel or Capaldi. Because the model posits the development of parallel excitatory and inhibitory S-R associations, partially reinforced cues acquire inhibitory S-R associations during acquisition, while continuously reinforced cues do not. During extinction, the rapid increase in the inhibitory S-R association for the continuously reinforced cue results in faster extinction compared to the partially reinforced cue which undergoes a comparatively slow increase in S-R inhibition due to already having acquired a strong inhibitory S-R association (cf. Rescorla, 2001). Therefore, according to the MLOE, the PREE can be diminished by preventing the participant from making the instrumental response on nonreinforced trials. This analysis was formally shown in simulations of the MLOE. In two experiments with two different animal

models, we found evidence for the MLOE predictions by omitting the response key on nonreinforced trials for a specific discriminative stimulus (Cue Y) that was only reinforced on 25% of trials. Both Amsel's and Capaldi's theories of the PREE do not predict that omitting the response key on nonreinforced trials will attenuate the PREE, because the cue remains partially reinforced which should cause frustration or a memory trace of nonreinforcement followed by reinforcement, and these processes should exert their effects on all cue types during extinction. In both experiments, we found no evidence of resistance to extinction to partially reinforced Cue Y which supports the MLOE's account of the PREE.

One potential limitation of this study is that during acquisition, nonreinforced trials of partially reinforced Cue Y may have been perceived differently than reinforced trials of this cue. This could serve as a context change and these contextual changes could be driving our effects. That said, Boughner and Papini (2006) report that the PREE persists even with changes in context from acquisition to extinction, suggesting this effect might be somewhat immune to context effects. Moreover, the purpose of including the key during the ITI for Experiment 1, and delaying the onset of the key during each trial in Experiment 2, was to reduce the likelihood that the key would modulate instrumental behavior, and instead only serve as the target of instrumental behavior. For instance, in Experiment 2, reinforced Y trials and nonreinforced Y trials were indistinguishable from each other for the first 3 seconds. Lastly, while a context-specific account of these data might provide an alternative to the MLOE's account, this is true of many phenomena in associative learning. The possibility of virtually any environmental or internal state change to be viewed as distinct context makes testing the context-specific account a protean endeavor. At the very least, the MLOE's explanation of these data provides the most parsimony and a-priori predictive power, whereas context-modulation accounts are typically more easily constructed post-hoc.

While Thorndike's (1911) original LOE has been the predominant model of habitual

responding, the MLOE increases its explanatory power. The model more accurately addresses the field's current understanding of extinction, explains why response variability increases under extinction (Blaisdell et al., 2016), and also provides a simple and parsimonious account of the PREE. Accumulating further evidence in support of the MLOE may grant insight into the computational rules central to habit formation, addiction-like behaviors, general reinforcement learning, and the neural circuitry that govern these processes.

Open Science Statement: The data for Experiments 1 & 2 and the MATLAB code for the simulations can be found here <https://osf.io/7dfh8/>. The study was not preregistered.

Acknowledgements: The author's thank Christina Indudhara for her help with data collection and study design.

Footnotes:

[‡]Four pigeons were initially run because we had only four functioning touchscreen-operated operant chambers at the time. After their completion, we conducted a post-hoc power analysis on our results. The critical finding was a PREE for partially reinforced Cue B relative to continuously reinforced Cue X but not for partially reinforced Cue Y relative to Cue X. Thus, it was most important to be powered to detect a reliable PREE. A repeated measures ANOVA including the 5 sessions of extinction data and 2 cues (B and X) revealed the size of the effect of cue was $\eta^2 = 0.260$. The correlation between responding to Cue B and Cue X was $r = 0.98$, and a power analysis in G*power (Faul, Erdfelder, Lang, & Buchner, 2007), revealed that even with our four subjects, we had 100% power to detect an effect of this size.

^{‡‡}Much like the Law of Effect itself, the procedure used in Experiment 2 was very much the result of a trial and error process. We piloted many different procedures, utilizing different rewards, responses, and different methods of presenting stimuli, until we finally observed a reliable PREE. This explains why we used points as the reward, delayed the onset of the key presentation, and omitted the autoshaping component of the procedure.

