

Separate Training Influences Relative Validity

Rick Mehta

Acadia University, Canada

Jamie-Lynne Dumont, Sharon Combiadakis, and Douglas A. Williams

University of Winnipeg, Canada

An appetitive conditioning experiment with rats assessed the predictions of a new performance-based account of associative learning called the computational comparator hypothesis (Murphy, Baker, & Fouquet, 2001a, 2001b). A between-subjects design was used in which the stimuli A or B were separately trained either as excitors or as inhibitors prior to and during a relative validity treatment. During relative validity training, X was reinforced when presented with A but was not reinforced when presented with B. In test, responding to X in extinction was lower when A or B had been separately trained as excitors than as inhibitors. Thus, contrary to the computational comparator hypothesis, responding to X was affected by more than just inhibitory training of A. Better fits to the data were obtained by Pearce's configural theory (Pearce, 1987, 1994) and the extended comparator hypothesis (Denniston, Savastano, & Miller, 2001) than by the elemental theory of Rescorla and Wagner (1972) or the computational comparator hypothesis.

In a classic experiment by Wagner et al. (1968), rats in a true discrimination (TD) group and in a pseudo-discrimination (PD) group were exposed to a light (X) that was presented in compound with one of two tones (A and B). In the TD group, a shock was administered on trials in which X was presented in compound with A, but no shocks were administered on trials in which X was presented in compound with B (XA+, XB-, where + and - stand for reinforced and unreinforced trials). In the PD group, on the other hand, shocks were administered on half of the XA trials and on half of the XB trials (XA+/-, XB+/-, where +/- stands for 50% reinforcement). The main result came from the test stage. Even though X had been reinforced on 50% of acquisition trials in both conditions, the animals responded more to X after PD training than after TD training. This finding suggested that X's ability to evoke a conditioned response (CR) was dependent on how often it was reinforced in comparison to A, the only other conditioned stimulus (CS) present on reinforced trials in the TD treatment. This early demonstration of what is called the relative validity effect had a profound impact on subsequent theory development by providing support for the assumption that a CS's relative ability to predict reinforcement, not its absolute validity, determines the magnitude of the CR. Similar results demonstrating the generality of the relative validity effect have since been found in pigeon autoshaping (Wasserman, 1974), and human contingency learning (Baker et al., 1993; Van Hamme & Wasserman, 1993; Wasserman, 1990).

Perhaps the most widely cited theoretical account of relative validity was developed by Rescorla and Wagner (1972; hereafter referred to as the R-W model). They suggested that CSs compete for association with the unconditioned stimulus (US) through an error-correction process, which limits the amount of associative strength

The research in this report was supported by a grant from the Natural Sciences and Engineering Research Council (NSERC) to Douglas A. Williams. Correspondence concerning this article should be addressed to Rick Mehta, Psychology Department, Acadia University, 18 University Ave., Wolfville, NS, Canada B4P 2R6 (rick.mehta@acadiau.ca).

acquired by a less valid cue. This competition between cues is formalized in Equation 1:

$$\Delta V = \alpha\beta (\lambda - \Sigma V) \quad (1)$$

In this equation, ΔV is the change in associative strength of an individual CS on a trial and is affected by three variables. The first is a learning rate parameter that represents the salience or associability of the CS (α). This parameter takes a maximum value of 1 when the CS is present and a value of 0 when the CS is absent. The degree of change is also affected by a learning rate parameter whose value is determined by the salience of the US (β). This parameter is assumed to be higher on reinforced trials than on unreinforced trials, and again varies between 0 and 1. Finally, changes in associative strength are influenced by the maximum associative strength that the US can support (λ), which typically takes a value of 0 when the US is absent and a value greater than 0 when the US is present. The difference between the US actually obtained (λ) and the amount of associative strength accrued to all of the cues present on a given trial (ΣV) determines whether the direction of change in associative strength is positive or negative.

Applied to relative validity, the R-W model predicts that A will acquire associative strength at the expense of X because it is a more valid predictor of reinforcement. In the TD group, both A and X are expected to gain associative strength on XA+ trials. Unlike A, however, X loses associative strength on the unreinforced XB- trials. At asymptote, X should have only modest excitatory strength compared to the highly excitatory A, while B should become a weak conditioned inhibitor because it is unreinforced in the presence of the excitatory X. In the PD group, A and B should acquire low levels of associative strength because they are reinforced on half the trials in which they are presented. On the other hand, X should gain more associative strength than A or B because it is reinforced twice as often as A or B. Thus, X is expected to accumulate more associative strength after PD training than after TD training, and this discrepancy explains the observed difference in responding.

In contrast to the R-W model, which postulates that relative validity is caused by a deficit in acquisition, performance-based accounts, such as the original comparator hypothesis (Miller & Matzel, 1988; but see Miller & Schachtman, 1985) and extended comparator hypothesis (Denniston, Savastano, & Miller, 2001), stipulate that organisms form associations between all available CSs present on a given trial and the US. At test, the target CS's associative strength is compared with the associative strength accumulated by other CSs to determine the strength of the CR. According to the original comparator hypothesis, the associative strength of the target CS is compared to the CS with the greatest associative strength. In this case, X would be compared to A. The extended comparator hypothesis allows for comparisons with cues that previously accompanied the target CS (e.g., A, B, and the background cues of the experimental chamber), and are called first-order comparators. In addition, the extended comparator hypothesis suggests that variations in the strength of other cues (called second-order comparators) should also modulate responding to X if they have an impact on the ability of a first-order comparator to evoke a CR. Thus, the extended comparator hypothesis removes the constraint that the associative strength of a target CS is compared with only the CS that has the most associative strength, and allows for an integration of first- and second-order comparators. However, it does not state the exact manner or rule (e.g., averaging, comparing ratios) by which this integration determines performance. With regard to relative validity, both ap-

proaches predict that X will evoke little responding after TD training if it is assumed that A serves as the most important comparator stimulus because it previously accompanied X and was the best predictor of reinforcement. Evidence consistent with performance-based accounts of relative validity comes from experiments in which responding to X increases when A is extinguished after TD training (Cole, Barnet, & Miller, 1995), and experiments in which post-training extinction of A increases the inhibitory properties of B (Blaisdell & Miller, 2001). The first result, increased responding to X after extinguishing A, is readily predicted by the original comparator hypothesis. The latter result is consistent with the extended comparator hypothesis, but not the original comparator hypothesis, because the extended comparator hypothesis assumes that extinction of A will increase the strength of the CR to X, which in turn has implications for responding to B. In this case, A is said to act as a second-order comparator stimulus because its effect on B is presumably mediated through up-modulation of X, which is a first-order companion of B.

