The content of compound conditioning

Justin A. Harris, Benjamin J. Andrew, & Evan J. Livesey

School of Psychology, University of Sydney

Abstract

In three experiments using Pavlovian conditioning of magazine approach, rats were trained with a compound stimulus, AB, and were concurrently trained with stimulus B on its own. The reinforcement rate of B, r_B, was either ½, ⅔, or ⅖ of r_{AB}. After extended training, the conditioning strength of A was assessed using probe trials in which A was presented alone. Responding during A was compared with that during AB, B, and a third stimulus, C, for which r_C = $r_{AB} - r_B$. In each experiment, the rats' response rate during A was almost identical to that during C (and during B, when $r_B = ½r_{AB}$). This suggests that, during AB conditioning, the rats had learned about r_A as being equal to $[r_{AB} - r_B]$, and implies that the content of their learning was a linear function of r. The findings provide strong support for rate-based models of conditioning (e.g., Gallistel & Gibbon, 2000). They are also consistent with the associative account of learning defined in the Rescorla-Wagner (1972) model, but only if the learning rate during reinforcement equals that during non-reinforcement.

Address correspondence to:

Justin Harris
School of Psychology
University of Sydney
Sydney 2006
Australia
Tel: +61 2 9351 2864
Email: justin.harris@sydney.edu.au

Some of the most important findings in the history of associative learning research have come from experiments demonstrating competition effects between conditioned stimuli (CSs). In broad terms, these cue-competition effects arise when two or more CSs are presented on the same trial and reinforced by presentation of an unconditioned stimulus (US). When each CS is subsequently presented on its own, conditioned responding is weak relative to that observed under control conditions where the CS had been reinforced on its own (Kamin, 1968; Kehoe, 1982; Mackintosh, 1976). This suggests that, when CSs are conditioned concurrently, they compete for conditioning strength. The nature of this competition is highlighted by demonstrations that, during concurrent reinforcement of two CSs, the amount that each CS conditions depends on the conditioning strength of the other CS. In a blocking design, conditioning of a stimulus is weaker if reinforced in compound with a previously conditioned stimulus than if reinforced in compound with a novel stimulus (Kamin, 1968; Mackintosh, 1971). Similarly, when a target CS is reinforced in compound with a second CS, the strength of conditioning of the target is affected by concurrent conditioning of the other CS: conditioning of the target CS is reduced if the other CS is well correlated with the US, whereas conditioning of the target CS is increased if the other CS is poorly correlated with the US (Wagner, 1969; Wagner, Logan, Haberlandt, & Price, 1968).

The most influential explanation of these cue-competition effects attributes them to a learning rule in which the amount of conditioning that accrues to each CS on a given trial is proportional to the difference between the total amount of conditioning that can be supported by the US, λ, and the sum of conditioning strengths already accrued to all CSs that are present on that trial (Rescorla & Wagner, 1972). This rule can be expressed formally as

$$\Delta V_X = \alpha\beta \left(\lambda - \sum_{i=1}^{n} V_i \right) \qquad \text{Equation 1}$$

In this equation, V_X is the conditioning strength (or "associative strength") of CS X; α and β are rate parameters related to properties of the CS and US, respectively. Because the change in strength of each CS is governed by the difference between λ and the summed strength of all (n) CSs present on that trial, V is ultimately shared among CSs that are conditioned concurrently. This means that, if two CSs are conditioned together to asymptote, the sum of their Vs will equal λ, and thus V of each CS will equal λ minus V of the other. Unfortunately, robust empirical test of this simple yet precise prediction has been elusive, due to the absence of any means by which to obtain an accurate estimate of V from conditioned responding.

Recently, Andrew and Harris (2011) have developed an experimental approach that allowed them to address a closely related question. They used a magazine approach paradigm with rats, in which the delivery of each food pellet was signaled by one of several CSs. One key feature of their procedure was that the length of each CS, and thus the CS-US interval, varied randomly from trial to trial according to a uniform distribution centered on a given mean. This ensured that the rats' response rates were steady as time elapsed during the trial (Harris, Gharaei, & Pincham, 2011), making it possible to draw meaningful comparisons between the magnitude of responding to different CSs that had different reinforcement rates, where reinforcement rate is defined by the inverse of the mean CS-US interval.

The particular question addressed by Andrew and Harris (2011) concerned the summation of conditioned responding when two separately conditioned stimuli are presented together as a compound. It is conventional among associative learning theories to assume that the strength of a conditioned response is given by the sum of the associative values of all stimuli present (e.g.,

Spence, 1936). In this respect, the rules that govern responding are assumed to parallel those that govern learning as specified in Equation 1. When two CSs are presented together, V of the compound will equal ΣV of the individual CSs, and responding to the compound should be proportional to ΣV. In three experiments, Andrew and Harris trained rats with three CSs: A, B, and X. The mean reinforcement rate, r, differed between the CSs, but in each case $r_X = r_A + r_B$. Once response rates to all stimuli had reached a stable level after extended training, occasional "probe" presentations of the compound AB were included in the conditioning sessions. Across all three experiments, the rats responded on these AB trials at the same rate as they responded to X. If we assume that the rats' responses to AB were based on the sum of what they had learned about A and B individually, then the content of that learning was linearly proportional to the rate at which those CSs had been reinforced.

Andrew and Harris (2011) concluded that their results provided strong support for a particular theoretical approach to conditioning, as exemplified by Rate Estimation Theory (RET), which assumes that r is the content of what animals learn during conditioning (Gallistel & Gibbon, 2000). They also noted that their results were not consistent with the specific predictions of the Rescorla-Wagner (1972) model since, according to Equation 1, the terminal V of a CS is a hyperbolic function of r, similar to that shown in Figure 1 (see Harris & Carpenter, 2011 for the derivation of this relationship from Equation 1). The shape of the hyperbolic function means that, for CSs A, B, and X in Andrew and Harris's experiments, $V_X < V_A + V_B$. Therefore, if $V_{AB} = V_A + V_B$ for compound AB, then $V_{AB} > V_X$, and responding to the compound AB should be greater than responding to X, which was not observed.

