Published in final edited form as:

Neurobiol Learn Mem. 2014 February ; 0: 119–135. doi:10.1016/j.nlm.2013.07.018.

The Error in Total Error Reduction

James E. Witnauer¹, Gonzalo P. Urcelay², and Ralph R. Miller³

¹State University of New York at Brockport

²Cambridge University

³State University of New York at Binghamton

Abstract

Most models of human and animal learning assume that learning is proportional to the discrepancy between a delivered outcome and the outcome predicted by *all* cues present during that trial (i.e., total error across a stimulus compound). This total error reduction (TER) view has been implemented in connectionist and artificial neural network models to describe the conditions under which weights between units change. Electrophysiological work has revealed that the activity of dopamine neurons is correlated with the total error signal in models of reward learning. Similar neural mechanisms presumably support fear conditioning, human contingency learning, and other types of learning. Using a computational modelling approach, we compared several TER models of associative learning to an alternative model that rejects the TER assumption in favor of local error reduction (LER), which assumes that learning about each cue is proportional to the discrepancy between the delivered outcome and the outcome predicted by that specific cue on that trial. The LER model provided a better fit to the reviewed data than the TER models. Given the superiority of the LER model with the present data sets, acceptance of TER should be tempered.

Keywords

within-compound associations; cue competition; associative learning; conditioned inhibition; Pavlovian conditioning; total error reduction; local error reduction

1. Introduction

One measure of the importance of a scientific framework is its influence and explanatory potential across different levels of analysis. In this sense, the associative tradition has been highly successful. Associative ideas have influenced neuroscientific theories (e.g., Schultz, 1998), theories of animal learning (e.g., Rescorla & Wagner, 1972), and connectionist models of cognition (e.g., Rumelhart, Hinton, & Williams, 1986). In parallel with the widespread use of the associative framework, total error reduction (TER) is often used to model changes in the associative structures that support behavioral control and cognitive functioning. The TER view asserts that learning is driven by (and functions to reduce) the difference between predicted and actual events. In recent decades, there have been several

© 2013 Elsevier Inc. All rights reserved.

Corresponding author: Ralph R. Miller, Department of Psychology, SUNY – Binghamton, Binghamton, NY 13902-6000, USA, Tel: (607) 777-2291, FAX: (607) 777-4890, rmiller@binghamton.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

reviews of the TER approach to learning (e.g., Gluck & Bower, 1988; Niv & Schoenbaum, 2008; Schultz, 1998). Each of these reviews has presented the TER approach favorably. They often cite neurophysiological correlates of total error signals (e.g., Schultz, Dayan, & Montague, 1997), behavioral tests of TER (e.g., Kamin, 1968), and the widespread use of TER in connectionist modelling (e.g., Nosofsky, Kruschke, & McKinley, 1992). The present review offers a different perspective. We begin by defining TER and reviewing some important applications of TER in behavioral, neuroscientific, and cognitive models. Learning in many influential associative models of Pavlovian conditioning as well as connectionist and neural network models of cognition seems to be fundamentally similar and based on TER mechanisms, which is also consistent with recordings of neural activity (e.g., Schultz & Dickinson, 2000) and responsiveness to local or systemic pharmacological manipulations (e.g., Lattal & Bernardi, 2007; McNally & Westbrook, 2006). In the present review, a computational modelling approach is used to test TER at the neuroscientific, behavioral, and human cognitive levels of analysis. We demonstrate that a simpler local error reduction (LER) model produces a better fit than several TER models with respect to several important sets of data.

1.1 TER in models of human and nonhuman animal cognition

Early models of classical conditioning (e.g., Bush & Mosteller, 1955) predominantly used LER learning algorithms, which assume that both conditioned stimulus (CS)-unconditioned stimulus (US) contiguity and local error are conjointly necessary for learning. Local error is the difference between the magnitude of the US that is received and the strength of the predicted US with the prediction being based on only the specific CS for which the change in associative strength is being calculated even if other CSs were present on that trial. Such models lacked a mechanism for explaining how cues interact with each other in associative learning situations. The first application of the TER approach in psychology was in the animal learning literature and was prompted by early experiments in cue interactions (e.g., blocking). Kamin (1968) found that when a target stimulus (X) is paired with the unconditioned stimulus (US) in the presence of a well-established signal (A) for the US, responding to X is reduced relative to a situation in which X is paired with the US in the presence of a previously neutral stimulus. Kamin postulated that so-called blocking reflects a fundamental limitation of simple contiguity theory because X enjoys equal contiguity with the US in the two conditions. Indeed, Kamin argued that surprise upon receiving the US, in addition to X-US contiguity, is necessary for learning about X. Thus, in blocking A should reduce the unexpectedness (i.e., surprisingness) of the US presented during AX-US pairings, and consequently reduce learning about X. The degraded contingency effect is related to blocking and also contributed largely to the subsequent development of TER models (Rescorla, 1968). In degraded contingency situations, unsignalled presentations of the US reduce responding to X when they are interspersed among X-US pairings relative to a control condition in which the unsignalled USs are omitted. Thus by the late sixties, both the degraded contingency effect and conventional (two-phase forward) blocking challenged existing models of learning by suggesting that good contiguity between a cue and its outcome and local error are not sufficient to explain all instances of associative learning.

Based on Kamin's (1968) discovery of blocking and Rescorla's (1968) discovery of the degraded contingency effect, Rescorla and Wagner (1972) developed the first computational model of TER in psychology. According to the Rescorla-Wagner model, contiguity and surprise (captured here as total error across all cues present on a given trial) are necessary conditions for changes in the strength of an association between X and the US. The critical development in the Rescorla-Wagner model was its error term. In this model, learning is directly related to the difference between the maximum associative strength that the US can support (λ) and the total strength of the US expected by the subject based on the sum of the

associative strengths of *all* stimuli present at the time of each X-US pairing (ΣV). Changes in associative strength function to reduce total error ($\lambda - \Sigma V$). This model also assumes that the saliences (i.e., associability) of X and the US are important determinants of the rate of associative learning. According to the model, the change in the strength of the X-US association on a given trial is:

$$\Delta V_{X-US} = \alpha_X * \beta_{US} * (\lambda_{US} - \Sigma V_{i-US}) \quad (1)$$

where α_X is a free parameter representing the salience of X, β_{US} similarly represents the salience of the US (typically two values: one for US present and a second lower one for US absent), λ_{US} represents the maximum associative strength supportable by the experienced US, and ΣV_{i-US} represents the sum (over all cues present) of the strengths of the CS-US associations. This model asserts that the X-US associative strength is updated trial-by-trial using the following equation:

$$V_{X-US}^{(n+1)} = \Delta V_{X-US}^{(n)} + V_{X-US}^{(n)} \quad (2)$$

where n indexes the last trial completed. In other words, the strength of the X-US association at the beginning of a trial is equal to the associative strength before the most recently completed trial plus the change in the associative strength on the most recently completed trial.