References

- Amsel, A. (1967). Partial Reinforcement Effects On Vigor and Persistence: Advances in Frustration Theory Derived from a Variety of Within-Subjects Experiments. *Psychology of Learning and Motivation*, 1, 1–65. [https://doi.org/10.1016/S0079-7421\(08\)60512-5](https://doi.org/10.1016/S0079-7421(08)60512-5)
- Antonitis, J. J. (1951). Response variability in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology*, 42(4), 273–281. <https://doi.org/10.1037/h0060407>
- Athalye, V. R., Santos, F. J., Carmena, J. M., & Costa, R. M. (2018). Evidence for a neural law of effect. *Science (New York, N.Y.)*, 359(6379), 1024–1029. <https://doi.org/10.1126/science.aao6058>
- Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37(4–5), 407–419. [https://doi.org/10.1016/S0028-3908\(98\)00033-1](https://doi.org/10.1016/S0028-3908(98)00033-1)
- Balleine, B. W., & O'Doherty, J. P. (2010). Human and Rodent Homologies in Action Control: Corticostriatal Determinants of Goal-Directed and Habitual Action. *Neuropsychopharmacology*, 35(1), 48–69. <https://doi.org/10.1038/npp.2009.131>
- Blaisdell, A. P., Stolyarova, A., & Stahlman, W. D. (2016a). The Law of Expect or a Modified Law of Effect? Outcome expectation and variation in learned behavior. *Conductual*, 4(2), 61–90. Retrieved from <http://conductual.com/content/law-expect-or-modified-law-effect>
- Blaisdell, A. P., Stolyarova, A., & Stahlman, W. D. (2016b). The Law of Expect or a Modified Law of Effect? Outcome expectation and variation in learned behavior 1, 2. *Conductual*, 4(2), 61–90. <https://doi.org/ISSN:2340-0242>
- Boughner, R. L., & Papini, M. R. (2006). Survival of the partial reinforcement extinction effect after contextual shifts. *Learning and Motivation*, 37(4), 304–323. <https://doi.org/10.1016/j.lmot.2005.09.001>
- Bouton, M. E. (2004). Context and Behavioral Processes in Extinction. *Learning & Memory*, 11(5), 485–494. <https://doi.org/10.1101/LM.78804>
- 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. <https://doi.org/10.1016/j.beproc.2013.08.006>
- Chan, C. K. J., & Harris, J. A. (2017, August 1). Extinction of Pavlovian conditioning: The influence of trial number and reinforcement history. *Behavioural Processes*, Vol. 141, pp. 19–25. <https://doi.org/10.1016/j.beproc.2017.04.017>
- Colwill, R. M. (1991). Negative discriminative stimuli provide information about the identity of omitted response-contingent outcomes. *Animal Learning & Behavior*, 19(4), 326–336. <https://doi.org/10.3758/BF03197893>
- Eckerman, D. A., & Lanson, R. N. (1969). Variability of response location for pigeons responding under continuous reinforcement, intermittent reinforcement, and extinction. *Journal of the Experimental Analysis of Behavior*, 12(1), 73–80. <https://doi.org/10.1901/jeab.1969.12-73>
- Epstein, R., & Skinner, B. F. (1980). Resurgence of responding after the cessation of response-independent reinforcement. *Proceedings of the National Academy of Sciences of the United States of America*, 77(10 II), 6251–6253. <https://doi.org/10.1073/pnas.77.10.6251>
- Epstein, Robert. (1985). Extinction-Induced Resurgence: Preliminary Investigations And Possible Applications. *The Psychological Record*, 35(2), 143–153. <https://doi.org/10.1007/bf03394918>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research*