The hyperbolic relationship between V and r described by Harris and Carpenter (2011), and thus the predictions that Andrew and Harris
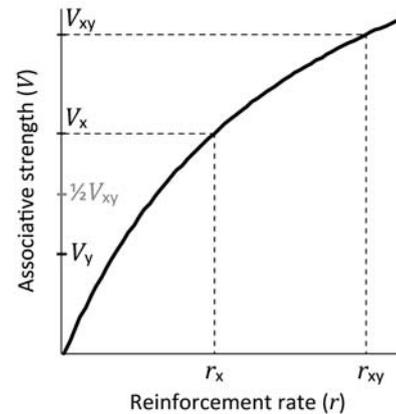

Figure 1. A rectangular hyperbola mapping the relationship between reinforcement rate, r, and associative strength (V), as derived from Equation 1 (see Harris & Carpenter, 2011 for derivation). In this example, a compound, XY, and one of its components, X, are conditioned concurrently such that $r_x = \frac{1}{2}r_{xy}$. By virtue of the shape of the hyperbolic function, $V_x > \frac{1}{2}V_{xy}$. If, as assumed by the Rescorla-Wagner model, $V_{xy} = V_x + V_y$, then $V_y < V_x$.

(2011) made for V_{AB}, were based on an assumption that the learning rate parameter, β, is smaller during non-reinforcement than during reinforcement. This assumption can be justified on a number of grounds. First, the assumption is important if one uses Equation 1 to model changes in V within a trial, as is necessary when explaining the effect of CS-US interval on conditioning strength. Without this assumption, within-trial extinction dominates over the acquisition of excitation whenever the major part of any trial is spent without reinforcement. The assumption is also necessary for Equation 1 to explain certain phenomena, such as relative signal validity (Wagner, et al., 1968), a fact that has led researchers, including Rescorla and Wagner (1972), to assume that the learning rate for non-reinforcement is substantially lower than for reinforcement (also Blough, 1975; Wagner, 1981). Indeed, Rescorla (2002) has provided experimental data comparing rates of associative change during conditioning and extinction, which have confirmed the assumed difference in learning rates between reinforce-

3

ment and non-reinforcement. Nonetheless, the assumption has important implications in the present context because a linear relationship between V and r can be derived from Equation 1 if β takes the same value for reinforcement and non-reinforcement (see Appendix 1 for derivation). As such, the results reported by Andrew and Harris can be taken to support the simplest trial-based version of the Rescorla-Wagner model[1].

In the present paper, we describe three experiments that use an approach similar to that developed by Andrew and Harris (2011) to investigate how much rats learn about CSs conditioned in compound. The experiments test the assumption, at the heart of the Rescorla-Wagner (1972) model, that learning is regulated by the sum of Vs of all CSs present on each conditioning trial, and therefore V of each CS from a compound will equal V of the compound minus V of the other CS. A comparable prediction is made by RET (Gallistel & Gibbon, 2000) that the value of r assigned to each CS in the compound will equal the r of the compound minus the r assigned to the other CS. Thus both models assume that learning about a compound is divided between the CSs that make up the compound, but the models differ in the content of that learning: V versus r. This difference in content can translate to a difference in predicted responding. To illustrate, consider the case in which a compound XY is reinforced at rate r_{xy}, and on intervening trials X is reinforced on its own at half this rate. If rats learn specifically about r (as argued by RET), or if their learning about V is linearly proportional to r (as predicted by the Rescorla-Wagner model using the same value of β for reinforcement and non-reinforcement), then what the rats learn about Y must equal what they learn about X. Therefore, responding to X should equal that to Y. However, if V is related to r by a hyperbolic function (as predicted by the Rescorla-Wagner model when β for non-

reinforcement is less than β for reinforcement), then $V_x > \frac{1}{2}V_{xy}$ (see Figure 1). Because $V_{xy} = V_x + V_y$, it follows that $V_y < \frac{1}{2}V_{xy} < V_x$. Therefore, under these assumptions, the Rescorla-Wagner model predicts responding to Y will be less than responding to X. The experiments presented here test these predictions. The data from all three experiments can be downloaded from the website:

http://sydney.edu.au/science/psychology/staff/justinh/downloads/

Experiment 1

Experiment 1 followed the logic described above (and illustrated in Figure 1). Rats were trained with a compound, AB, and three single CSs, B, C, and D, each of which varied in length from trial to trial around a mean of 30 s. Presentations of AB were reinforced on 100% of trials. Presentations of B and C were reinforced on 50% of trials, while D was reinforced on 25% of trials. It was expected that the rats would respond most to AB, at similar rates to B and C, and least to D. Training was continued for 44 days, but the last 10 days included non-reinforced[2] "probe" trials on which A was presented alone. If the rats learned about reinforcement rates, or if their learning about V was linearly proportional to r, we would expect that what they learned about A would equal what they had learned about B and C. In this case, the rats should respond to A at the same level as they respond to B and C. If, however, the rats learned about the V, but the rate of learning was less during non-reinforcement than during reinforcement, then V_A should be less than V_B and V_C, and therefore the rats should respond less to A than to B and C.

---

[1] The authors thank Daniel Gottlieb for pointing this out.

[2] We chose to present the probes without reinforcement because this would, if anything, lead to a reduction in the observed response on those trials. We reasoned that this is the more conservative approach when it comes to testing the hypothesis that A = B/C (rather than A < B/C).

Methods

Subjects

Sixteen experimentally naive male Hooded Wistar rats (Rattus norvegicus; 8 to 10 weeks of age at the start of the experiment) were obtained from the Laboratory Animal Services breeding unit at The University of Adelaide, South Australia. During the experiment, they were housed in groups of 8 in large white plastic tubs, measuring 26 x 59 x 37cm (height x length x depth), located in the animal colony maintained by the School of Psychology at the University of Sydney. They had unrestricted access to water in the home tubs. Three days prior to commencement of the experiment, their access to food was restricted to 2 hr per day (to commence half an hour after the end of the daily training sessions).

Apparatus

Rats were trained and tested in 16 Med Associates™ conditioning chambers measuring 28.5 x 30 x 25 cm (height x length x depth). The end walls of each chamber were made of aluminum; the sidewalls and ceiling were Plexiglas™. The floor of the chamber consisted of stainless-steel rods, 0.5 cm in diameter, spaced 1.5 cm apart. Each chamber had a recessed food magazine in the center of one end wall, with an infra-red LED and sensor located just inside the magazine to record entries by the rat. A small metal cup measuring 3.5 cm in diameter and 0.5 cm deep was fixed on the floor of each food magazine. Attached to the food magazine was a dispenser delivering 45 mg food pellets (Noyes Formula P; Research Diets Inc, New Brunswick, NJ). 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 70dB. Experimental events were controlled and recorded automatically by computers and relays located in the same room.

Two auditory and two visual stimuli were presented from four spatially separated sources. A tone (2.9 kHz) was produced from a piezo buzzer positioned on the floor of the sound-attenuating shell behind each conditioning chamber. White noise (78dB) was presented from a speaker mounted on the wall of each conditioning chamber above and to the right of the food magazine. A flashing light (2 Hz; 3.0cd/m^2) was emitted by a 3x3 array of white LEDs, located on the floor of the sound-attenuating shell in front of the conditioning chamber. A steady light (30cd/m^2) was produced by an incandescent bulb mounted high on the back wall of the sound-attenuating shell.