Results favourable to performance-based accounts encouraged Baker and colleagues (Baker et al., 2001; Murphy, Baker, & Fouquet, 2001a, 2001b) to develop the computational comparator hypothesis (hereafter referred to as CCH). Unlike other performance-based approaches (e.g., Denniston et al., 2001), the associative strength of the target cue is compared with only the strongest cue presented in acquisition, and thus is analogous to the original comparator hypothesis. The conditional probability of US occurrence in the presence of a given cue, $P(US|C)$, is first determined. The conditional probabilities for all cues are then rank-ordered in order to determine which cue is the strongest, $P(US|S)$. To compare the response strength (RS) of a target CS to the strongest cue, $P(US|C)$ is multiplied with itself and divided by $P(US|S)$. This approach is summarized by the formula:

$$RS = P(US|C) \times P(US|C)/P(US|S) = P(US|C)^2/P(US|S) \quad (2)$$

In the TD treatment, $P(US|A)$ is 1 because A is always reinforced; $P(US|B)$ is 0 because B is never reinforced; and $P(US|X)$ is 0.5 because X is partially reinforced. When the cues are rank-ordered, A becomes the comparator stimulus because it is the strongest cue, $P(US|A) = P(US|S) = 1$. When the RS for each cue is calculated via Equation 2, the values obtained are 1, 0, and 0.25 for A, B, and X. In the PD treatment, the conditional probabilities for A, B, and X are all 0.5 because each of these cues is reinforced on 50% of trials. Because the conditional probabilities are all equal [$P(US|A) = P(US|B) = P(US|X) = P(US|S) = 0.5$], any of the other cues can serve as the comparator stimulus for the target cue. For example, either B or X could serve as the comparator stimulus for A. Similarly, A or B can be the comparator stimulus for X. When the RS for X is calculated by plugging these conditional probabilities into Equation 2, the value obtained is 0.5 [$P(US|X)^2/P(US|S) = 0.5^2/0.5 = 0.5$]. In this case, the RS for X is larger after PD than after TD training (0.5 versus 0.25).

One strength of CCH is that it correctly predicts the magnitude of differences observed in contextual conditioning. Consider the experiments of Murphy et al. (2001b) that examined the relative validity of contextual and discrete cues in an appetitive conditioning procedure. In this series of experiments, rats were first trained with either a TD or PD discrimination. At test, responding during the periods just prior to CS presentation (pre-CS) and during CS presentation was assessed in either the same context (Experi-

ments 1-3) or in a different context (Experiments 2-3). The first measure, pre-CS responding, was used as an assay of contextual conditioning. The two main findings from these experiments were (a) that there was less responding to X after TD training than after PD training, and (b) that there was less contextual conditioning after TD training than after PD training.

In this series of experiments, the R-W model, the original and extended comparator hypotheses, and CCH were able to account for the differences in responding to X, that is, the basic relative validity effect. However, differences in contextual conditioning were most readily anticipated by CCH. For brevity, we do not discuss the detailed predictions of the original or extended comparator hypotheses, but instead restrict our focus to CCH.

An important feature of some comparator theories, such as the extended comparator hypothesis or CCH, is that they account for contextual conditioning by requiring the assumption that the strength of the context is compared to that of discrete cues; that is, a discrete cue serves as the comparator stimulus for the context. According to CCH, contextual conditioning is determined by comparing the strength of the context in relation to the strongest cue. The first value required is the conditional probability of reinforcement in the presence of the context, $P(\text{US}|\text{context})$. To calculate this value, Murphy et al. (2001a, 2001b) rely on two assumptions: (a) that the probability of reinforcement in the presence of the context is less than that for discrete CSs, and (b) that the animals have equal experience with context extinction and discrete cues. Given these assumptions, $P(\text{US}|\text{context})$ is assumed to have the same value in the TD and PD treatments. The next step involves comparing the context to the strongest cue. In the TD condition, the context would be compared to A, the consistently reinforced cue [$P(\text{US}|A) = P(\text{US}|S) = 1$]; in the PD condition, the context would be compared to A, B, or X, all of which have equal conditional probabilities and are of smaller magnitude than that of A in the TD condition [$P(\text{US}|A) = P(\text{US}|B) = P(\text{US}|X) = P(\text{US}|S) = 0.5$]. When the context is compared to the strongest cue, CCH predicts that the RS for the context is lower after TD training than after PD training. The R-W model is less able to predict this finding because it assumes that the associative strength accrued to the context should be small at asymptote in both the TD and PD conditions. Hence, any observed differences in contextual conditioning should be small in magnitude.