The Rescorla-Wagner model explains both the blocking and degraded contingency effects in a straight forward manner. In blocking situations, when X is paired with the US in the presence of a well-established blocking cue (A), total error ($\lambda - \Sigma V_{i-US}$) is expected to be low relative to a control condition in which a neutral cue replaces the blocking cue because the blocking cue contributes to the expectation of the US, whereas a previously neutral control cue fails to evoke US expectation. Thus, learning about X should be disrupted in the blocking condition because expectation of the US is supported by A. In degraded contingency situations, unsignalled presentations of the US in the training context presumably increase the strength of the context-US association, which then increases US expectation during X-US pairings in the same context and consequently blocks learning about X. The Rescorla-Wagner model has proven to be a powerful model of Pavlovian conditioning that correctly predicts a number of important animal learning phenomena (e.g., Rescorla, 1970, 1971, for reviews, see Miller, Barnet, & Grahame, 1995; Siegel & Allan, 1995). The Rescorla-Wagner model also inspired theoretical and empirical developments across divergent levels of analysis. For example, the Rescorla-Wagner model explains some phenomena within human cognition (e.g., category learning). The observation of blocking in human contingency learning prompted researchers to consider the contribution of total error to learning of contingent relationships by humans (Shanks, 1985). This prompted the development of models of human contingency learning based directly on the Rescorla-Wagner model (e.g., Van Hamme & Wasserman, 1994; detailed below). In the category learning literature, one of the earliest applications of TER was Gluck and Bower's (1988) adaptive network (see also McClelland & Rumelhart, 1985). In their influential paper, Gluck and Bower noted the formal similarity between TER as implemented in the Rescorla-Wagner model and TER in connectionist models of human cognition. While the details of stimulus representation in their adaptive network model differed slightly from that of the Rescorla-Wagner model (largely in order to encompass category learning), the learning rule (TER) was identical. Based on mechanisms similar to those used to explain blocking, this model correctly predicts that participants often neglect baseline rates of categories in judgements of a feature's likelihood of predicting a rare category.

The PDP (Parallel Distributed Processing) research group (Rumelhart et al., 1986) provided important theoretical contributions to cognitive science that were instrumental in the widespread acceptance of connectionist models of cognition. One such contribution was the application of TER to multilayered connectionist models of human cognition. These so-called backpropagation models of cognition are a powerful class of connectionist models that are capable of learning complex, nonlinear (e.g., exclusive-or) input-output mappings. Backpropagation models are characterized by the use of one or more layers of hidden units between input (sensory) and output (response) units. When a type of TER learning algorithm is used, these backpropagation models are capable of learning to solve complex discriminations (e.g., Pearce, 1994). The combination of TER and hidden layers makes these models considerably more powerful and complex than their two-layer connectionist counterparts.

Many backpropagation models are intended to be brain-like in several ways and are therefore sometimes referred to as artificial neural networks. First, units in backpropagation networks are assumed to be neuron-like in terms of their activation functions. The sigmoid activation function that is fundamental to backpropagation models loosely corresponds to the relationship between stimulation and firing rate in real neurons. Second, unlike models of cognition that ordinarily use discrete representation of stimuli (e.g., Gluck & Bower, 1988; Rescorla & Wagner, 1972), stimuli are assumed to be represented by distributed patterns of activation across the hidden layer of the network. This corresponds to the distributed neural representational systems that are presumably used by the brain and rejects the local representational approach. Third, these models assume feedforward processing, meaning that activation of input units spreads through the network as a result of unidirectional weighted connections analogous to the unidirectional flow of activation in an actual neural network. Note that not all applications of the backpropagation approach are artificial neural networks (e.g., Kruschke, 1992), so some backpropagation models make different assumptions about stimulus representations and activation functions. The generalized delta rule is the learning algorithm used in most connectionist systems (Rumelhart et al., 1986). This learning algorithm is used to change the weights of connections in connectionist models. Most important to the present discussion is that the generalized delta rule uses TER and it is at the heart of artificial neural networks and other backpropagation systems.

The generalized delta rule was developed by Rumelhart and his colleagues (1986) as a way to implement TER in connectionist networks with two or more layers. It is used to determine the change in the weights (ΔWeight) for each of the connections between the hidden (j) and output (k) units, which is specified in the following equation:

$$\Delta\text{Weight}_{k-j} = \text{Rate}_{\text{Out}} * \text{Activation}_j * (\text{Teacher}_k - \text{Activation}_k) \quad (3)$$

where Rate_{Out} is a free parameter that determines the learning rate, Activation_j is a variable that represents the activation of the hidden unit j and $(\text{Teacher}_k - \text{Activation}_k)$ is the total error signal for output unit k . Teacher_k is the teaching signal, which represents the desired activation of output unit k (similar to λ in the Rescorla-Wagner model). Activation_k represents the actual activation of the output unit. Because Activation_k is determined by the summed, weighted inputs to output unit k from all of the hidden units, Activation_k is similar to ΣV in the Rescorla-Wagner model. The major challenge that Rumelhart and his colleagues overcame in developing the generalized delta rule was in specifying how error should be calculated and used to modify the weights between the input and hidden nodes in three (or more) layer networks. This situation is important because there is not a teaching signal (i.e., a desired output) for hidden units. Thus, these systems must use total error at the output nodes to modify the weights between input and hidden units. Because hidden units do

not directly interact with teaching signals, error at the hidden layer is calculated differently from error at the output layer. The total error signal applied to hidden unit j (Error_j) is determined using the following equation:

$$\text{Error}_j = \sum [\text{Weight}_{k-j} * (\text{Teacher}_k - \text{Activation}_k)] \quad (4)$$

where, again ($\text{Teacher}_k - \text{Activation}_k$) is the difference between the actual and desired activation of the output unit k . Critically, to model the propagation of error backwards through the weighted connections in the network, error at the hidden units is determined by summing (across all output units) the associatively weighted error signal. The error signal at the output unit (Error_j) is then used to influence changes in the weights between input and hidden units. Notice that the ideal relationship between error at the output layer and changes in the strength of connections between input and hidden units is not obvious because of the distributed nature of the hidden layer. That is, all input-hidden weights can potentially influence all of the error signals at the output layer. This constitutes a credit assignment problem. How should credit for error at the output layer be assigned in changing the connections between the input and hidden layer? Rumelhart et al. suggest that total error is reiteratively propagated backwards from the output units that generated the error signal to other layers of units. Moreover, total error is weighted by an amount that is proportional to the activation of the hidden unit on the trial that resulted in an error signal.

Critical to the present discussion, the learning rules in backpropagation models are fundamentally similar to the Rescorla-Wagner model in that they use TER. Of course, the Rescorla-Wagner model makes different assumptions about representation and is a simpler example of a TER model, but the use of TER in both models is formally equivalent (Gluck & Bower, 1988) and both capture the psychological intuition that learning is driven by prediction violations, with predictions being derived from all available information on a given trial.