Procedure

On the day before training began, the rats received a single 20-min magazine training session during which 20 food pellets were presented on a variable-time (VT) 1 min schedule, with no stimulus presentations. The rats then received daily conditioning sessions 5 days per week for a total of 44 days. On Days 1 to 34, each session consisted of 16 intermixed presentations of three CSs (B, C, and D) and one compound (AB), for a total of 64 presentations. The allocation of the stimuli was counterbalanced across rats, with the constraint that A and B were always in different modalities. This means that A (our probe) was in the same modality as C for 8 of rats, and in the same modality as D for 8 rats. Since D was reinforced less than C (and produced less responding), the "average" generalization could have biased responding to A to be the less than C. Thus once again, in this respect our experiment is a more conservative test of the hypothesis that A=C. The order of presentations was randomized within each quarter of the session (i.e., each quarter contained 4 presentations of each CS and compound). The duration of the stimuli varied from trial to trial, with a mean length of 30 s (range 2 to 58 s). Every presentation of the compound, AB, ended with the delivery of a single food pellet that coincided with the termination of the compound. Separate presentations of B were reinforced with a food pellet on 50% of trials.

Presentations of C were also reinforced on 50% of trials, while D was reinforced on 25% of trials. Care was taken to ensure that the reinforced trials were sampled uniformly from the full range of CS durations, so that the mean length of CS presentations was the same on reinforced and non-reinforced trials. Days 35 to 44 continued with these CS and compound presentations, but additionally included two 30-s probe presentations of A on its own. These occurred on Trials 22 and 43 (of 66 trials). Across all 44 days, the average inter-stimulus-interval was 90 s (the intervals varied randomly according to a uniform distribution, but with a minimum of 35 s). The timing of each photo-beam interruption by head entry into the magazine was recorded during each CS, and each 30-s pre-CS period. Each session lasted approximately 130 min.

Results & Discussion

The left side of Figure 2 plots the mean response rate during the compound, AB, and during each of the single CSs, B, C, and D, across the 44 days of Experiment 1. Responding increased over the first 20 days, and remained relatively steady thereafter. A repeated measures ANOVA, with Greenhouse-Geisser correction, was used to analyze response rates during AB, B, C and D across the whole experiment. This revealed significant main effects of stimulus, $F(2.5,37) = 19.21$, $p < .001$, and of Day, $F(4,60) = 15.32$, $p < .001$, and a significant interaction between stimulus and Day, $F(5,75) = 3.24$, $p = .011$. Further analyses revealed there to be a significant interaction between the linear trend across stimuli (i.e., as a function of reinforcement rate) and the linear trend across days, $F(1,15) = 15.58$, $p = .001$, as well as between the quadratic trends across stimulus and across days, $F(1,15) = 8.15$, $p = .012$. These confirm what is evident in Figure 2, that the slope and curvature of the acquisition curves are themselves a function of each CS's reinforcement rate, as would be expected given the differences in the terminal response rates between CSs.

From Days 35 to 44, each session included two probe trials with A, as shown in Figure 2. The right side of the figure shows the mean response rates during these probe trials, and during presentations of AB, B, C and D, averaged over the last 10 days. An ANOVA of these data revealed a significant main effect of CS, $F(2.7,40) = 11.94$, $p < .001$. This was followed by a series of contrasts that compared responding to A with responding to each of the other CSs. This showed that response rates during A were significantly lower than during AB, $F(1,15) = 7.37$, $p = .016$, and were significantly greater than during D, $F(1,15) = 20.39$, $p < .001$. However, response rates during A did not differ from those during B or C, both $Fs(1,15) < 1.22$, $ps > .287$. As is evident in Figure 2, mean response rate during A (0.74) was marginally below that during B (0.76) and marginally above the mean response rate during C (0.68). In other words, the rats responded to A as if it had been reinforced at the same rate as B and C.

The probe presentations of A occurred on Trials 22 and 43, approximately one third and two thirds the way through each test session. If there were a consistent bias in response rate at either the start or end of each session, this would affect response rates measured for the trained CSs but not the probe trials with A. To take account of this possibility, we repeated the above analysis but excluded all trials in the first and last quarter of each session. There was no substantive difference between this analysis and the one reported above: response rates to A were still significantly different from AB, $F(1,15) = 10.54$, $p = .005$, and D, $F(1,15) = 10.70$, $p = .005$, but were not different from B, $F(1,15) = 2.57$, $p = .130$, or C, $F(1,15) < 1$.

The results of this experiment clearly suggest that what rats learn about A is equivalent to what they learn about B or C. However, this result cannot be subjected to the standard statistical test because the evidence rests on acceptance of the null hypothesis. To address this type of problem, Gallistel (2009) has recent-
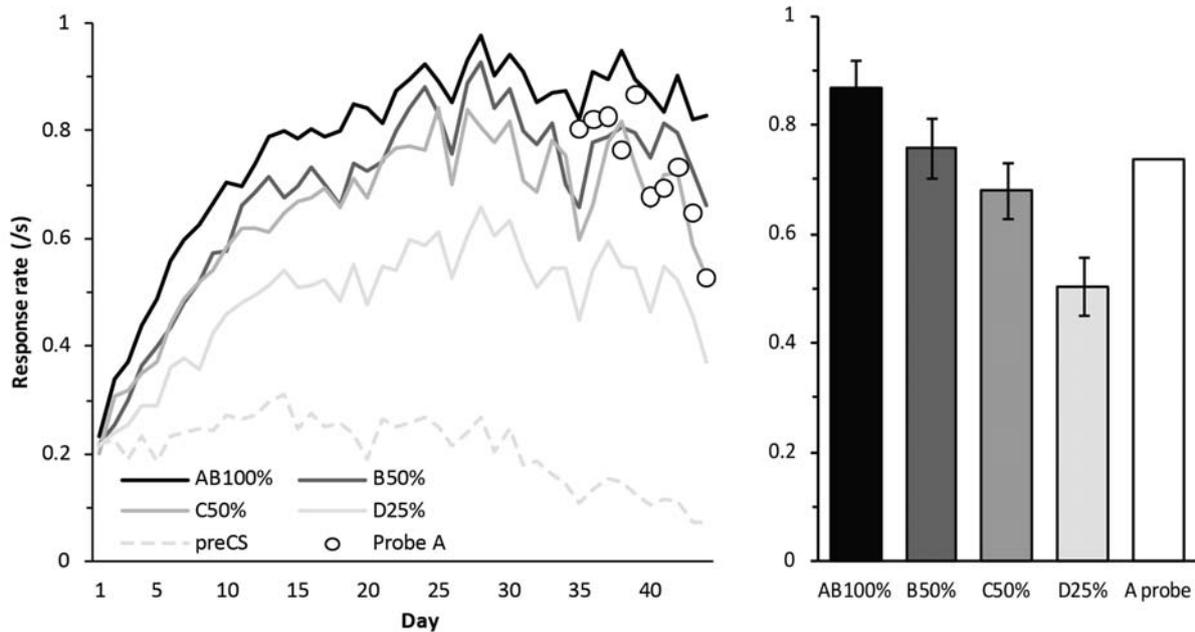
Figure 2. **Left**: Mean response rates (responses per second) to the trained compound (AB) and the individual CSs (B, C, and D) on each day of Experiment 1. Stimulus durations varied from trial to trial around a mean of 30 s, and were reinforced with a food pellet on either 100% (AB), 50% (B and C), or 25% (D) of trials. On Days 35 to 44, each session included probe trials of stimulus A. **Right**: Response rates to the compound and all CSs, averaged over the final 10 days. Vertical bars show within-subject SEMs for the difference between response rates to each trained stimulus and the probe presentations of A.