- Methods*, 39(2), 175–191. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17695343>
- Gallistel, C. R. (2012). Extinction from a rationalist perspective. *Behavioural Processes*, 90(1), 66–80. <https://doi.org/10.1016/j.beproc.2012.02.008>
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107(2), 289–344. <https://doi.org/10.1037/0033-295X.107.2.289>
- Herrick, R. M., & Bromberger, R. A. (1965). Lever displacement under a variable ratio schedule and subsequent extinction. *Journal of Comparative and Physiological Psychology*, 59(3), 392–398. <https://doi.org/10.1037/h0022063>
- Humphreys, L. G. (1939). The effect of random alternation of reinforcement on the acquisition and extinction of conditioned eyelid reactions. *Journal of Experimental Psychology*, 25(2), 141–158. <https://doi.org/10.1037/h0058138>
- Neuringer, A., Kornell, N., & Olufs, M. (2001). Stability and variability in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, 27(1), 79–94. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11199517>
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., ... Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195–203. <https://doi.org/10.3758/s13428-018-01193-y>
- Rescorla, R. A. (2001). Are associative changes in acquisition and extinction negatively accelerated? *Journal of Experimental Psychology: Animal Behavior Processes*, 27(4), 307–315. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11676082>
- Rescorla, Robert A. (1993). Inhibitory associations between S and R in extinction. *Animal Learning & Behavior*, 21(4), 327–336. <https://doi.org/10.3758/BF03197998>
- Rescorla, Robert A. (2000). Associative changes in exciters and inhibitors differ when they are conditioned in compound. *Journal of Experimental Psychology: Animal Behavior Processes*, 26(4), 428–438. <https://doi.org/10.1037/0097-7403.26.4.428>
- Rescorla, Robert A. (1999). Within-Subject Partial Reinforcement Extinction Effect in Autoshaping. *Quarterly Journal of Experimental Psychology*, 52b(1), 75–87. Retrieved from <https://journals.sagepub.com/doi/pdf/10.1080/713932693>
- Rescorla, Robert A., & Wagner, A. R. (1972). A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement. In A. . Black & W. . Prokasy (Eds.), *Classical conditioning II: current research and theory* (pp. 64–99). Retrieved from <https://pdfs.semanticscholar.org/afaf/65883ff75cc19926f61f181a687927789ad1.pdf>
- Smith, K. S., & Graybiel, A. M. (2013). A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron*, 79(2), 361–374. <https://doi.org/10.1016/j.neuron.2013.05.038>
- Stahlman, W. D., Young, M. E., & Blaisdell, A. P. (2010). Response variability in pigeons in a Pavlovian task. *Learning & Behavior*, 38(2), 111–118. <https://doi.org/10.3758/LB.38.2.111>
- Stahlman, W. David, & Blaisdell, A. P. (2011). The modulation of operant variation by the probability, magnitude, and delay of reinforcement. *Learning and Motivation*, 42(3), 221–236. <https://doi.org/10.1016/J.LMOT.2011.05.001>
- Stahlman, W. David, Roberts, S., & Blaisdell, A. P. (2010). Effect of reward probability on spatial and temporal variation. *Journal of Experimental Psychology: Animal Behavior Processes*, 36(1), 77–91. <https://doi.org/10.1037/a0015971>
- Thorndike, E. (1911). *Animal intelligence : experimental studies*. Retrieved from <https://www.worldcat.org/title/animal-intelligence-experimental-studies/oclc/844297>
- Todd, T. P., Vurbic, D., & Bouton, M. E. (2014). Behavioral and neurobiological mechanisms of extinction in Pavlovian and instrumental learning. *Neurobiology of Learning and Memory*, 108, 52–64. <https://doi.org/10.1016/J.NLM.2013.08.012>
- Wood, W., & Neal, D. T. (2007). A new look at habits and the habit-goal interface. *Psychological*

Review, 114(4), 843–863. <https://doi.org/10.1037/0033-295X.114.4.843>

*key absent on nonreinforced trials

Table 1: Reinforcement schedule for Experiment 1 and simulations. The experiment was entirely within subject and each bird was exposed to 25 trials of each cue. Trials lasted a total of 20s. The first presentation of each cue was always reinforced, and the remaining 24 trials were either reinforced 100% of the time or 25% of the time (i.e. 6 reinforced trials). During nonreinforced Y trials, the key was not presented. During extinction, each cue was presented 24 times and the key was always present for all cues.

Condition	Cue	Acquisition	Extinction
AB	A	100%	0%
AB	B	25%	0%
XY	X	100%	0%
XY	Y	25% *key absent on nonreinforced trials	0%
AB & XY	C	0%	100%

Table 2: Reinforcement schedule for Experiment 2. Participants were randomly assigned to condition AB or condition XY. Trials lasted a total of 7s, but the key was only present on the last 4s. During acquisition, the first presentation of each cue was always reinforced, and the remaining 16 trials were either reinforced 100% of the time or 25% of the time (i.e. 4 reinforced trials). During nonreinforced Y trials, the key was not presented. During extinction, each cue was presented 24 times and the key was always present for all cues. Cue C was included for both condition AB and XY to make the switch from acquisition to extinction more salient and to provide incentive to continue responding in extinction.

Figure 1: The associative structure of the Modified Law of Effect. Excitatory and inhibitory associations form between S-R following either positive or negative outcomes respectively. In addition, there are fixed lateral-inhibitory links between all possible responses that can occur in the presence of the stimulus. Responding to R1 is governed by the net strengths of the excitatory and inhibitory S-R1 associations, as well as lateral inhibition received from other responses. As the activation of S-R1 increases, so too does its lateral inhibition directed at all other responses with which it shares lateral inhibitory connections.

Figure 2: Example of trial screen displays used in Experiments 1 and 2. Colored rectangles served as discriminative stimuli (cues) and a white circle served as the response key. **a)** An example of a trial with one cue and the response key present. **b)** An example of an acquisition-phase trial with partially-reinforced Cue Y and the response key absent.