1.2. TER in the brain

One function of midbrain dopamine neurons is often assumed to be related to reward. In recent years, the dominant view has been that these cells are important for implementing TER in reward learning situations. Specifically, dopamine neurons in the substantia nigra and the ventral tegmental area presumably encode a signal for total error. Mathematically, Schultz (2002) represented this as:

$$\text{Dopamine Response} = \text{Reward Occurred} - \text{Reward Predicted} \quad (5)$$

This corresponds to the total error signal in the Rescorla-Wagner and backpropagation models. There are several candidate structures that project onto midbrain dopamine neurons that could encode the reward prediction signal (i.e., ΣV), which is necessary for computing total error. Schultz (1998) argued that the direct inhibitory effect (and indirect excitatory effect through double inhibition) of GABAergic striatal neurons on midbrain dopamine neurons could contribute to the midbrain total error signal. Moreover, dopamine neurons receive inputs from the amygdala, brainstem, and prefrontal cortex, all of which could contribute to the total error signal. Midbrain dopamine neurons project onto both the striatum and the prefrontal cortex, where dopamine efflux is presumed to modulate long-term potentiation (Schultz, 1998).

Notably, the mathematical details of total error, as implemented in most dopamine-based reinforcement learning theories, differ slightly from the original Rescorla-Wagner model and from the simplified model presented above. These models are usually based on the

temporal difference model, which was developed by Sutton and his colleagues (e.g., Sutton 1988) in order to improve the temporal precision of total error calculations. Critically, the temporal difference model is based on the Rescorla-Wagner model, but makes more sophisticated (and complicated) assertions about the way that the total error signal can change within a given trial. The temporal difference model assumes that animals process events in real time, which is surely more realistic than the trial-wise information processing assumed in the Rescorla-Wagner model. Specifically, the temporal difference model assumes that total error is updated continuously throughout a trial, whereas the Rescorla-Wagner model asserts that total error is only updated at the end of the trial. According to the temporal difference model, total error (Error) at time $t + 1$ is given by the following equation:

$$\text{Error}(t+1) = \text{Outcome}(t+1) - \text{Prediction}(t+1) - \text{Prediction}(t) \quad (6)$$

Because predictions are based on all present cues, this model assumes TER (instead of LER). This computation of total error offers the advantage of updating within a trial and is different from the trial-wise computations because it is sensitive to changes in reward prediction that occur within a trial (e.g., Montague, Hyman, & Cohen, 2004; Niv & Schoenbaum, 2008). In most versions of this model, future events are discounted by applying a temporal decay parameter to the future outcome ($\text{Outcome}[t + 1]$) and prediction ($\text{Prediction}[t + 1]$). In most delay conditioning situations, the total error predicted by a temporal difference model will be similar to the simpler, trial-wise computations proposed by Rescorla and Wagner. The predictions of this model diverge from those of the Rescorla-Wagner model concerning the asymptotic distribution of behavioral control across continuously reinforced serial compound stimuli (e.g., A-X-Reward). In this situation, the temporal difference model predicts that the first cue in the sequence (A) will gain the most behavioral control, whereas the Rescorla-Wagner model fails to make predictions about the distribution of associative strength within a serial compound. Experiments designed to test this prediction have generally supported the temporal difference view of learning (e.g., Cole & McNally, 2007).

How does the total error signal influence learning at a neural level? Schultz (1998, 2002) asserted that total error signals at the neural level operate similarly to total error signals in connectionist and artificial neural network systems. Accordingly, changes in the weights (i.e., synaptic efficacy) of cortical-cortical or cortical-striatal synapses should be modelled using the following equation:

$$\Delta \text{Weight}_{pre-post} = \text{Rate} * \text{Activation}_{pre} * \text{Activation}_{post} * \text{Error} \quad (7)$$

where Activation_{pre} represents the activation of the presynaptic unit and Activation_{post} represents the activation of the postsynaptic unit. Notice that this learning rule is formally similar to the Rescorla-Wagner model and the delta rule as implemented in artificial neural networks (Rumelhart et al., 1986). The critical difference in both cases is that Schultz's model uses real-time updating of error (and consequently associative strength), whereas the delta rule and Rescorla-Wagner model use trial-wise updating (see the previous section for differences between the delta rule and Rescorla-Wagner).

Perhaps the most compelling empirical support for Schultz's model comes from the electrophysiological data reported by Schultz and his colleagues (1997; see also Waelti, Dickinson, & Schultz, 2001; Tobler, Dickinson, & Schultz, 2003). The electrophysiological activity of neurons that project from the midbrain (ventral tegmental area and substantia nigra) to the striatum have been well characterized and correspond very closely to the total

error signal in temporal difference models of learning. There are four results that are considered to be supportive of the idea that midbrain-striatum neurons encode total error. First, these neurons emit phasic responses to unpredicted primary rewards. This is explained by the temporal difference model because the value of Error ($t + 1$) should be large based on the reward [Outcome ($t + 1$)] being presented in the absence of expectation based on a CS [Prediction (t)]. Second, neurons fail to respond to predicted (e.g., by an appetitive CS) primary rewards. In this situation, Error ($t + 1$) is presumably low because the reward is predicted by the CS. Consequently, the Outcome ($t + 1$) and Prediction(t) take similar values at asymptote, making the difference between them relatively small. Third, firing rate is phasically depressed upon the omission of a predicted reward (i.e., a negative prediction error). This occurs because Error ($t + 1$) is negative based on values of 0 for Outcome ($t + 1$) and Prediction ($t + 1$) as well as a relatively large value of Prediction (t). Fourth, in a sequence of events that predicts reward, the earliest predictive event (e.g., the onset of an appetitive CS) controls an increase in the firing rate of these units because it is essentially an unpredicted secondary reward. Thus, in this situation Error ($t + 1$) is large because Prediction ($t+1$) is large and Prediction (t) is small. These four critical predictions were confirmed using electrophysiological recordings of dopamine neurons in the midbrains of behaving monkeys (Schultz et al., 1997). Thus, the dopamine system has received attention as a candidate system that could encode a total error signal, which serves to modulate neural communication in a way that is consistent with formal models of learning.

In addition to the application of the TER approach to the study of neurobiological basis of reward learning, the TER approach has also been applied to learning in aversive, fear conditioning situations. Behaviorally, fear conditioning is a relatively well-characterized system. Moreover, the amygdala has been identified as the brain structure that centrally controls emotional behavior. Therefore, it is not especially surprising that the neurobiological basis of fear conditioning has been the focus of extensive theorizing. Fanselow's (1998) model is one influential, neurobiologically informed model of fear conditioning. Central to the present discussion, this model assumes that learning serves to reduce total error across a stimulus compound. Before the development of his model, it was almost universally accepted that long-term potentiation (LTP) in the basolateral nucleus of the amygdala was important for fear conditioning. However, it was unclear how LTP could account for many of the behavioral phenomena of fear conditioning. Temporal contiguity is necessary and sufficient for LTP, but does not seem to be sufficient for learning (or at least the expression of learning) at the behavioral level (e.g., Kamin, 1968). Fanselow's model reconciled the view that LTP in the amygdala (for which contiguity is necessary and sufficient) is necessary for learning with the view that some degree of predictive error is also necessary for activation of US specific amygdala cells. Computationally, this model is similar to the Rescorla-Wagner model because it asserts that conjoint activation of the CS and the US are necessary for learning and includes a TER mechanism.