ly described how a simple Bayesian analysis can be used to compare the likelihood of the null hypothesis being correct, given the available data, versus the likelihood of some alternative hypothesis. We have applied this analysis to the present data to calculate the likelihood that response rates to A and B are equal ($R_A = R_B$), and compare this with the likelihood that $R_A < R_B$. We have assumed that the difference between R_A and R_B is normally distributed with a variance equal to the unbiased estimate of the variance of the sample. Not surprisingly, the likelihood function is maximal when the difference between R_A and R_B is close to zero.

To calculate the marginal likelihood of the hypothesis that $R_A < R_B$, it is necessary to specify the possible minimum and maximum size of this difference. According to the Rescorla-Wagner model, the predicted associative strength of A should equal [$V_{AB} - V_B$], but to specify R_A, we must define the relationship between R and V.

This was recently attempted by Harris and Carpenter (2011) for the same conditioning paradigm used here. They found the relationship between R and V to be well described as a rectangular hyperbola, $R = aV / [cV + d]$ (where a, c, and d are non-negative), like that shown in Figure 1. With respect to R_A, the predicted value will depend on the curvature of the hyperbolic function, which is determined by parameter c. At one extreme, when $c = 0$, the relationship is linear ($R = aV/d$), and so $R_A = [R_{AB} - R_B]$. This defines the minimum value of R_A. (In the analyses we have conducted here, we set this minimum at [$R_{AB} - R_B + R_{pre}$], which assumes that responding to A is added to the base response rate measured during the pre-CS interval. This assumption improves the marginal likelihood of this hypothesis.) For values of $c > 0$, the curvature of the hyperbolic function increases the predicted value for R_A, towards an upper limit set by R_B (this is the limit because $V_A < V_B$).

Therefore, we can define the alternative hypothesis as stating that R_A will fall somewhere within the range from $[R_{AB} - R_B + R_{pre}]$ up to R_B. If we allow this range to include R_B, this alternative hypothesis covers the full scope of predictions from the Rescorla-Wagner model (i.e., including the prediction that $R_A = R_B$ when β does not differ between reinforced and non-reinforced trials). The odds ratio of the marginal likelihoods of this hypothesis versus the null hypothesis favors the null hypothesis 5:1. In other words, based on the data from Experiment 1, the specific prediction that $R_A = R_B$ is five times more likely to be correct than the looser prediction that $[R_{AB} - R_B + R_{pre}] \leq R_A \leq R_B$. The equivalent analysis of the difference between R_A and R_C led to a similar conclusion, that the odds for the null hypothesis ($R_A = R_C$) versus the alternative hypothesis favored the null 10:1.

Experiment 2

In Experiment 1, rats were trained with a compound, AB, that was reinforced on 100% of presentations, while they were concurrently trained with B alone reinforced on 50% of presentations. After extended training, the rats were tested for responding to A. This revealed that the conditioning strength of A was equal to that of B, consistent with the proposal that what rats learned about A could be estimated by the difference in reinforcement rates between AB and B (i.e., that $r_A = r_{AB} - r_B$). Experiment 2 provided a further test of this prediction, using different proportional reinforcement rates. Rats were again trained with AB reinforced on 100% of trials, while they were concurrently trained with B reinforced on 66% of trials, as well as two other CSs, C and D, that were reinforced on 33% and 17% of trials, respectively. After extended training, we tested the rats for responding to A. If they had learned an estimate of the reinforcement rate of A based on the difference in reinforcement rates between AB and B, then their response rate to A should equal that to C.

Methods

Subjects and apparatus

Sixteen experimentally naive male Hooded Wistar rats (7 to 8 weeks of age at the start of the experiment) were obtained from the same source and were housed in the same manner as described for Experiment 1. The rats were trained and tested in the same 16 Med Associates™ conditioning chambers, using the same four counterbalanced stimuli, as in Experiment 1.

Procedure

Rats first received a single 20-min magazine training session. They then received daily conditioning sessions 5 days per week for a total of 40 days. On Days 1 to 30, each session consisted of 12 intermixed presentations of three CSs (B, C, and D) and a compound (AB), for a total of 48 presentations. The order of presentations was randomized within each quarter of the session (i.e., each quarter contained 3 presentations of each of the CSs). The duration of each CS or compound varied from trial to trial, with a mean length of 20 s (range 2 to 38 s). Every presentation of the compound, AB, ended with the delivery of a single food pellet. Separate presentations of B were reinforced with a food pellet on 66% of trials. Presentations of C were reinforced on 33% of trials, while D was reinforced on 17% of trials (2 out of 12). Care was taken to ensure that the reinforced trials were sampled uniformly from the full range of CS durations, so that the mean length of CS presentations was the same on reinforced and non-reinforced trials. Days 31 to 40 continued with these CS and compound presentations, but additionally included two probe presentations of A on its own at trials 16 and 32. Across all 40 days, the average inter-stimulus-interval was 100 s (the intervals varied randomly according to a uniform distribution, but with a minimum of 35 s). The timing of each photo-beam interruption by head entry into the magazine was recorded during each CS, and each 20-s pre-CS period. Each session lasted approximately 110 min.