Figure 3: Simulated key pecks averaged across the 1000 agents following the training (upper panel) and extinction (lower panel) treatments of Experiment 1. Cues A and X were continuously reinforced and Cues B and Y were reinforced on 25% of trials but the key was not present on nonreinforced Y trials. Blocks represent two daily sessions in acquisition and one session in extinction. **a)** Mean number of pecks to the key on rewarded trials during presentation of each cue during acquisition. **b)** Proportion of pecks to the key during the presentation of each cue as a function of the number of pecks made on the first day of extinction.

Figure 4: Key pecks averaged across the 4 pigeons in Experiment 1. Cues A and X were continuously reinforced and cues B and Y were reinforced on 25% of trials but the key was not present on nonreinforced Y trials. Blocks represent two daily sessions. **a)** Mean number of pecks made to the key on rewarded trials during presentation of each cue during acquisition. **b)** Mean number of pecks made to the key during the presentation of each cue during extinction. **c)** Proportion of pecks made to the key during the presentation of each cue as a function of the number of pecks made on the first day of extinction.

Figure 5: Key presses on rewarded trials averaged across 20 human participants per condition in Experiment 2. Cues A and X were continuously reinforced and Cues B and Y were reinforced on 25% of trials but the key was not present on nonreinforced Y trials. Blocks include the mean value of key presses from 4 consecutive trials, except for Block 1 of acquisition which only includes the mean value of the first trial. **a)** Mean number of presses during rewarded trials to the key during presentation of each cue during acquisition for Condition AB. **b)** Mean number of presses made during rewarded trials to the key during presentation of each cue during acquisition for Condition XY. **c)** Mean number of pressed made during extinction trials for Condition AB. **d)** Mean number of presses made during extinction trials for Condition XY.

Supplementary Material

Figure S1: Acquisition data for Experiments 1 and 2 including pecks/presses during nonreinforced B and null Y trials. **a)** Mean number of key pecks per trial type during acquisition in Experiment 1, averaged across 2 session blocks. Reinforced trials are annotated as Cue_r. All of the A and X trials were reinforced but only 25% of B and Y trials were reinforced. B_nr refers to nonreinforced trials while the B stimulus was present. The large number of pecks to the key during presentation of this cue suggests the learning of a strong S-R inhibitory association. That more pecks were made during B_nr trials compared to A_r and X_r trials is likely due to the fact that eating from the food hopper detracted from key pecking on reinforced trials. Y_null refers to nonreinforced trials while the Y stimulus was present but the key was absent. Thus, what is recorded here are pecks made to the location of where the key would have been. It is an empirical question whether this is a uniquely independent response compared to pecking the key, but even if one considers this response to be the same as key pecking, the S-R inhibitory association for Cue Y is significantly smaller than that for Cue B. **b)** Mean number of presses made to the key during each trial type for condition AB. Cue A was continuously reinforced. Cue B was reinforced on 25% of trials. Reinforced trials are annotated as Cue_r. B_nr refers to nonreinforced trials while Cue B was present. Key pressing during this trial should facilitate the learning of a S-R inhibitory association. **c)** Mean number of presses made to the key during each trial type during acquisition for condition XY. Cue X was continuously reinforced. Cue Y was reinforced on 25% of trials but the key was not present on nonreinforced trials for Cue Y (referred to as Y_null trials). We still recorded key presses to the area where the key would have been but very few presses were made suggesting a weak or no S-R inhibitory association was formed during presentation of Cue Y.

Figure S2: Key presses to Cue C averaged across 20 human participants per condition in Experiment 2. Blocks include the summed value of 4 subsequent trials, except for block 1 of acquisition which only includes responding to the first trial. Cue C was not reinforced during acquisition but was continuously reinforced during extinction. This served to make the change from acquisition to extinction for the other cues more salient. **a)** Mean number of presses made to the key per trial to Cue C during the acquisition period for the other cues. There was no significant of Block, $F(4, 152) < 1.0$, or Condition, $F(1,38) < 1.0$. **b)** Mean number of presses made to the key per trial to Cue C during the extinction period for the other cues. There was no significant effect of Block, $F(5, 190) = 1.78$, $p = 0.12$, $\eta^2 = 0.01$, or Condition, $F(1,38) < 1.0$.

Figure Captions:

Figure 1: The associative structure of the Modified Law of Effect. Excitatory and inhibitory associations form between S-R following either positive or negative outcomes respectively. In addition, there are fixed lateral-inhibitory links between all possible responses that can occur in the presence of the stimulus. Responding to R1 is governed by the net strengths of the excitatory and inhibitory S-R1 associations, as well as lateral inhibition received from other responses. As the activation of S-R1 increases, so too does its lateral inhibition directed at all other responses with which it shares lateral inhibitory connections.