Fanselow's (1998) fear conditioning model asserts that Pavlovian fear conditioning is driven by the combined action of ascending and descending pain and sensory pathways that ultimately converge in the amygdala. According to this framework, the sensory information corresponding to the CS is processed in the thalamus, which projects to the amygdala. The thalamus-amygdala CS signal corresponds to associability of the CS (α) in the Rescorla-Wagner model. Peripheral information concerning the painful US converges in the dorsal horn of the spinal cord. The nociceptive information received by the dorsal horn corresponds to λ in the Rescorla-Wagner model. This information projects into the raphe nucleus and then into the periaqueductal gray. According to this model, the periaqueductal gray is critical because it both sends information to and receives information from the amygdala. The information sent to the amygdala from the periaqueductal gray corresponds to the total error signal (i.e., $\lambda - \Sigma V$) and to associability of the US (β). That is, the periaqueductal gray sends

information to the amygdala about the magnitude of the perceived US, which is related to both β and λ . The cells in the amygdala project onto opioid receptors in the periaqueductal gray. The amygdala-periaqueductal gray opioid signal is assumed to encode the expectation of the US (i.e., ΣV). The opioid system in the periaqueductal gray presumably diminishes the effect of the US through conditioned analgesic mechanisms. That is, opioid release in the periaqueductal gray (i.e., ΣV) diminishes the effectiveness of the US (β and λ). Thus, the US-related output of the periaqueductal gray (to the basolateral amygdala) is thought to constitute the total error signal. Other models assume that the prefrontal cortex contributes to the generation a total error signal based on a pathway involving periaqueductal gray-thalamus-prefrontal cortex-amygdala projections, which explains why temporary inactivation of the prefrontal cortex reduces the blocking effect (e.g., Furlong, Cole, Hamlin, & McNally, 2010).

Fanselow's (1998) description of the effect of conditioned analgesia on US processing has been supported by studies of the effect of opioid antagonists on fear conditioning. It has been shown that administration of an opioid antagonist (naloxone) enhances simple elemental conditioning produced through CS-US pairings (e.g., McNally, Pigg, & Weidemann, 2004; Young & Fanselow, 1992). Young and Fanselow administered training that consisted of pairings of a CS with no footshock, a low intensity, medium intensity, or high intensity footshock. Orthogonal to the footshock intensity manipulation, they manipulated the administration of naloxone (i.e., subjects received naloxone or saline). Within the saline condition, the CS failed to control the CR when no footshock was administered or when the footshock was of low intensity. Critically, naloxone (relative to saline) increased the CS's behavioral control at low, medium, and high footshock levels. Fanselow's model explains this finding by asserting that the administration of naloxone decreases the conditioned analgesic response. Thus, naloxone increases the effectiveness of the footshock by increasing total error ($\lambda - \Sigma V$) through attenuated conditioned analgesia. Total error is presumably increased by naloxone because it reduces the conditioned analgesic effect (i.e., ΣV). Moreover, naloxone increases context conditioning, decreases the rate of extinction, and modulates cue interactions (McNally et al., 2004). Such studies are consistent with the view that US expectation provokes an analgesic response that diminishes the effectiveness of the US in a way that is identical to the way that the total error signal in the Rescorla-Wagner model operates. Interestingly, Fanselow (1998) has argued that total error reduction (negative feedback) may operate in all neural systems that contribute to associative learning. Consistent with this view, TER mechanisms have been used to explain the function of neural circuits in the cerebellum (controlling eyeblink conditioning [Fanselow, 1998]) and prefrontal cortex (controlling human causal learning [Corlett et al., 2004]), among other systems.

Above we have described data that are widely viewed as supportive of total error driving learning, which functions to reduce total error. We later (in the General Discussion) return to these data and explain why we feel that they fail to make a compelling case for TER mechanisms. Prior to doing so, we examine several types of behavioral situations in which a LER model makes predictions distinctly different from TER models.

2. The error in TER

Given the widespread influence and the seemingly successful empirical tests of the TER approach, it seems reasonable to expect that a TER model of associative learning would better fit the results of neuroscientific, animal learning, and human cognitive experiments than a LER model. However, recent discoveries by animal learning researchers appear to challenge the premises of the TER approach and are consistent with a retrieval-focused, LER model (reviewed by Stout & Miller, 2007; Wheeler & Miller, 2008). The purpose of

the second part of this review was to use a computational modelling approach to contrast TER and local error reduction (LER) approaches as applied to some of these new cognitive behavioral findings. To accomplish this, we compared the statistical fits of models that differed in their assumptions about error correction. The TER models that were considered were based on that proposed by Rescorla and Wagner (1972). As previously mentioned, this model's assumptions about total error allow it to successfully explain many aspects of animal learning (e.g., Kamin, 1968), human learning (e.g., Gluck & Bower, 1988), and neural learning (e.g., McNally et al., 2004). But it is also clear, based on a rather extensive literature, that the basic Rescorla-Wagner model fails in many important ways (reviewed by Miller et al., 1995). In the present simulations, we examined the fit of several contemporary associative models that are based on the TER approach as well as a LER model. The TER models were selected in an effort to maximize the explanatory power of the TER approach. Because of their representational flexibility, we simulated a standard, feedforward backpropagation model. Moreover, because some of the data that were simulated included examples of retrospective revaluation (i.e., a change in behavioral control by a cue as a result of associative revaluation of another cue that was previously presented in compound with the target cue), we simulated a variant of Van Hamme and Wasserman's (1994) model of retrospective revaluation. It has also been argued that positive mediation (e.g., second-order conditioning) through within-compound associations could contribute to some of the phenomena that are often interpreted as problematic for the TER approach (e.g., Esber, Pearce, & Haselgrove, 2009; Pineno, 2007). For example, sometimes training a target stimulus in the presence of an excitatory companion stimulus increases excitatory responding to the target stimulus (positive mediation) instead of decreasing excitatory responding or promoting inhibitory responding (negative mediation). Most positive mediation effects are problematic for the TER approach because cues are predicted to compete for limited associative strength; thus, training a target with another cue should reduce learning about the target. In response to this, several models of associative learning have modified the TER approach in an effort to explain positive mediation effects. Consequently, we simulated Pineno's TER model of positive mediation.