Results & Discussion

The left side of Figure 3 plots the mean response rate during the compound, AB, and during each of the single CSs, B, C, and D, across all 40 days of Experiment 2. Responding increased over the first 20 days, and remained relatively steady thereafter. A repeated measures ANOVA, analyzing response rates during AB, B, C and D across the whole experiment, revealed significant main effects of stimulus, $F(2,30) = 53.72$, $p < .001$, and of Day, $F(3.8,57) = 12.16$, $p < .001$, and a significant interaction between stimulus and Day, $F(8.2,123) = 5.32$, $p < .001$. Further analyses revealed there to be a near-significant interaction between the linear trend across stimuli and linear trend across days, $F(1,15) = 4.29$, $p = .056$, and significant interactions between linear trend across stimuli and quadratic trend across days, $F(1,15) = 23.90$, $p < .001$, as well as between the quadratic trend across stimuli and both linear and quadratic trends across days, $F(1,15) \geq 27.58$, $ps < .001$. These confirm that the slope and curvature of the acquisition curves are themselves a function of reinforcement rate.

From Days 31 to 40, each session included two probe trials with A, as shown in Figure 3. The right side of Figure 3 shows mean response rates during these probe trials, and during presentations of AB, B, C and D, averaged over the last 10 days. An ANOVA on these data showed there to be a significant main effect of CS, $F(2.3,34) = 29.56$, $p < .001$. This was followed by contrasts comparing response rates between A and each of the other CSs. The analysis showed that response rates during A were significantly lower than during AB, $F(1,15) = 6.56$, $p = .022$, and B, $F(1,15) = 8.62$, $p = .010$, and were significantly higher than during D, $F(1,15) = 56.24$, $p < .001$. Response rates during A did not differ from those during C, $F(1,15) < 1$.
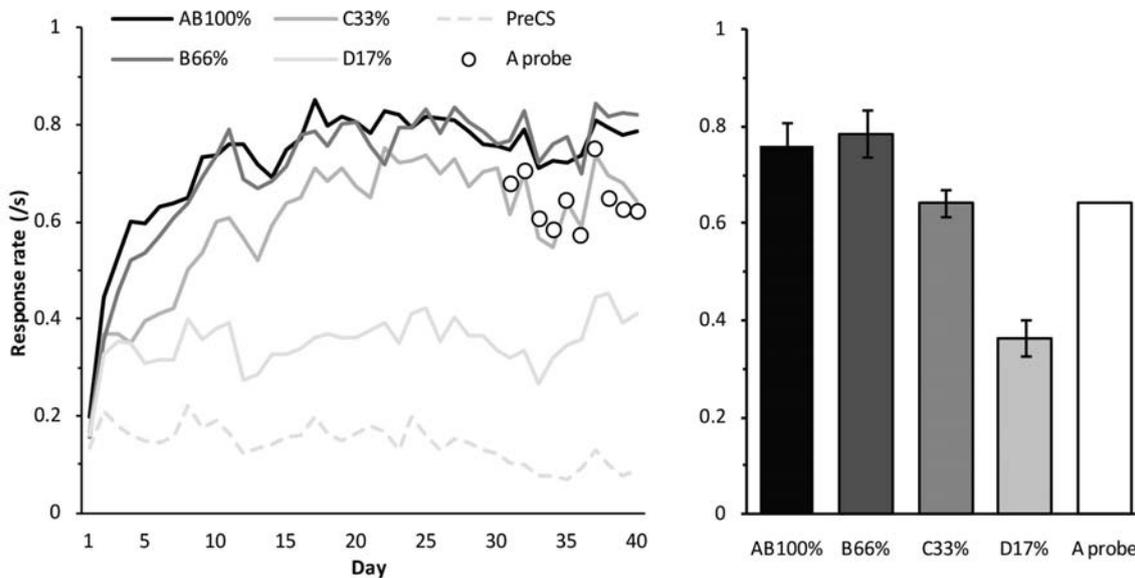

Figure 3. **Left**: Mean response rates (responses per second) to the trained compound (AB) and the individual CSs (B, C, and D) on each day of Experiment 2. Stimulus durations varied from trial to trial around a mean of 20 s, and were reinforced with a food pellet on either 100% (AB), 66% (B), 33% (C), or 17% (D) of trials. On Days 31 to 40, each session included probe trials of stimulus A. **Right**: Response rates to the compound and all CSs, averaged over the final 10 days. Vertical bars show within-subject SEMs for the difference between response rates to each trained stimulus and the probe presentations of A.

As can be seen in Figure 3, these were almost identical – the mean response rates during A and C were both 0.64. In other words, the rats responded to A as if it had been reinforced at the same rate as C.

As for Experiment 1, we repeated the above analysis but excluded all trials in the first and last quarter of each session. There was no substantive difference between this analysis and the one reported above: response rates to A were still significantly different from AB, $F(1,15) = 5.42$, $p = .034$, from B, $F(1,15) = 5.84$, $p = .029$, and from D, $F(1,15) = 45.13$, $p < .001$, but were not different from C, $F(1,15) < 1$.

We again explored further the comparison between R_A and R_C using the Bayesian analysis proposed by Gallistel (2009). As to be expected from the fact that R_A and R_C were almost identical, the likelihood function for the difference between R_A and R_C was centered very close to zero. We compared the likelihood of the null hypothesis, that $R_A = R_C$, with the likelihood of the hypothesis that $R_A \leq R_C$. For this second hypothesis, we specified the predicted value of R_A as falling within a range between $[R_{AB} - R_B + R_{pre}]$ and R_C. The odds ratio of the marginal likelihoods of this hypothesis versus the null hypothesis favors the null by 15:1.

Experiment 3

Experiments 1 and 2 both showed that, when rats are trained with a compound AB, what they learn about A can be deduced from the difference between the reinforcement rates of AB and B. Experiment 3 extends this investigation by using a different means to manipulate reinforcement rate. The proportion of trials that were reinforced was kept at 100% for all CSs, while the reinforcement rates of the compound and single CSs were manipulated by varying the mean duration of their CS-US intervals. Thus, rats were trained with a compound, AB, that had a mean CS-US interval of 30 s, and with three single CSs that had mean

CS-US intervals of 50 s (C), 75 s (B), and 150 s (D). Note that the reinforcement rates of B and C sum to equal that of AB (i.e., 1/75 + 1/50 = 1/30). In effect, $r_B = \frac{2}{5}r_{AB}$, and $r_C = \frac{3}{5}r_{AB}$. If the rats learn to estimate the reinforcement rate of A as equal to the difference between AB and B, then they should respond to A at the same rate as they respond to C.

Methods

Subjects and apparatus

Sixteen experimentally naive male Hooded Wistar rats (8 to 10 weeks of age at the start of the experiment) were obtained from the same source and were housed in the same manner as described for Experiment 1. The rats were trained and tested in the same 16 Med Associates™ conditioning chambers, using the same four counterbalanced stimuli, as in Experiment 1.