In another series of experiments that investigated relative validity using both Pavlovian and instrumental conditioning preparations, Murphy et al. (2001a) provided further evidence for CCH. In these experiments, rats were exposed to one of three treatments. Group TD1 acquired a true discrimination in which three compounds were presented, but only one was reinforced. Rats in this group acquired an XA^- , XB^- , 2XC^+ discrimination, where XC^+ was presented twice as many times as XA^- or XB^- (hence, 2XC^+). This ensured that X was reinforced on 50% of trials. Group TD2 acquired an XA^+ , XB^+ , 2XC^- discrimination, in which XA and XB were reinforced whereas XC was unreinforced. The control group, Group PD, was presented with an $\text{XA}^{+/-}$, $\text{XB}^{+/-}$, $2\text{XC}^{+/-}$ discrimination, in which all three compounds were reinforced on 50% of trials. In this situation, the R-W model predicts the basic relative validity effect, namely that there should be more responding to X after PD training than after TD training. More important is the prediction that there should be more responding to X in Group TD2 than in Group TD1. This prediction arises because, in the TD2 treatment, X is followed by reinforcement at twice the rate of A or of B, and hence is able to acquire more associative strength

than it would otherwise. On the other hand, C should become a conditioned inhibitor that protects X from extinction on unreinforced XC trials. By contrast, in Group TD1, X must compete against the more predictive C for excitatory strength on the reinforced XC+ trials, and loses strength on unreinforced XA- and XB- trials. Thus, according to the R-W model, X should acquire up to twice as much additional associative strength after a TD2 than a TD1 treatment.

Whereas the R-W model predicts large differences between the TD1 and TD2 groups in responding to X, performance-based approaches generally predict no differences between the two treatments because X, which is reinforced on 50% of trials in both conditions, would be compared to C in TD1, or to either A or B in TD2. Since the comparator cues are consistently reinforced, X's associative strength relative to the comparator stimulus remains similar. Accordingly, there should be no differences in responding to X.

Responding to X was similar in the TD1 and TD2 groups, and both groups exhibited less responding to X than the PD group, a result correctly anticipated by performance-based approaches. However, CCH was able to account for one result that the original comparator hypothesis could not. That is, there was more contextual conditioning (as assessed by pre-CS responding) in the PD group than in the TD1 and TD2 groups, which did not differ from one another. In this case, CCH anticipates this result because the probability of reinforcement in the presence of the context is the same across all three conditions. In the PD condition, the context would be compared to any of the three discrete cues, all of which have equal conditional probabilities of 0.5 because they were partially reinforced [$P(US|A) = P(US|B) = P(US|C) = P(US|S) = 0.5$]. In both of the TD conditions, the context would be compared to the strongest cue. In the TD1 condition, the context would be compared to the consistently reinforced C [$P(US|C) = P(US|S) = 1$]; in the TD2 condition, X would be compared to either A or B, both of which are consistently reinforced [$P(US|A) = P(US|B) = P(US|S) = 1$]. In this experiment, the strongest cue is weaker after PD than after either TD1 or TD2 training. Hence, there should be more contextual conditioning observed in the PD group than in the TD1 or TD2 groups. According to Murphy et al. (2001a), the original comparator hypothesis does not anticipate higher contextual conditioning to the PD group because of its assumption that a stimulus will be less likely to elicit a response if a more valid cue is present. In the PD, TD1, and TD2 treatments, X was a stronger predictor of reinforcement than the context; hence, the context should have remained behaviorally silent and there should have been no differences observed between the PD and two TD conditions. Although the original comparator hypothesis could not predict the observed differences in contextual conditioning, the extended comparator hypothesis could. However, Murphy et al. (2001a, p. 66) preferred the account provided by CCH because it could "be viewed as a subset of the mechanisms of the extended comparator theory...[and could therefore provide]... an account that is consistent with that theory but is more parsimonious."

In summary, measures of contextual conditioning provide the strongest evidence suggesting that relative validity may be accounted for by a computational version of a performance-based model. However, the view that CCH provides the best comprehensive account is open to question. For example, it provides no explanation for the finding of Blaisdell and Miller (2001) that B should gain inhibitory strength if A is extinguished after XA+, XB- discrimination training. Hence, further research is required to resolve this issue. Separate training of the stimuli involved in relative validity is another method by

which predictions derived from CCH may be compared with those of the R-W model and other comparator theories. To contrast their predictions, we report an appetitive conditioning experiment using rats as experimental subjects, and head entries into a food magazine as the CR. In this experiment, A or B was separately trained as an excitor (A-Excitor: Y-, YA+; B-Excitor: Y-, YB+) or as an inhibitor (A-Inhibitor: Y+, YA-; B-Inhibitor: Y+, YB-) prior to and during relative validity training (Table 1). At the end of the experiment, animals were presented with X, A, and B in test. In this situation, the R-W model predicts that separate training of A or B should have a large impact on responding to X at test, whereas CCH predicts a modest effect in only one of the conditions examined. The response pattern anticipated by the R-W model will be described first followed by that of CCH.

Table 1
Design of Experiment.

Group	Stages of Experiment		
	Stage 1	Stage 2 (Relative Validity Training)	Test
A-Excitor	Y-, YA+	Y-, YA+, XA+, XB-	X-, A-, B-
A-Inhibitor	Y+, YA-	Y+, YA-, XA+, XB-	X-, A-, B-
B-Excitor	Y-, YB+	Y-, YB+, XA+, XB-	X-, A-, B-
B-Inhibitor	Y+, YB-	Y+, YB-, XA+, XB-	X-, A-, B-

Note. Stimulus Y was a high-pitch tone, and stimuli X, A, and B were counterbalanced as clicker, low-pitch tone, and noise. The + and - denote reinforced and unreinforced trials. In each stage, the trial types were presented in randomized blocks.

As shown in Table 2, the R-W model predicts that separate training of A or B should have a large impact on responding to X, A, and B. In the A-Excitor group, A should have 1λ of associative strength before X is introduced; thus, A should block X from gaining strength on the XA+ trials (Kamin, 1969) of the relative validity discrimination. If X remains neutral, there would be no basis for changes to either X or B on the unreinforced XB- trials. By contrast, if A were trained as an inhibitor (-1λ , A-Inhibitor group), X would receive large increments in associative strength on reinforced XA+ trials, bringing its final strength to 2λ . On the XB- trials, B would then become strongly inhibitory (-2λ) to bring the quantity $\lambda - \Sigma V$ to zero. A similar pattern of results is predicted if B is trained as an excitor or an inhibitor, albeit for different reasons. In the B-Excitor group, B should initially acquire and then maintain 1λ of associative strength as a result of the Y-, YB+ training. On XB- trials, X should then ultimately gain a terminal level of inhibitory strength of -1λ in order to bring the quantity $\lambda - \Sigma V$ to zero. On XA+ trials, reinforcement of A in the presence of the inhibitory X should cause A to become superexcitatory (2λ). Finally, if B is separately trained as an inhibitor (B-Inhibitor group), X should acquire excitatory properties on the unreinforced XB- trials, and block A from becoming excitatory on XA+ trials. To summarize, the R-W model predicts that separate training of A or B as excitors should lead to decreases in responding to X, rela-

tive to separate training of A or B as inhibitors.