Figure 2: Example of trial screen displays used in Experiments 1 and 2. Colored rectangles served as discriminative stimuli (cues) and a white circle served as the response key. **a)** An example of a trial with one cue and the response key present. **b)** An example of an acquisition-phase trial with partially-reinforced Cue Y and the response key absent.

Figure 3: Simulated key pecks averaged across the 1000 agents following the training (upper panel) and extinction (lower panel) treatments of Experiment 1. Cues A and X were continuously reinforced and Cues B and Y were reinforced on 25% of trials but the key was not present on nonreinforced Y trials. Blocks represent two daily sessions in acquisition and one session in extinction. **a)** Mean number of pecks to the key on rewarded trials during presentation of each cue during acquisition. **b)** Proportion of pecks to the key during the presentation of each cue as a function of the number of pecks made on the first day of extinction.

Figure 4: Key pecks averaged across the 4 pigeons in Experiment 1. Cues A and X were continuously reinforced and cues B and Y were reinforced on 25% of trials but the key was not present on nonreinforced Y trials. Blocks represent two daily sessions. **a)** Mean number of pecks made to the key on rewarded trials during presentation of each cue during acquisition. **b)** Mean number of pecks made to the key during the presentation of each cue during extinction. **c)** Proportion of pecks made to the key during the presentation of each cue as a function of the number of pecks made on the first day of extinction.

Figure 5: Key presses on rewarded trials averaged across 20 human participants per condition in Experiment 2. Cues A and X were continuously reinforced and Cues B and Y were reinforced on 25% of trials but the key was not present on nonreinforced Y trials. Blocks include the mean value of key presses from 4 consecutive trials, except for Block 1 of acquisition which only includes the mean value of the first trial. **a)** Mean number of presses during rewarded trials to the key during presentation of each cue during acquisition for Condition AB. **b)** Mean number of presses made during rewarded trials to the key during presentation of each cue during acquisition for Condition XY. **c)** Mean number of presses made during extinction trials for Condition AB. **d)** Mean number of presses made during extinction trials for Condition XY.

Figure S1: Acquisition data for Experiments 1 and 2 including pecks/presses during nonreinforced B and null Y trials. **a)** Mean number of key pecks per trial type during acquisition in Experiment 1, averaged across 2 session blocks. Reinforced trials are annotated as Cue_r. All of the A and X trials were reinforced but only 25% of B and Y trials were reinforced. B_nr refers to nonreinforced trials while the B stimulus was present. The large number of pecks to the key during presentation of this cue suggests the learning of a strong S-R inhibitory association. That more pecks were made during B_nr trials compared to A_r and X_r trials is likely due to the fact that eating from the food hopper detracted from key pecking on reinforced trials. Y_null refers to nonreinforced trials while the Y stimulus was present but the key was absent. Thus, what is recorded here are pecks made to the location of where the key would have been. It is an empirical question whether this is a uniquely independent response compared to pecking the key, but even if one considers this response to be the same as key pecking, the S-R inhibitory association for Cue Y is significantly smaller than that for Cue B. **b)** Mean number of presses made to the key during each trial type for condition AB. Cue A was continuously reinforced. Cue B was reinforced on 25% of trials. Reinforced trials are annotated as Cue_r. B_nr refers to nonreinforced trials while Cue B was present. Key pressing during this trial should facilitate the learning of a S-R inhibitory association. **c)**

Mean number of presses made to the key during each trial type during acquisition for condition XY. Cue X was continuously reinforced. Cue Y was reinforced on 25% of trials but the key was not present on nonreinforced trials for Cue Y (referred to as Y_null trials). We still recorded key presses to the area where the key would have been but very few presses were made suggesting a weak or no S-R inhibitory association was formed during presentation of Cue Y.

Figure S2: Key presses to Cue C averaged across 20 human participants per condition in Experiment 2. Blocks include the summed value of 4 subsequent trials, except for block 1 of acquisition which only includes responding to the first trial. Cue C was not reinforced during acquisition but was continuously reinforced during extinction. This served to make the change from acquisition to extinction for the other cues more salient. **a)** Mean number of presses made to the key per trial to Cue C during the acquisition period for the other cues. There was no significant of Block, $F(4, 152) < 1.0$, or Condition, $F(1,38) < 1.0$. **b)** Mean number of presses made to the key per trial to Cue C during the extinction period for the other cues. There was no significant effect of Block, $F(5, 190) = 1.78$, $p = 0.12$, $\eta^2 = 0.01$, or Condition, $F(1,38) < 1.0$.