Procedure

As in Experiment 1, rats received a single 20-min magazine training session before commencing daily conditioning sessions 5 days per week for a total of 42 days. On Days 1 to 35, each session consisted of 12 intermixed presentations of three CSs (B, C, and D) and one compound (AB), for a total of 48 presentations. The order of presentations was randomized within each quarter of the session (i.e., each quarter contained 3 presentations of each of the CSs). The duration of each CS or compound varied from trial to trial, but regardless of length each CS terminated with delivery of a food pellet. Presentations of AB had a mean duration of 30 s (2 to 58 s); presentations of B on its own had a mean duration of 75 s (range 2 to 148 s); presentations of C had a mean duration of 50 s (2 to 98 s); and presentations of D had a mean of 150 s (2 to 298 s) for the first 15 days of training, but this was then reduced to a mean of 100 s (2 to 198 s) in the hope that increasing the overall reinforcement rate would lift response rates to all CSs. Days 36 to 42 continued with these CS and compound presentations, but additionally included two

non-reinforced probe presentations of A on its own, at trials 17 and 27. Across all 42 days, the average inter-stimulus-interval was 100 s (the intervals varied randomly according to a uniform distribution, but with a minimum of 35 s). Head entries into the magazine were recorded during each CS and the 30-s pre-CS period. Each session lasted approximately 2 hr.

Results & Discussion

The left side of Figure 4 plots the mean response rate during the compound, AB, and during each of the single CSs, B, C, and D, across all 42 days of Experiment 3. Responding increased gradually over the first 30 days, and remained relatively steady thereafter. A repeated measures ANOVA, with Greenhouse-Geisser correction, was used to analyze response rates during AB, B, C and D across the whole experiment. For this analysis, response rates were averaged over 2-day bins in order to make up for instances of missing data due to computer failure for four boxes (affecting Rats 9 to 12 on Day 30) and a fault with the photo-beam in one box (Rat 8 on Day 8). The ANOVA revealed significant main effects of stimulus, $F(1.6,24) = 35.82$, $p < .001$, and of Day, $F(3.4,51) = 11.28$, $p < .001$, and a significant interaction between stimulus and Day, $F(5.6,84) = 2.98$, $p = .013$. Further analyses revealed there to be a near-significant interaction between the linear trend across stimuli and linear trend across days, $F(1,15) = 4.29$, $p = .056$, and significant interactions between linear trend across stimuli and quadratic trend across days, $F(1,15) = 23.90$, $p < .001$, as well as between the quadratic trend across stimuli and both linear and quadratic trends across days, $F(1,15) \geq 27.58$, $ps < .001$. These confirm that the slope and curvature of the acquisition curves are themselves a function of reinforcement rate.
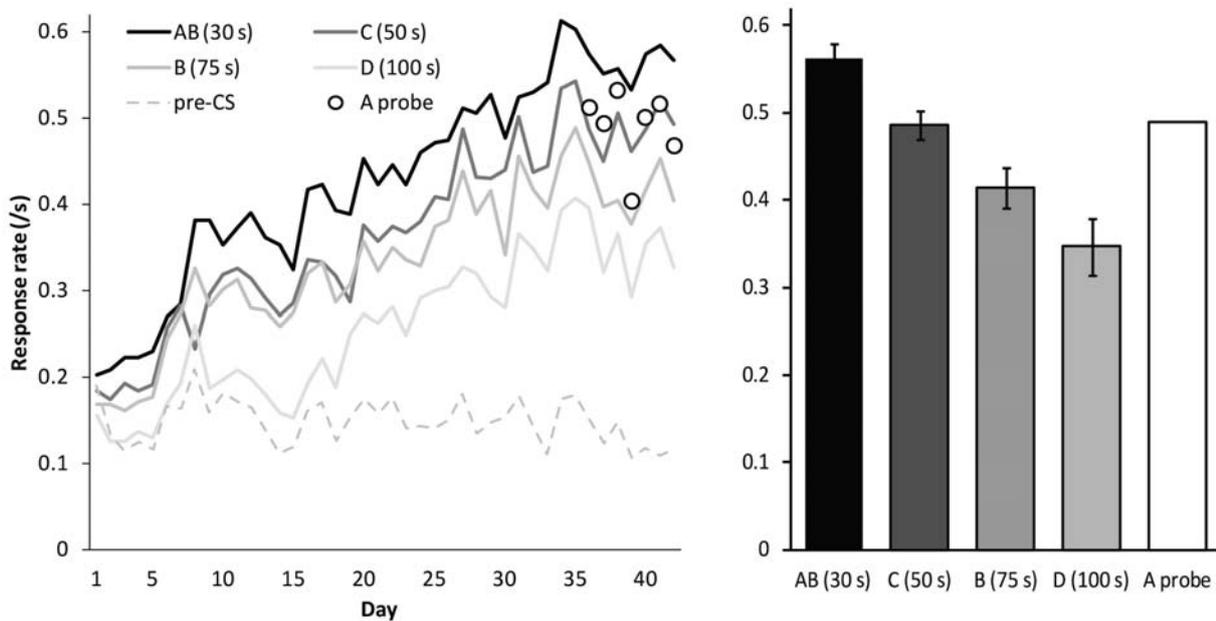

Figure 4. **Left**: Mean response rates (responses per second) to the trained compound (AB) and the individual CSs (B, C, and D) on each day of Experiment 3. All stimulus presentations terminated with the delivery of a food pellet, and stimulus durations varied from trial to trial around a mean of 30 s (AB), 75 s (B), 50 s (C), or 100 s (D) of trials. On Days 36 to 42, each session included probe trials of stimulus A. **Right**: Response rates to the compound and all CSs, averaged over the final 7 days. Vertical bars show within-subject SEMs for the difference between response rates to each trained stimulus and the probe presentations of A.

From Days 35 to 42, each session included two probe trials with A, as shown in Figure 4. The right side of the figure shows mean response rates during these probe trials, and during presentations of AB, B, C and D, averaged over the last 7 days. An ANOVA established that there was a significant main effect of CS, $F(1,15) = 20.98$, $p < .001$. Follow-up contrasts compared responding to A with responding to the other CSs. Response rates during A were significantly lower than during AB, $F(1,15) = 18.84$, $p < .001$, but significantly higher than during B, $F(1,15) = 11.94$, $p = .004$, or D, $F(1,15) = 20.56$, $p < .001$. However, response rates during A did not differ from those during C, $F(15) < 1$. Indeed, as can be seen in Figure 4, these were almost identical, with the mean response rates during A and C both equal to 0.49. In other words, the rats responded to A as if it had been reinforced at the same rate as C.