Table 2
Predictions from the Rescorla-Wagner (1972) Model.

Relative Validity Training	Terminal Associative Strength of Cues Presented in Test		
	X	A	B
(A-Excitor) Y-, YA+, XA+, XB-	0λ	1λ	0λ
(A-Inhibitor) Y+, YA-, XA+, XB-	2λ	-1λ	-2λ
(B-Excitor) Y-, YB+, XA+, XB-	-1λ	2λ	1λ
(B-Inhibitor) Y+, YB-, XA+, XB-	1λ	0λ	-1λ

In contrast, CCH predicts minimal effects of separate training of A or B on X (Table 3). For each condition, we have calculated the probability of US occurrence given the presence of each cue over the course of the entire experiment (i.e., Stages 1 and 2). Inspection of Table 3 reveals that only separate training of A as an inhibitor is expected to impact responding to X. This prediction arises because A is the strongest cue in all of the other experimental conditions because it is always reinforced, $P(US|A) = P(US|S) = 1$; hence, A's response strength is 1 ($1^2/1 = 1$). But when A is separately trained as an inhibitor, it is unreinforced on the YA- trials of Stages 1 and 2 and reinforced on XA+ trials of Stage 2. Accordingly, its overall conditional probability is one-third, $P(US|A) = 0.33$. Because X's probability of being reinforced remains at 0.5 in all conditions, X becomes the strongest cue in the A-Inhibitor treatment, $P(US|X) = P(US|S) = 0.5$. Under these conditions, X's response strength increases from 0.25 to 0.5 ($0.5^2/0.5 = 0.5$). This prediction contrasts with that of the original and extended comparator hypotheses, which putatively allow for a greater impact of separate training of A or B. As an example, we consider the B-Excitor condition. For both theories, excitatory training of B might be expected to attenuate responding to X even though B should be a less influential comparator stimulus than A. According to the original comparator hypothesis, separate excitatory training of B as an excitor may impact responding to X because X should be weak when compared to A or B (both of which are more strongly associated with the US relative to X), which could make X behaviorally silent. Similarly, the extended comparator hypothesis also allows for separate training of B to impact X. In this case, the model assumes that B should serve as a first-order comparison stimulus for X, and that X will be down-modulated by separate excitatory training of B. Thus, unlike the original and extended comparator hypotheses which allow for separate training of B to influence responding to X, CCH makes the unique prediction that only inhibitory training of A should impact responding to X.

Table 3
Predictions from the Computational Comparator Hypothesis (Baker et al., 2001).

Stimulus	Relative Validity Training							
	A-Excitor		A-Inhibitor		B-Excitor		B- Inhibitor	
	P(US C)	RS	P(US C)	RS	P(US C)	RS	P(US C)	RS
X	0.5	0.25	0.5	0.5	0.5	0.25	0.5	0.25
A	1	1	0.33	0.22	1	1	1	1
B	0	0	0	0	0.67	0.45	0	0

Note. P(US|C) represents the probability of US occurrence given the presence of the cue being evaluated. Response Strength (RS) is calculated by the formula $P(US|C)^2/P(US|S)$, where P(US|S) represents the probability of the US given the presence of the strongest cue and is displayed in boldface font.

Method

Subjects

The experimental subjects were 48 experimentally naïve male rats (*Rattus norvegicus*) of Sprague-Dawley descent (Charles River Canada, Québec, Canada). The animals were 90 days old and weighed 250-275 g upon arrival at the University of Winnipeg. The rats were housed in pairs and were kept in a colony room that operated on a 16-hr light: 8-hr dark cycle. After a 2-week adaptation period, during which they were handled daily, the rats were placed on a food-restricted diet and maintained at 80% of their free-feeding weight for the duration of the experiment. Rats had *ad libitum* access to water in their home cages. Twelve rats were randomly assigned to each of the A-Excitor, A-Inhibitor, B-Excitor, and B-Inhibitor groups.

Apparatus

The apparatus consisted of eight identical chambers (30.8 x 22.0 x 27.5 cm; MED Associates, Georgia, VT) that were enclosed in chests that attenuated light and sound. Each chest contained an exhaust fan that provided masking noise of 70 dB. The side walls and ceiling of each chamber were made of clear Perspex, and the stimulus panel and back wall were made of aluminum. The floor consisted of 18 stainless steel rods (5 mm diameter), running parallel to the front wall, spaced 11 mm apart. A houselight (24 V) located on the back wall, 3.0 cm from the top of the chamber, was illuminated for the entire duration of each experimental session. In the middle of the ceiling was a speaker that delivered three auditory cues: a 0.5 kHz low tone, an 8.0 Hz clicking sound, and a white noise. A sonalert module located 1.0 cm to the right of the speaker was used to present a 2.9 kHz high tone. All auditory stimuli were 8 dB louder than the 70 dB masking noise provided by the ventilation fans. The low tone, clicker, and white noise were counterbalanced and served the stimulus roles of X, A, and B, and the high tone served the stimulus role of Y. In the center of the stimulus panel, 2.0 cm above the grid floor, was a 5.0 x 5.0 cm food tray from which food pellets (45 mg, Formula 21, Bio-Serv, New Jersey, U.S.A.) were delivered as the reinforcer. When a rat placed its head in the food tray, it interrupted an infrared photobeam; this interruption was detected by a 386sx computer equipped with MED-PC software that controlled the experiment.