Like positive mediation, retrospective revaluation is problematic for traditional TER theories. As previously stated, retrospective revaluation occurs when, after training of a target stimulus as part of a compound of two cues, the response potential of the target stimulus changes as a result of further training (reinforced or nonreinforced) with nontarget stimulus in the compound. Similar to positive mediation, some later variants of the original TER view explain retrospective revaluation. In most situations, both positive mediation and retrospective revaluation contribute to the overall behavior of organisms (despite the opposing nature of these two phenomena). Hence, we additionally constructed and simulated a hybrid of Pineno's and Van Hamme and Wasserman's models of associative learning. This hybrid model would presumably have great flexibility in explaining data and would therefore be a particularly powerful TER model that explains many of the well-known failures of simple TER models (Miller et al., 1995).

The fits of the various TER models were compared to the fit of Stout and Miller's (2007) sometimes competing retrieval (SOCR) model. SOCR is a model of associative learning that accounts for many of the phenomena that the basic TER model fails to explain as well as most of TER's successes. SOCR was selected because it directly rejects TER, unlike most models of Pavlovian conditioning. SOCR is similar to signal detection models of performance (Hannah & Allan, 2011) but assumes that associations are the basic units of information encoded by animals. Thus, in many ways, SOCR is well-matched to most conventional TER models because it assumes an associationist structure, elemental representations of stimuli, and that learning is trial-wise in nature. SOCR is different from TER models in its rejection of TER learning algorithms and its use of a competitive

response rule that is similar in spirit to the performance rules instantiated in signal detection and contingency theory (see below). Thus, SOCR assumes that all cue interactions are driven by information processing that occurs at the moment of testing rather than during acquisition. That is, SOCR assumes that acquisition is noncompetitive (i.e., acquisition occurs identically whether cues are trained together (i.e., in compound) or elementally (Bush & Mosteller, 1955). Instead of a competitive acquisition rule (e.g., Rescorla & Wagner, 1972), SOCR uses a competitive retrieval rule, which assumes that responding is based on the information that is retrieved at test and was stored based on contiguity and LER.

2.1 Simulation Methods

In order to compare the TER models with the SOCR model, we conducted simulations in which we contrasted the fit of different models to data from experiments in Pavlovian conditioning. We compared the models by using an optimization technique (hillclimbing) that searches for the best fitting set of parameters for each of the models that we simulated. This should result in unbiased comparisons between models because this procedure theoretically equates the models with respect to the goodness of the values used for free parameters. We elected to use a different set of parameters (i.e., a new hillclimbing search) for each of the series of experiments that were simulated because, across different series of experiments, procedural and stimulus parameters often varied widely, making it unlikely that a single set of parameters will be appropriate across multiple series of experiments. In contrast, within a series of experiments, in which procedural and stimulus parameters were held constant, model parameters were held constant. This approach assured a fair comparison among models because it avoids the possibility of better parameters for one model than another which exists when parameters are transferred across experiments with nonidentical procedures. Such optimization techniques often suffer from local minima. Towards reducing the likelihood of reporting local minima, we repeated the optimization algorithm using several (4) widely separated starting points. Despite this, we cannot categorically reject the possibility that the optimization algorithms discovered local (instead of global) minima.

2.11 SOCR simulations

The observations by Kaufman and Bolles (1981) and Matzel, Schachtman, and Miller (1985) that the excitatory response potential of the target cue can increase as a function of posttraining associative deflation (i.e., extinction of a stimulus that was previously presented in compound with the target CS; a.k.a. a comparator stimulus) prompted Miller and Matzel (1988) to develop their comparator hypothesis. The extended comparator hypothesis of Denniston, Savastano, and Miller (2001) modified the original comparator hypothesis so that it might be applied to situations in which the target stimulus has multiple comparator stimuli. More recently, Stout and Miller (2007) formalized the extended comparator hypothesis. Their SOCR model of cue interaction makes three critical assumptions: (1) learning is based on a LER mechanism in which cues do not compete for associative strength, (2) cue interaction, including positive and negative mediation, is driven by within-compound associations between the target and so-called comparator stimuli, and (3) responding to a target stimulus is determined by comparison processes that occur at the time of test. Increases in the strength of the association between two stimuli (Stim1 and Stim2; e.g., X and a US, X and A) are modeled using the following equation:

$$\Delta V_{\text{Stim1-Stim2}} = \text{Salience}_{\text{Stim1}} * \text{Salience}_{\text{Stim2}} * (\lambda - V_{\text{Stim1-Stim2}}) \quad (8)$$

where Salience is a free parameter representing salience of the stimulus in question (a function of stimulus intensity among other variables), λ is a fixed parameter (set at 1 by convention) that represents the maximum associative strength supportable on a single trial

by Stim2, and $V_{\text{Stim1-Stim2}}$ is the pretrial strength of the Stim1-Stim2 association. The update equation for SOCR (and other models) in the present simulations takes the familiar form:

$$V_X^{n+1} = \Delta V_X^n + V_X^n \quad (9)$$

Decrements in the strength of the Stim1-Stim2 association occur when Stim1 is presented in the absence of Stim2 and are modeled by the equation:

$$\Delta V_{\text{Stim1-Stim2}} = \text{Salience}_{\text{Stim1}} * -k1 * V_{\text{Stim1-Stim2}} \quad (10)$$

where $k1$ ($0 < k1 < 1$) is a free parameter representing the rate of extinction. Note that these learning rules are simpler than the Rescorla-Wagner learning rule because they are insensitive to the associative status of stimuli that are outside of the pair being modeled (in other words, it uses LER instead of TER). SOCR further assumes that responding to X (R_X) is given by the following equation:

$$R_X = V_{X-US} - k2 * (\sum \text{OP}_{X-i-US} * rV_{X-i} * rV_{i-US}) \quad (11)$$

where $k2$ is a free parameter that weights comparison processes so they have less impact ($0 < k2 < 1$) than the direct CS-US association (i.e., Link 1, the X-US association), OP_{X-i-US} (termed the operator switch) is a variable that represents the degree of discrimination between directly and indirectly activated US representations (used to model positive mediation effects like second-order conditioning with few trials and negative mediation effects like conditioned inhibition with many trials), and r which is explained following Equation 13. rV_{X-i} represents X's potential to activate a representation of a comparator stimulus (i) and rV_{i-US} represents the potential of a comparator stimulus (i) to activate a representation of the US. This response rule captures the psychological intuition that responding to X is determined by a comparison through subtraction of the US representation indirectly activated by X (through the effective X-comparator stimulus [rV_{X-i}] and effective comparator stimulus-US [rV_{i-US}] associative linkage) from the US representation directly activated by X (through the X-US association [V_{X-US}]). Basically, SOCR states that responding to X is not a function of the absolute value of V_{X-US} . Rather it reflects the value of V_{X-US} relative to other cues that were present when X was trained. Thus, the above equations are intended to reflect the principle that behavioral control by a target CS (X) is determined not by the absolute strength of the direct X-US association but by the strength of that association *relative* to the relationship between background cues and the US (rV_{i-US}) where background cues are retrieved through within-compound associations (rV_{X-i}).