As for Experiments 1 and 2, we repeated the above analysis but excluded all trials in the first and last quarter of each session. There was little difference between this analysis and the one reported above: response rates to A were still significantly different from AB, $F(1,15) = 8.80$, $p = .010$, from B, $F(1,15) = 11.38$, $p = .004$, and from D, $F(1,15) = 42.52$, $p < .001$, but were not different from C, $F(1,15) = 2.33$, $p = .148$.

We have again used the Bayesian analysis proposed by Gallistel (2009) to explore further the comparison between R_A and R_C. The likelihood function for the difference between R_A and R_C was centered very close to zero. We compared the likelihood that $R_A = R_C$ to the likelihood that R_A falls within a range between $[R_{AB} - R_B + R_{pre}]$ and R_C. The odds ratio of the marginal likelihoods of this hypothesis ($[R_{AB} - R_B + R_{pre}] \leq R_A \leq R_C$) versus the null hypothesis ($R_A = R_C$) favors the latter by 16:1.

General Discussion

The three experiments presented here replicate the common observation that, if two CSs are conditioned simultaneously as a compound, conditioned responding is weaker when either CS is subsequently presented on its own than when the two are presented together. The most popular explanation for this effect is of mutual "overshadowing" between the CSs, such that what is learned about the compound is divided between the two CSs (Kamin, 1968). It is typically assumed that this division is based on a simple arithmetic principle that learning about the compound is the sum of what is learned about each CS (Gallistel & Gibbon, 2000; Rescorla & Wagner, 1972). The three experiments described here confirm this arithmetic principle, but only as applied to the reinforcement rates, r, of the CSs and their compound. In each experiment, rats were trained with a compound, AB, reinforced at a particular rate, r_{AB}, while being concurrently trained with one of the two CSs, B, reinforced on its own at a lower rate, r_B. When the rats were eventually presented with CS A on its own, they responded to A as if it had been reinforced at a rate equal to the difference between r_{AB} and r_B. This pattern of results for A was obtained in all three experiments, even though they used different reinforcement rates for AB and B. Thus, when $r_B = \frac{1}{2}r_{AB}$ in Experiment 1, the rats responded to A as if $r_A = \frac{1}{2}r_{AB}$; when $r_B = \frac{2}{3}r_{AB}$ in Experiment 2, they responded to A as if $r_A = \frac{1}{3}r_{AB}$; and when $r_B = \frac{2}{5}r_{AB}$ in Experiment 3, they responded to A as if $r_A = \frac{3}{5}r_{AB}$. This rule held both when r_{AB} and r_B were manipulated by changing the proportion of trials that were reinforced (Experiments 1 and 2) and by varying the mean length of the CS-US interval (Experiment 3).

The findings reported here complement our recent report showing that, when rats were presented with a compound composed of two separately conditioned stimuli, they responded to the compound as if it had been reinforced at a rate equal to the sum of rs of the individual CSs (Andrew & Harris, 2011). Both sets of findings show that what rats learn about individual stimuli during compound conditioning, and what they infer about a

compound of two individually conditioned stimuli, preserves the additivity of the reinforcement rates of those stimuli. This means that the content of their learning is linearly related to reinforcement rates, as argued by rate-based models of conditioning (Gallistel & Gibbon, 2000).

The present observations provide some qualified support for associative models that identify the content of learning with V, as defined in Equation 1, and which assume that learning about CSs in a compound is limited by the sum of what has already been learned about each CS (e.g., Rescorla & Wagner, 1972). As shown in Appendix A, if the learning rate does not differ between reinforcement and non-reinforcement, Equation 1 predicts that V is linearly proportional to r, which is perfectly consistent with the findings reported here and by Andrew and Harris (2011). However, there are important consequences if we commit the Rescorla-Wagner model to the assumption that β does not change between reinforcement and non-reinforcement. First, the model is no longer able to explain the relative validity effects (Wagner, et al., 1968) that have been taken as key evidence in support of the model since its inception (Rescorla & Wagner, 1972). The assumption is also directly contradicted by Rescorla's (2002) evidence that associative change is slower during non-reinforcement than reinforcement. Finally, if Equation 1 is used to model changes in conditioning strength at the sub-trial level, which is necessary to explain how conditioning strength is systematically affected by the CS-US interval (as shown here in Experiment 3; see also Harris & Carpenter, 2011), then the V of any prolonged CS will at best be a small fraction of λ because extinction during each trial will dominate over excitatory conditioning at the end of the trial. Reasons such as these have given rise to the common assumption that the learning rate during non-reinforcement is much lower than during reinforcement (Blough, 1975; Rescorla & Wagner, 1972; Wagner, 1981).

The present findings are problematic for an associative mechanism based on Equation 1 if, for the reasons just described, the learning rate for reinforcement is higher than for non-reinforcement. Harris and Carpenter (2011) showed that, under these conditions, the relationship between V and r is hyperbolic. This relationship means that the difference between V_{AB} and V_B will always be smaller than the value of V that is established when $r = r_{AB} - r_B$ (see Figure 1). In other words, V_A should be less than what was observed in each of the present experiments. This conclusion may appear to overstate the falsifiability of the model's prediction since that prediction does not specify how much lower the response rate to A should have been. One might also imagine that responding to A could be elevated by its within-compound association with B (although the repeated presentations of B alone throughout training should have greatly reduced this association; Rescorla, 1983; Rescorla & Freberg, 1978). Moreover, one could remain agnostic about the difference in learning rates between reinforcement and non-reinforcement, preferring to allow the rates to differ or not depending on circumstances. As such, one could argue that the Rescorla-Wagner model predicts that the V established when $r = [r_{AB} - r_B]$ is the upper-limit on the range of values that V_A could take, and as such, the present data do not constitute strong evidence against the model since its predictions can include (or approximate) the observed data. Nonetheless, the data reported here offer at best weak support for this rather broad hypothesis, but represent much stronger support for the more specific hypothesis that learning is directly proportional to reinforcement rate. Using a Bayesian analysis (Gallistel, 2009) to estimate the marginal likelihoods of the two hypotheses, we found that the odds strongly favored the specific hypothesis in all three experiments (by 5:1 or 10:1 in Experiment 1, by 15:1 in Experiment 2, and by 16:1 in Experiment 3).