Procedure

All animals were exposed to a sample of the food pellets in their home cages. They then experienced one 30-min session of magazine-training in which food pellets were delivered independently of behavior on a variable time 30-s schedule. Over the next 15 days (Stage 1), rats underwent one of four training

conditions: A-Excitor (Y-, YA+), A-Inhibitor (Y+, YA-), B-Excitor (Y-, YB+), or B-Inhibitor (Y+, YB-). In each session, there were 20 reinforced and 20 unreinforced trials. A single food pellet was delivered on reinforced trials; none was delivered on unreinforced trials. Over the next 30 sessions (Stage 2), XA+ and XB- trials were added to the Stage 1 discrimination. There were 10 presentations of each trial type, resulting in a total of 40 trials per session. In test, animals were first presented with four warm-up trials from Stage 2: one XA+ trial, one XB- trial, and one trial each of the originally reinforced and unreinforced trials of Stage 1. After the warm-up trials, they received 16 unreinforced presentations of each of X, A, and B in randomized blocks of three trials. The mean intertrial interval (ITI) was 60 s for all stages of the experiment, with ITIs ranging from 20 s to 100 s.

For each trial, the total amount of time that the rat had its head in the food tray was measured in the 10 s period just prior to presentation of the CS (pre-CS) and in the 10 s period during the presentation of the CS. The dependent variable was the corrected head entry duration, defined as the difference between the amount of time the rat had its head in the food tray during the CS and pre-CS (CS period minus pre-CS period). A rejection criterion of 0.05 was used for all statistical analyses.

Results

The data were first examined to determine whether pre-CS durations were similar across groups. First, the pre-CS durations from the last five sessions of Stage 2 were averaged across stimuli, and were analysed with a 2 (Target: A or B) x 2 (Training: Excitor or Inhibitor) between-subjects analysis of variance (ANOVA). Neither of the two main effects nor the interaction attained significance, largest $F(1, 44) = 2.99$. The means for pre-CS responding across the last five sessions, along with their corresponding standard error of the means (*SEM*), were 3.53 s (A-Excitor; *SEM* = 0.65s), 4.42 s (A-Inhibitor, *SEM* = 0.53 s), 4.59 s (B-Excitor, *SEM* = 0.52 s), and 3.44 s (B-Inhibitor, *SEM* = 0.65 s). Similar to the Stage 2 data, the pre-CS durations across the test session were averaged across trials and examined. A 2 x 2 ANOVA on these data found that there were no group differences in pre-CS responding, largest $F(1, 44) = 2.90$. Mean pre-CS durations in test were 2.51 s (A-Excitor, *SEM* = 0.47 s), 2.08 s (A-Inhibitor, *SEM* = 0.48 s), 2.45 s (B-Excitor, *SEM* = 0.38 s), and 1.47 s (B-Inhibitor; *SEM* = 0.31 s). Similar levels of pre-CS responding across groups suggests that contextual conditioning was similar, and that the use of corrected head entry durations as the main dependent variable is appropriate.

The Stage 1 discriminations were readily mastered. Mean corrected head entry durations on the final session of Stage 1 were 3.65 s (YA+) and 0.24 s (Y-) for the A-Excitor group, 4.12 s (Y+) and 0.64 s (YA-) for the A-Inhibitor group, 3.88 s (YB+) and 0.78 s (Y-) for the B-Excitor group, and 3.29 s (Y+) and 0.44 s (YB-) for the B-Inhibitor group. The average *SEM* for these data was 0.39 s.

The acquisition curves for Stage 2 (i.e., the 30 sessions in which the Stage 1 discriminations were joined by XA+ and XB- trials), displayed across six blocks of five sessions, are shown in Figure 1. The reinforced and unreinforced trials retained from Stage 1 are symbolized by O+ and O-. Responding to O+ and O- was analysed with a 2 (Target: A or B) x 2 (Training: Excitor or Inhibitor) x 2 (Stimulus: O+ or O-) x 6 (Block) mixed ANOVA. Target and Training were the between-subjects factors, and Stimulus and Block were the within-subjects factors. This analysis revealed main effects for Block, $F(5, 220) = 25.61$, and Stimulus, $F(1, 44) = 277.60$, and Stimulus x Block, $F(5, 220) = 3.10$, Training x Stimulus, $F(1, 44) = 4.04$, $p < .051$, and Target x Training x Block, $F(5, 220) = 3.55$, interactions. The Stimulus x Block interaction arose because all groups exhibited decreased responding to O- over blocks whereas responding to O+ remained high; and the Training x Stimulus interaction arose because the discrimination of O+ versus O-

was better when A or B were separately trained as excitators than as inhibitors. However, an analysis of Block 6 revealed that all groups responded more to O+ than to O-, $F(1, 44) = 248.91$, and that all groups learned the discrimination equally well, largest $F(1, 44) = 2.61$.