Just as first-order comparison processes determine responding to X, the potential of X to activate representations of first-order comparator stimuli (rV_{X-i}) and the potential of comparator stimuli to activate representations of the US (rV_{i-US}) are determined by comparison processes involving second-order (j) comparator stimuli. Second-order comparator stimuli are other stimuli that have associations with first-order comparator stimuli. Specifically, the following equation was used to calculate X's potential to activate comparator stimuli:

$$rV_{X-i} = V_{X-i} - k2 * (\text{Op}_{X-j-i} * V_{X-j} * V_{j-i}) \quad (12)$$

This equation captures the psychological intuition that the effective first-order comparator stimulus representation is determined by the representation of the first-order comparator

activated directly through the association between X and the first-order comparator (V_{X-i}), compared to the representation of the first-order comparator activated through higher-order linkage (i.e., V_{X-j} and V_{j-i}). The strength of the effective indirectly activated representation of the US (rV_{i-US}) is determined by similar higher-order comparator processes:

$$rV_{i-US} = V_{i-US} - k2 * (Op_{i-j-US} * V_{i-j} * V_{j-US}) \quad (13)$$

Note that constraining $k2$ to be less than unity assures that higher order comparator effects are always weaker than lower order comparator effects. Importantly, the operator switch (Op_{X-i-US}) changes with experience (i.e., number of trials) in activating a representation of the US. Changes in the operator switch are modeled using the equation:

$$\Delta Op_{X-i-US} = k3 * Salience_X * V_{X-i} * V_{i-US} * (1 - Op_{X-i-US}) \quad (14)$$

where $k3$ ($0 < k3 < 1$) is the rate with which changes occur in the operator switch. Note that this equation applies only when $V_{X-US} = 0$. Otherwise, $\Delta Op_{X-i-US} = 1 - Op_{X-i-US}$. This rule for the operator switch captures the psychological intuition that subjects must learn to discriminate between directly activated representations and indirectly activated representations unless X is directly paired with the US. Initial values of operator switches were -1 , which allowed for facilitative interactions among stimuli (e.g., second-order conditioning).

2.12 Rescorla-Wagner simulations

For the Rescorla-Wagner (1972) model, all associative changes (ΔV) that were simulated (including both CS-US and within-compound associations) were modeled using the following equation:

$$\Delta V_{Stim1-Stim2} = \alpha_{Stim1} * \beta_{Stim2} * (\lambda_{Stim2} - \sum V_{i-Stim2}) \quad (15)$$

where Stim1 and Stim2 represent two different contiguous stimuli, α_{Stim1} represents the associability (or salience, see below for equations) of Stim1, β_{Stim2} represents the associability of Stim2 (see below), and the quantity $(\lambda_{Stim2} - \sum V_{i-Stim2})$ represents the discrepancy between the magnitude of Stim2 that occurred (λ_{Stim2}) and the magnitude of Stim2 that was expected based on all cues present ($\sum V_{i-Stim2}$). To reduce the number of free parameters, all λ 's were assumed to be fixed at 1. In the present simulations, when Stim2 was present:

$$\beta_{Stim2} = Salience_{Stim2} \quad (16)$$

and when Stim2 was absent:

$$\beta_{Stim2} = k1 * Salience_{Stim2} \text{ and } \lambda_{Stim2absent} = 0 \quad (17)$$

where $k1$ is a free parameter ($0 < k1 < 1$) that represented the diminished associability of absent stimuli. This parameter allows the model to anticipate several important phenomena, including the relative stimulus validity effect and the observation that extinction ordinarily requires more trials than acquisition.

2.13 Van Hamme and Wasserman simulations

Van Hamme and Wasserman (1994, Wasserman & Castro, 2005) asserted that the associability of a cue (Stim1) that is absent, but expected based on within-compound associations to a cue which is present, is proportional to the salience (a.k.a. associability) of

the stimulus when it is present, but of negative value. In their model, the strength of within-compound associations is not considered. In the present simulations, we slightly altered the Van Hamme and Wasserman model to better capture the essential psychological intuition that absent stimuli are retrieved through within-compound associations (Witnauer & Miller, 2011b). Thus, in the present simulations, when Stim1 was present,

$$\alpha_{\text{Stim1}} = \text{Salience}_{\text{Stim1}} \quad (18)$$

and when Stim1 was absent:

$$\alpha_{\text{Stim1}} = k2 * \text{Salience}_{\text{Stim1}} * \sum V_{i-\text{Stim1}} \quad (19)$$

where $k2$ is a free parameter ($-1 < k2 < 0$) that represents the diminished (and negative) activation of a retrieved but absent stimulus. $\sum V_{i-\text{Stim1}}$ is the activation of the representation of the absent stimulus induced by within-compound associations. This last mechanism allows the model to explain several critical phenomena, such as differential retrospective revaluation as a function of within-compound associations. Notably, this aspect of the model is different from the original Van Hamme and Wasserman model. However, it is congruent with the spirit of the model, and previous simulation work revealed that it permits explanation of some phenomena that elude the original model without interfering with successes of the original model (Witnauer & Miller, 2011b). Responding in this model was assumed to equal the CS-US associative strength (i.e., $R_X = V_{X-US}$).

2.14 Pineno and Hybrid simulations

Many researchers have asserted that within-compound associations support positive mediation in select situations (e.g., Pineno, 2007; Rescorla, 1982). We chose to add Pineno's (2007) response rule (which was developed to explain second-order conditioning) to the above version of Van Hamme and Wasserman's (1994) revision of the Rescorla-Wagner model in order to create a version of the Van Hamme and Wasserman model capable of explaining positive mediation. Pineno's model was chosen because, when combined with the Van Hamme and Wasserman model, it provides a better fit than many other models based on the Rescorla-Wagner framework (e.g., Witnauer & Miller, 2013). According to the original Pineno model, responding to X is given by:

$$R_X = V_{X-US} + N_X * (V_{X-i} * V_{i-US}) \quad (20)$$

where N_X is a variable representing X's novelty. This equation effectively anticipates the effects of both first-order conditioning (i.e., through V_{X-US}) and second-order conditioning (through $V_{X-i} * V_{i-US}$). However, it is possible that activation of the US could spread through associative networks that are higher than second-order. Thus, we assumed that for each possible order of activation (e.g., through second-, third-, and fourth-order associative structures):

$$\text{Activation}_{US}^n = \sum (k3^n * V_{X-i} * V_{i-j} * \dots * V_{k-US}) \quad (21)$$

where $k3$ ($0 < k3 < 1$) in the Hybrid model ($k2$ in our implementation of the original Pineno model) effectively dampens the activation of higher-order stimuli in a manner similar to that assumed by SOCR. This damping parameter was not present in Pineno's original model, but we reasoned that the model would be a better match to SOCR if this parameter was included. Presumably, if adding the parameter reduced the quality of the model's fit, the hill climbing algorithm we used to identify parameters would find a very high value near unity for the best fitting value of this parameter. Here, n is the order of activation and i , j , and k