The Rescorla-Wagner (1972) model's prediction for V_A arises from two core assumptions: (1)

that V is updated by the algorithm shown in Equation 1, and (2) that Vs sum, such that $V_{AB} = V_A + V_B$. A number of more recent associative models have been developed that retain the same learning algorithm, or something functionally equivalent, but which relax the second assumption by using non-linear rules to describe the generalisation of V between CSs and their compounds (e.g., Harris, 2006; Harris & Livesey, 2010; McLaren & Mackintosh, 2002; Pearce, 1987, 1994; Wagner, 2003; Wagner & Brandon, 2001). These models anticipate a generalization decrement between AB and A or B, and thus, with respect to the present experiments, these models are not constrained to predict that $V_A = V_{AB} - V_B$. Indeed, most of the models[3] can anticipate that $V_A > V_{AB} - V_B$, and thus avoid the prediction that V_A will be less than the value of V when $r = r_{AB} - r_B$. But while this might seem to save these models from falsification by the present data, it comes at the expense of their usefulness because they are unable to make empirically-testable predictions about the strength of conditioning of A. Indeed, the Bayesian analysis we have used to test the predictions of the Rescorla-Wagner model would also penalise these models heavily for the lack of precision in their predictions. In effect, the analysis takes account of the fact that, when a given result is consistent with the predictions of two models, the model that was more specific in predicting that result is more likely given the evidence than the model that only vaguely predicted the same result. Thus, the very specific (and thus highly falsifiable) predictions made by RET (Gallistel & Gibbon, 2000), or by the Rescorla-Wagner model using the same β for reinforcement and non-reinforcement, are more likely given the present data than the much looser set of

predictions that are otherwise offered by associative models.

In sum, the three experiments reported here have consistently shown that when rats were trained with a compound, AB, what they learned about A could be predicted by the difference in reinforcement rates of AB and B. That is, they appeared to learn about the reinforcement rate of A, r_A, as being equal to $[r_{AB} - r_B]$. This means that the content of their learning is directly related to reinforcement rates. This conclusion is consistent with what has been proposed by rate-based models of conditioning (Gallistel & Gibbon, 2000), and is also consistent with the Rescorla-Wagner (1972) model but only if that model assumes the same learning rate for reinforcement and non-reinforcement.

Reference

Andrew, B. J., & Harris, J. A. (2011). Summation of reinforcement rates when conditioned stimuli are presented in compound. *Journal of Experimental Psychology: Animal Behavior Processes, 37*, 385-393.

Blough, D. S. (1975). Steady state data and a quantitative model of operant generalization and discrimination. *Journal of Experimental Psychology: Animal Behavior Processes, 104*(1), 3-21.

Bush, R. R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review, 58*, 313-323.

Gallistel, C. R. (2009). The importance of proving the null. *Psychological Review, 116*, 439-453.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review, 107*, 289-344.

Harris, J. A. (2006). Elemental representations of stimuli in associative learning. *Psychological Review, 113*, 584-605.

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.

---

[3] The exception to this statement is the Replaced Elements Model (Wagner, 2003; Wagner & Brandon, 2001) which predicts that V_A should be even lower. This occurs because any replacement of A elements will reduce V_A when A is presented alone on test compared to the value acquired by A during AB conditioning.

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., & Livesey, E. J. (2010). An attention modulated associative network. *Learning & Behavior, 38*, 1-26.

Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the prediction of behavior: aversive stimulation* (pp. 9-31). Miami: Miami University Press.

Kehoe, E. J. (1982). Overshadowing and summation in compound stimulus conditioning of the rabbit's nictitating membrane response. *Journal of Experimental Psychology: Animal Behavior Processes, 8*, 313-328.

Mackintosh, N. J. (1971). An analysis of overshadowing and blocking. *Quarterly Journal of Experimental Psychology, 23*, 118-125.

Mackintosh, N. J. (1976). Overshadowing and stimulus intensity. *Animal Learning & Behavior, 4*, 186-192.

McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning & Behavior, 30*, 177-200.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review, 94*, 61-73.

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

Rescorla, R. A. (1983). Effect of separate presentation of the elements on within-compound learning in autoshaping. *Animal Learning and Behavior, 11*, 439-446.

Rescorla, R. A. (2002). Comparison of the rates of associative change during acquisition and extinction. *Journal of Experimental Psychology: Animal Behavior Processes, 28*, 406-415.

Rescorla, R. A., & Freberg, L. (1978). The extinction of within-compound flavor associations. *Learning and Motivation, 9*, 411-427.

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

Spence, K. W. (1936). The nature of discrimination learning in animals. *Psychological Review, 43*, 427-449.

Wagner, A. R. (1969). Stimulus selection and a "modified continuity theory.". In G. H. Bower & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 3, pp. 1-41). New York: Academic Press.

Wagner, A. R. (1981). SOP: a model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: memory mechanisms* (pp. 5-47). Hillsdale, NJ: Erlbaum.

Wagner, A. R. (2003). Context-sensitive elemental theory. *Quarterly Journal of Experimental Psychology, 56B*, 7-29.

Wagner, A. R., & Brandon, S. E. (2001). A componential theory of Pavlovian conditioning. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories.* (pp. 23-64). Mahwah NJ, USA: Lawrence Erlbaum Associates, Inc.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology, 76*, 177-186.

(Appendix follows)

Appendix 1: Derivation of linear relationship between *V* and *r*

Here we describe how, according to Equation 1, V becomes a linear function of r when β takes the same value during reinforcement and non-reinforcement. Let $n = 1/r$, such that reinforcement occurs every n^{th} trial. The loss of associative strength, ΔV^-, that occurs across all n-1 non-reinforced trials[4], and the eventual increment in strength, ΔV^+, on trial n, are given by

$$\Delta V^- = (-\alpha\beta V_1 + -\alpha\beta V_2 + \ldots + -\alpha\beta V_{n-1})$$

$$\Delta V^+ = \alpha\beta(\lambda - V_n)$$

where V_i is the associative strength on trial i. With extended conditioning, changes in V will approach an equilibrium at which $\Delta V^+ = -\Delta V^-$ (i.e., each increment in V on reinforced trials is matched by the loss of V across non-reinforced trials). Thus,

$$\alpha\beta(\lambda - V_n) = (\alpha\beta V_1 + \alpha\beta V_2 + \ldots + \alpha\beta V_{n-1})$$

After dividing through by αβ, we can rearrange the above equation as

$$\lambda = V_1 + V_2 + \ldots + V_{n-1} + V_n$$

If $\overline{V}$ is the mean of V_i over all n trials, then

$$\lambda = \overline{V}n$$

and therefore,

$$\overline{V} = \lambda r$$

---

[4] If the processes that drive extinction also operate during the reinforced trial, n, as is assumed in other similar models (Bush & Mosteller, 1951; Wagner, 1981), the relationship between $\overline{V}$ and r is no longer linear as derived here, but is instead hyperbolic, even if β takes the same value for both reinforcement and non-reinforcement.