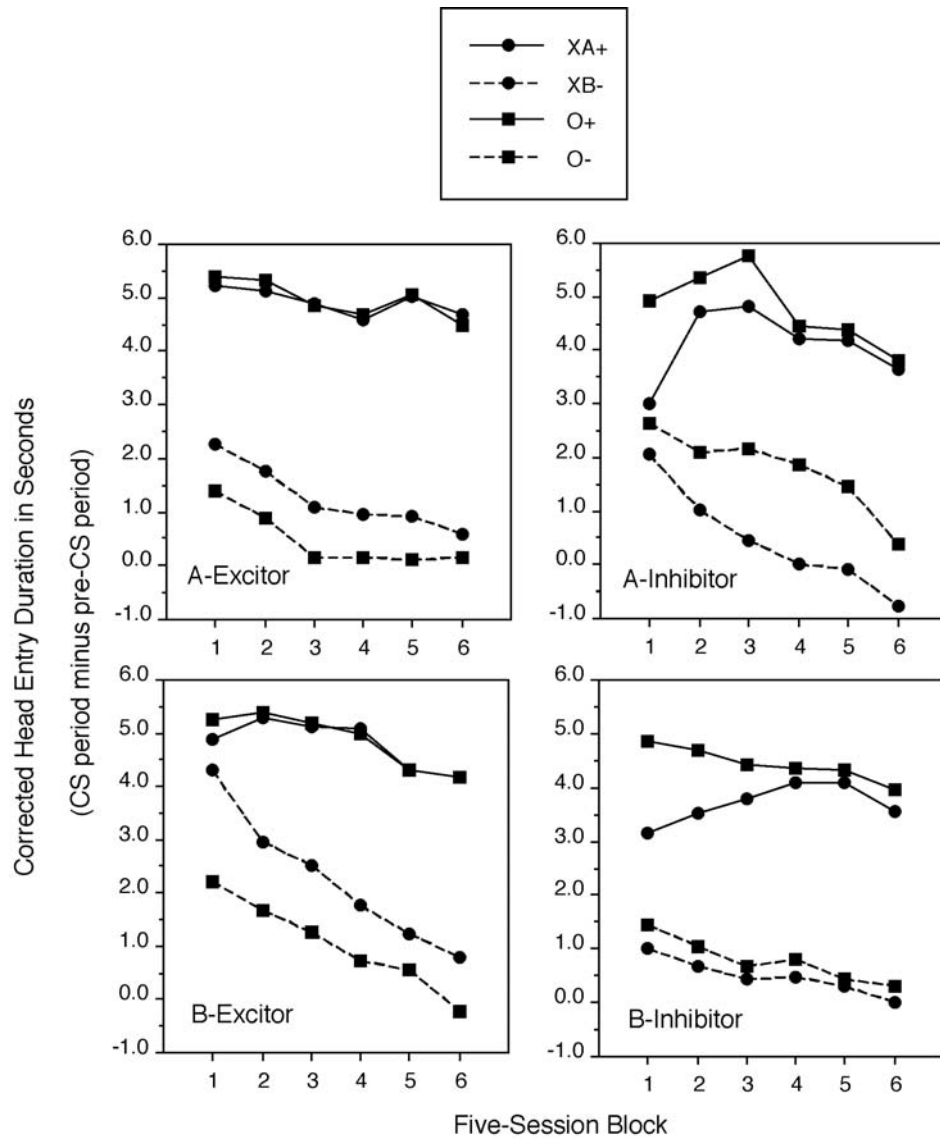


Figure 1. Mean duration of head entries (in seconds), corrected for pre-CS responding (CS period minus pre-CS period), during Stage 2 (Relative Validity Training). Data are averaged across blocks of five sessions. O+ and O- represent the reinforced and unreinforced stimuli that were presented in Stages 1 and 2, whereas XA+ and XB- represent the reinforced and unreinforced stimuli that were unique to Stage 2. Stimulus Y was a high-pitch tone. The clicker, low-pitch tone, and noise were counterbalanced for their stimulus roles as X, A, and B.

Responding to XA+ and XB- was examined with a 2 (Target) x 2 (Training) x 2 (Stimulus: XA or XB) x 6 (Block) ANOVA. The analysis revealed main effects for Training, $F(1, 44) = 5.11$, Block, $F(5, 220) = 13.28$, and Stimulus, $F(1, 44) = 129.25$, as well as Training x Block, $F(5, 220) = 2.93$, Stimulus x Block, $F(5, 220) = 45.54$, Target

x Training x Block, $F(5, 220) = 3.59$, and Target x Training x Stimulus x Block, $F(5, 220) = 7.90$, interactions. The Training x Block interaction attained significance because the A-Excitor and B-Excitor groups initially exhibited higher responding than the A-Inhibitor and B-Inhibitor groups; responding was somewhat higher in the B-Excitor group than in the A-Excitor group, resulting in the Target x Training x Block interaction. Of most interest, however, is the four-way interaction: There were large differences in responding to XA and XB in the first block caused by prior excitatory or inhibitory conditioning of A and B. All groups responded more to XA than to XB, $F(1, 44) = 35.48$; however, this difference was more pronounced for the A-Excitor and B-Inhibitor groups combined, $F(1, 22) = 28.96$, than for the A-Inhibitor and B-Excitor groups combined, $F(1, 22) = 6.92$. Subsequent analyses confirmed that the A-Excitor and B-Inhibitor groups, and that the A-Inhibitor and B-Inhibitor groups, did not differ from each other, all $F_s < 1$. Although there were initial group differences in responding to XA and XB, all groups responded more to XA than to XB in Blocks 2-6; furthermore, the magnitude of this difference was similar across groups, all $F_s < 1$. The average *SEM* for the data depicted in Figure 1 was 0.55 s.




Figure 2. Mean duration of head entries (in seconds), corrected for pre-CS responding (CS period minus pre-CS period), during the test session. Responding is shown to X, A, and B for all groups. The clicker, low-pitch tone, and noise were counterbalanced for their stimulus roles as X, A, and B. Error bars represent standard error.