are different stimuli (i.e., we did not allow a stimulus to activate itself). $\text{Activation}_{\text{US}}$ sums across all possible permutations (without replacement) of $i, j, \dots k$. The number of associations in each link is determined by n . Responding to X is given by the following equation:

$$R_X = V_{X-\text{US}} + N_X * \sum (\text{Activation}_{\text{US}}) \quad (22)$$

where N_X is the novelty of the stimulus and summation occurs across all orders of activation (i.e., summation across n). This is a linear autoassociator with the addition of novelty (Pineno, 2007) and a free parameter ($k3$) that controls the extent to which activation spreads through the associative network. Decreases in the novelty of a stimulus were modeled by the equation:

$$\Delta N_X = \text{Salience}_X * k4 * (0 - N_X) \quad (23)$$

where $k4$ ($k3$ in the original Pineno model; $0 < k4 < 1$) controls the rate of changes of novelty. N_X on Trial 1 is unity and is updated in a manner similar to associative strength (i.e., $N_X^{n+1} = \Delta N_X^n + N_X^n$). The novelty variable is an important addition in that it facilitates the model addressing the observation that second-order conditioning is observed with relatively few A+ / AX- trials and conditioned inhibition is observed with relatively many of the same trials (e.g., Stout, Escobar, & Miller, 2004). The Hybrid model also assumes that changes in the strengths of CS-US associations proceed slower than changes in the strength of within-compound associations. This assumption is based on the differences in contiguity between the simultaneous pairings responsible for within-compound associations and the ordinarily sequential pairings of CSs with USs. Thus, we assumed that changes in the strength of CS-US associations were weighted by a factor of $k5$ ($0 < k5 < 1$) in the hybrid model ($k4$ in the original Pineno model).

$$\Delta V_{\text{Stim1-US}} = k5 * \alpha_{\text{Stim1}} * \beta_{\text{US}} * (\lambda_{\text{US}} - \sum V_{i-\text{US}}) \quad (24)$$

2.15 Backpropagation simulations

The backpropagation model was like those simulated previously (Rumelhart et al., 1986). Its architecture consisted of three layers: input, hidden, and output. The hidden layer size varied across simulations. In some, we assumed that the number of hidden units matched the input layer. In others, we directly compared the performance of a backpropagation model with four units to a model with two hidden units. Also, when simulations were particularly hardware intensive, simulations included only two units. The purpose for varying the number of units in the hidden layer was to find the optimal size for the hidden layer. In practice, this had little effect. The output layer had one unit that encoded expectation of the US. The network learned through gradient descent on error (e.g., Rumelhart et al., 1986), which is similar to the Rescorla-Wagner learning rule. The hidden layer used a sigmoidal activation function and the output unit's activation was linear. The model was implemented using Matlab's Neural Network Toolbox®. Inputs to the model were binary, with absent stimuli being represented by zero and present stimuli being represented by one. Each trial was represented by a single iteration of input to the network. That is, we did not allow the model to rehearse each input pattern until it reached an error criterion.

2.2 Scaling

In the present simulations, we were interested in differences between models' predictions and empirical data. Consequently, we needed to scale the predictions of the models to match

the scales used in the empirical research. The critical measurement in some of the present simulations was the $\log_{10}$ latency to consume water for five cumulative seconds in the presence of the test CS. Following Larrauri and Schmajuk (2008), we assumed that lick suppression multiplied by a scaling factor is equal to the response potential of X (R_X) as predicted by the models. In addition, we assumed that the lick suppression predicted by such proportional scaling is added to the scale's minimum (i.e., $\log_{10} 5$) set by our measuring latency to complete 5 s of drinking in the presence of the CS. Thus, in the present simulations, we approximated the lick suppression (LS) scale using the equation:

$$LS = \log_{10} 5 + (\text{Scaling} * R_X) \quad (25)$$

where Scaling is a free parameter used to approximate the $\log_{10}$ latency scale measured in lick suppression situations.

To approximate the suppression ratio (SR) scale (Annau & Kamin, 1961) used as the critical measure in simulations of lever press suppression situations, the response outputs of the models used in the present simulations were scaled using the following equation:

$$SR = (\text{Scaling} - R_X) \div [(\text{Scaling} - R_X) + (\text{Scaling} - R_{\text{Context}})] \quad (26)$$

where Scaling is a free parameter that was used to represent baseline levels of lever pressing (i.e., in the absence of any behaviorally relevant stimuli). Following Larrauri and Schmajuk (2008), in the experiments simulated using this scaling technique, testing was conducted in a context that was not differentially excitatory, so R_{Context} was assumed to be 0 in all of the suppression ratio simulations. All other simulated scales were assumed to be proportional to the cue's behavioral control such that

$$\text{Behavior} = \text{Scaling} * R_X \quad (27)$$

2.3 Model selection

The Bayesian information criterion (BIC) was calculated for each model's fit. We assumed that:

$$BIC = n * \log (SSE/n) + P * \log (n) \quad (28)$$

where n represents the number of groups that were simulated, SSE represents the sum of the squared differences between observed and predicted group means, and P is the number of free parameters used to fit the data (Waekliem, 2004). Effectively, $n * \log (SSE/n)$ represents a badness-of-fit measurement and the far right term ($P * \log (n)$) represents a penalty for model complexity. The BIC is useful because it controls for differences in the number of free parameters, which was necessary here because we compared models that used different numbers of them. Lower values indicate a better BIC score. Because we optimized the values of cue and outcome salience, we treated these values as free parameters. Arguably, such values are not free parameters because, ideally, they are constrained to be consistent with subjects' perception (e.g., more intense stimuli are constrained to be of higher salience than less intense stimuli). Our decision to treat these values as free parameters influenced the absolute BIC values, but this decision did not influence the differences between models with respect to their fits because all models considered the salience of an equal number of stimuli. Importantly, differences among BIC values must be qualified for several reasons. First, we chose phenomena that seemingly falsify simple TER models, which lack well-developed response rules (e.g., Rescorla &

Wagner, 1972). If we fit an expanded set of phenomena, such models would likely have an advantage in terms of BIC due to their use of a smaller number of free parameters. Second, the models that we simulated were not always well-matched in terms of how much explanatory value each free parameter provided. This might unfairly punish models with more free parameters. Thus, we also report the uncorrected fit (SSE) for each model. Differences among models in SSE usually matched differences in BIC.

Thus, each simulation included a free parameter for each non-counterbalanced stimulus and a scaling parameter. The models differed with respect to the number of free parameters beyond that. SOCR included three free parameters: k_1 [extinction rate], k_2 [comparator weight], and k_3 [operator rate]. Variants of the Rescorla-Wagner model included one (simple Rescorla-Wagner and Backprop), two (Van Hamme and Wasserman), four (Pineno), or five (Hybrid) free parameters in addition to those used to model stimulus salience. Note that, by convention, the Backprop model did not include a US node in the input layer, which resulted in one fewer free parameter (for US salience).