The main finding in test was that responding to X was impacted by separate training of A and B. These data are summarized in Figure 2, which shows responding to X, A, and B, averaged across the 16 trials of the test session. A 2 (Target) x 2 (Training) x 3 (Stimulus: A, B, or X) x 16 (Trial) mixed ANOVA revealed main effects of Stimulus, $F(2, 88) = 27.51$, and Trial, $F(15, 660) = 2.09$, as well as Stimulus x Trial, $F(30,$

1320) = 1.47, Training x Stimulus, $F(2, 88) = 6.69$, and Training x Stimulus x Trial, $F(30, 1220) = 1.89$, interactions. Of most importance is the Training x Stimulus interaction, which indicates a different pattern of responding to X, A, and B, depending on whether the separately trained stimulus was an excitator or inhibitor. Further examination of the interaction confirmed that responding to X was lower after A or B was separately trained as an excitator (A-Excitor and B-Excitor) than as an inhibitor (A-Inhibitor and B-Inhibitor), $F(1, 44) = 4.73$. There was no impact of the identity of the concurrently trained target (A or B) on responding to X, $F < 1$, nor did Target interact with the effect of excitatory or inhibitory training. As expected, responding to A was higher in the A-Excitor group than in the A-Inhibitor group, $F(1, 22) = 5.72$, whereas responding to B was higher in the B-Excitor group than in the B-Inhibitor group; however, this last difference was nonsignificant when averaged across all 16 trials of the test session, $F(1, 22) = 3.43$, $p < 0.08$. Initially, the presentation of B evoked head entry responding in the B-Excitor group, but responding dropped rapidly, producing a Training x Stimulus x Trial interaction. Averaged over Trials 1-4, responding to B was higher in the B-Excitor group (mean = 2.01 s; $SEM = 0.85$ s) than in the B-Inhibitor group (mean = -0.73 s; $SEM = 0.69$ s), $F(1, 22) = 6.28$, which confirms that B initially evoked different levels of head entry behavior when it was separately trained as an excitator or inhibitor. In summary, the effects of separate training of A or B as an excitator or inhibitor had a large impact on responding to X, but the identity of the trained stimulus (A or B) was unimportant.

Discussion

In this experiment, responding to X in test after XA+, XB- discrimination training was lower when A and B were separately trained as excitators than as inhibitors. This pattern of responding to X is consistent with the prediction of a number of acquisition- and performance-based models that separate training of A or B as excitators or as inhibitors should impact conditioning to X. However, we found no support for the unique prediction of CCH that only separate inhibitory training of A should increase responding to X. These results suggest that both A and B interact with X to determine the magnitude of the CR either directly in acquisition or through a comparison process at test.

Although all but CCH made a correct directional prediction, one might ask which of the remaining theories provided the best overall fit. The data were much less extreme than predicted by the R-W model. For example, the model predicts that X should have become superexcitatory with a final strength of 2λ when A was concurrently maintained as a conditioned inhibitor through nonreinforcement with the excitatory Y (A-Inhibitor group). As shown in Figure 2, X did not exhibit superexcitatory properties by evoking a stronger CR than any other stimulus. It could be argued that our results were preasymptotic, and that with further training X might have acquired superexcitatory properties. This claim seems unlikely for two reasons. First, all groups acquired the Stage 2 discrimination to an equally high level despite initial differences. Second, Pearce and Wilson (1990) have previously demonstrated that a CS reinforced in compound with a concurrent inhibitor does not necessarily endow that CS with superexcitatory properties.

The R-W model makes such extreme predictions because it is an elemental model that views compound stimuli as being a simple additive function of their constituent elements (i.e., $XA = X + A$). Acquisition-based theories that do not subscribe to this harsh assumption make less extreme predictions. Consider the configural theory of

Pearce (1987, 1994). According to Pearce, a stimulus complex acquires associative strength when reinforced, and also receives generalized associative strength from physically similar stimuli. For example, the amount of associative strength accumulated by the XA stimulus complex (ΔE_{XA}) of a relative validity discrimination would be determined by Equation 3:

$$\Delta E_{XA} = \beta (\lambda - V_{XA}) \tag{3}$$

where β is a rate-learning parameter that varies between 0 and 1, λ is the maximum amount of associative strength that the US can support, and V is the associative strength commanded by XA. Finally, V_{XA} represents the sum of a CS's current associative strength arrived at through the application of Equation 3, and the CS's generalized associative strength from physically similar stimuli, namely XB. In a relative validity discrimination, XA receives excitatory associative strength when it is reinforced but also receives generalized inhibition from XB. The amount of generalization received by XA from XB (${}_{XB}e_{XA}$) is determined by Equation 4:

$${}_{XB}e_{XA} = ({}_{XB}S_{XA}) (E_{XB}) \tag{4}$$

Here, the amount of generalized associative strength received from XB is determined by the product of the similarity between XA and XB, which varies between 0 and 1, and by XB's associative strength (E_{XB}).

Table 4
Predictions from the Configural Theory of Pearce (1987).

Relative Validity Training	Terminal Associative Strength of Cues Presented in Test		
	X	A	B
(A-Excitor) Y-, YA+, XA+, XB-	0.29 λ	0.93 λ	-0.10 λ
(A-Inhibitor) Y+, YA-, XA+, XB-	0.51 λ	0.12 λ	-0.17 λ
(B-Excitor) Y-, YB+, XA+, XB-	0.24 λ	0.59 λ	0.44 λ
(B-Inhibitor) Y+, YB-, XA+, XB-	0.45 λ	0.52 λ	-0.38 λ

When Pearce's theory is applied to the experiment reported here, X is expected to receive generalized excitation from the reinforced XA and generalized inhibition from the unreinforced XB. Similar to the R-W model, Pearce's theory predicts that separate training of A or B as excitors should decrease the amount of associative strength received by X relative to separate training of A or B as inhibitors. Computer simulations of the predictions from Pearce's model for our experiment are shown in Table 4. As can be seen, Pearce's model makes the same directional predictions as the R-W model, but the

differences in associative strength are not as large as those predicted by the R-W model. In this instance, the predictions are less extreme because Pearce's model allows for a generalization decrement between compounds and elements, a property that enables this particular acquisition-based account to provide an excellent fit to the data. It would seem that any acquisition-based theory, be it elemental (e.g., McLaren & Mackintosh, 2002; Wagner & Brandon, 2001) or configural (Pearce, 1987), would provide a good fit as long as it included a mechanism for generalization decrement.

We now turn to a consideration of the predictions from performance-based models. Although CCH could not account for our data, one might ask if a modification would increase its predictive power. Does CCH fail to account for our result because (a) the algorithm specified relies on comparing independent conditional probabilities, or (b) it assumes that the associative strength of a CS is compared to the strongest cue? In the first instance, CCH might fail to make the correct qualitative predictions about X because each cue's response strength is independently determined, whereas the cues in the acquisition-based models rely on the notion of cue interaction. With an acquisition-based approach, when associative strength accrues to one CS, it has implications for the associative strength accrued to other CSs. Altering the assumption of CCH to assume cue-interaction during acquisition is inconsistent with the spirit of the theory. Thus, any modification of the theory must focus on the details of the comparison process.

In the second instance, comparing a CS's strength relative to the strongest cue minimizes any effect of separate training because the response strength will not vary as long as its relation to the strongest cue remains unchanged. Of the three comparator models (original comparator hypothesis, extended comparator hypothesis, and CCH), the extended comparator hypothesis is best able to accommodate our results if it is assumed that the associative strength of X is compared to a weighted average of both A and B, rather than only the strongest cue. However, one might question the detailed fit here because the results were not statistically impacted by whether A or B served as the concurrent excitator or inhibitor: To account for the basic relative validity effect, the extended comparator hypothesis must necessarily assume that A plays a larger role than B in the comparison process. Thus, excitatory or inhibitory training of A should have produced a much larger impact than that of B in the current experiment. This last shortcoming, however, should not be read as an endorsement of acquisition-based theories because null results do not provide a firm basis for rejecting a theory.

Some comment should be made about the impact of separate training on A and B themselves. All of the models considered correctly anticipate that responding to the concurrently trained cues should increase or decrease, depending on whether they are trained as excitators or inhibitors. With regard to the effects of A and B on each other, the R-W model predicts that excitatory training of A or B will increase A's associative strength and decrease B's associative strength, whereas Pearce's configural theory and CCH predict that separate training of A will have little impact on B and vice versa (see Tables 2, 3, and 4). Inspection of Figure 2 suggests that the results are consistent with those predicted by Pearce's model and CCH. Responding to A was not influenced by separate training of B, and responding to B was not influenced by separate training of A. Our results could be seen as inconsistent with previous findings. Others have found that extinction of A after a TD treatment led to increased responding to A (Cole et al., 1995) and to increased inhibition to B (Blaisdell & Miller, 2001). Given that the current experiment was not designed to detect both excitation and inhibition (e.g., by use of summation and

retardation tests), it is difficult to evaluate whether our results are truly inconsistent with previous results.

In closing, the present experiment found that separate training of A and B as excitators or inhibitors influenced responding to X. These results suggest that the computational version of a performance-based model described by Murphy et al. (2001a, 2001b) must undergo modification to incorporate cues other than the strongest cue in the comparison process. Further research into the mechanism responsible for B's influence on X might be especially valuable. With no baseline condition in the current experiment, it is unclear whether both excitatory and inhibitory conditioning of B altered responding to X, or perhaps just excitatory conditioning of B. If concurrent inhibitory training of B were to alter responding to X, this would complicate application of performance-based theories to relative validity.

References

- Baker, A. G., Mercier, P., Vallée-Tourangeau, F., Frank, R., & Pan, M. (1993). Selective associations and causality judgments: The presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **19**, 414-432.
- Baker, A.G., Murphy, R.A., Vallée-Tourangeau, F., & Mehta, R. (2001). Contingency learning and causal reasoning. In R.R. Mowrer & S.B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 255-306). Mahwah, NY: Erlbaum.
- Blaisdell, A. P. & Miller, R. R. (2001). Conditioned inhibition produced by extinction-mediated recovery from the relative stimulus validity effect: A test of acquisition and performance models of empirical retrospective reevaluation. *Journal of Experimental Psychology: Animal Behavior Processes*, **27**, 48-58.
- Cole, R. P, Barnet, R. P., & Miller, R. R. (1995). Effect of relative stimulus validity: Learning or performance deficit? *Journal of Experimental Psychology: Animal Behavior Processes*, **21**, 293-303.
- Denniston, J. C., Savastano, H. I., & Miller, R. R. (2001). The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 65-117). Mahwah, NY: Erlbaum.
- Kamin, L. J. (1969). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning*. (pp. 42-64). Halifax, Canada: Dalhousie University Press.
- McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representation: Generalization and discrimination. *Animal Learning & Behavior*, **30**, 177-200.
- Miller R. R, & Matzel L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation: Vol. 22* (pp. 51-92). San Diego, CA: Academic Press.
- Miller, R. R., & Schachtman, T. R. (1985) Conditioning context as an associative baseline: Implications for response generation and the nature of conditioned inhibition. In R.R. Miller & N.E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 51-88). Hillsdale, NJ: Erlbaum.
- Murphy, R. A., Baker, A. G., & Fouquet, N. (2001a). Relative validity effects with either one or two more valid cues in Pavlovian and instrumental conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, **27**, 59-67.
- Murphy, R. A., Baker, A. G., & Fouquet, N. (2001b). Relative validity of contextual and discrete cues. *Journal of Experimental Psychology: Animal Behavior Processes*, **27**, 137-152.
- 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.
- Pearce, J. M., & Wilson, P. N. (1990). Configural associations in discrimination learning. *Journal of Experimental Psychology: Animal Behavior Processes*, **16**, 250-261.
- 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 theory and research* (pp. 64-99). New York: Appleton-Century-Crofts.

Van Hamme, L. J., & Wasserman, E. A. (1993). Cue competition in causality judgments: The role of manner of information presentation. *Bulletin of the Psychonomic Society*, **31**, 457-460.

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 theory* (pp. 301-336). Mahwah, NY: Erlbaum.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, **76**, 171-180.

Wasserman, E. A. (1974). Stimulus-reinforcer predictiveness and selective discrimination learning in pigeons. *Journal of Experimental Psychology*, **103**, 284-297.

Wasserman E. A. (1990). Attribution of causality to common and distinctive elements of compound stimuli. *Psychological Science*, **1**, 298-302.

Received April 22, 2003.

Revision received December 4, 2003.

Accepted December 5, 2003.