3.1 Simulation 1 – Changes in performance as a function of number of trials

Most TER models predict that behavioral control by a target stimulus after elemental training is a monotonic function of the number of target stimulus-US pairings. Specifically, associative strength is assumed to be a negatively accelerated function of number of trials. This allows these models to anticipate the prototypical acquisition curve in which increments in behavioral control become smaller as a direct function of amount of training. Moreover, the TER approach assumes that cue competition (e.g., overshadowing and blocking) influences learning by reducing both the rate of learning (by decreasing error) and the asymptotic amount of associative strength that can be allocated to a target cue (because the target and nontarget cues compete for a limited associative strength). Thus, the TER approach predicts that cue competition should be independent of the amount of compound training (provided that more than one compound trial is administered). The purpose of Simulation 1a was to compare TER models with SOCR as applied to data concerning the effect of number of trials on acquisition in elemental conditioning. Simulation 1b compared the models' fits to data concerning the role of number of trials in overshadowing.

3.1.1 Simulation 1a—One of the most important challenges to the TER approach concerns the effect of amount of training on behavioral control by a target cue. In simple elemental conditioning, responding often reaches a peak after a relatively small number of target stimulus-US pairings. Further training usually results in a reduction in responding to the target stimulus (Pavlov, 1927). This so called overtraining effect has been observed in many different preparations and has been the subject of some recent empirical work (e.g., Bouton, Frohardt, Sunsay, Waddeil, & Morris, 2008; Urcelay, Witnauer, & Miller, 2012).

In an effort to determine the role of competition between the target stimulus and the training context in the overtraining effect, Urcelay et al. (2012) conducted the experiments summarized in Table 1. In their Experiment 1, subjects received 5, 10, 20, or 50 X-US trials. In Experiment 2, subjects received either 5 or 50 trials followed or not followed by posttraining context extinction which was designed to reduce competition between the training context and the target stimulus at the time of testing (rather than during training). Suppression was strong after 5, 10, or 20 training trials and weak after 50, thereby replicating the overtraining effect. Moreover, responding recovered in Group 50-Ctx Ext as a result of posttraining context extinction, which suggests that competition between the training context and the target stimulus at the time of testing produces (or at least appreciably contributes to) the overtraining effect.

The best fitting predictions and observations for critical groups in Urcelay et al.'s (2012) series of experiments are reproduced in Table 2. Table 3 summarizes the performance of the models in terms of both SSE and BIC. Notably, the BIC and SSE values in Table 3 reflect simulations of all groups in Experiments 1–3, whereas Table 2 summarizes only the critical groups from Experiments 1 and 2. Interestingly, only SOCR and models based on Pineno (2007) were able to explain the basic overtraining effect. The Van Hamme and Wasserman, Rescorla and Wagner, and backpropagation models all anticipate that behavioral control by X will not change appreciably as a result of increasing the number of trials from 20 to 50. This was contradicted by Urcelay et al.'s observations. Among those models that anticipated the basic overtraining effect, only SOCR correctly predicted that associative extinction of the context would result in more responding. According to the Pineno-based models, the context functions to positively mediate responding to X through a process akin to that used to explain second-order conditioning. This should wane as X's novelty decreases; thus, extensive training should result in less responding to X, as the empirical data suggest. However, this account is challenged by the observation that deflation of the context results in more behavioral control by X, because the excitatory status of the context should be positively related to the response potential of X. SOCR explains the basic overtraining effect by assuming that the rate with which X gains behavioral control is higher than the rate with which the context gains the potential to compete with X's response potential (i.e., the context's potential as a comparator stimulus for X). This is due to the context's potential as a comparator stimulus being equal to the product of two fractions that typically vary between 0 and 1. Thus, SOCR assumes that the comparator impact of the context is a sigmoidal function of the number of trials rather than a negatively accelerated function. Based on this assumption, SOCR predicts that the comparator impact of the context should reach asymptote later in training than the target stimulus. Also, SOCR predicts that extinguishing the context increases responding to X because the context competes with X only at the time of testing. Based on the best fitting predictions, SOCR provides a better account of the qualitative pattern of data than did either of the Pineno-based models. In addition, differences among models in terms of SSE and BIC indicate that SOCR provides a better account of the overtraining effect than any of the TER models. Thus, the TER models failed to explain either the basic overtraining effect or the recovery from overtraining observed after posttraining associative deflation of the acquisition context.

3.1.2 Simulation 1b—Simulation 1a showed that the TER approach is challenged by the overtraining effect and the role of the acquisition context in producing the overtraining effect in elemental conditioning. Other researchers have investigated the effect of overtraining on cue competition (e.g., Stout, Arcediano, Escobar, & Miller, 2003; Azorlosa & Cicala, 1998). It is well-documented that cue competition diminishes as a function of the number of training trials (beyond a very small number). This is problematic for basic TER models (e.g., Rescorla & Wagner, 1972) because it implies that cue competition does not reflect competition among cues for a limited amount of asymptotic associative strength, as assumed by the TER approach. An important contribution to this literature is the observation that the training context functions to reduce the effectiveness of competing cues (e.g., an overshadowing cue) when large amounts of training are administered. For example, Stout et al. conducted the experiment summarized in Table 4. In their Experiment 1, all subjects received either compound or elemental training of X and Y, which were trained with many (36) and few (4) trials, respectively. Stout et al. observed an overshadowing deficit in the compound condition only when testing Y. That is, overshadowing was only observed when the stimulus was trained with relatively few reinforced trials; the stimulus that was trained extensively was not overshadowed. Alternatively stated, overtraining and compound training counteracted each other in that the combined effects of the two treatments produced less response attenuation than either treatment alone. In their second experiment, Stout et al.

extinguished the training context subsequent to compound training of X with many trials. Context extinction resulted in a reduction in behavioral control by X, which supports SOCR's assertion that the context reduces the potential of the overshadowing cue to compete with X at the time of testing.

The best-fitting predictions and critical group means are summarized in Table 5 and the fits of each of the models are summarized in Table 3. In Simulation 1b, we applied all models to all data points from Experiments 1 and 2 of Stout et al. (2003) and computed SSE and BIC based on those predictions. In order to simplify presentation of the data, we presented only the best fitting predictions and data from the critical groups. With respect to the best-fitting predictions, SOCR was the only model simulated that could explain all of the critical results of Stout et al. It explained both the basic overshadowing effect (less responding to Y in the compound condition than in the elemental condition) and the lack of overshadowing after many training trials (no difference between responding to X in the elemental and compound conditions). Moreover, SOCR correctly predicted that posttraining extinction of the context would attenuate behavioral control by X following many reinforced compound trials. In contrast, the TER models only predicted either the basic overshadowing deficit or an increase in responding as a function of increasing trial number, but not both. The backpropagation model accurately predicted a slight attenuation in responding to X and Y as a result of posttraining context extinction, which was driven by generalization between the X, Y, and the training context. However, the backpropagation model did not explain the critical effect of increasing the number of trials. Thus, SOCR provided a much better account of the critical group differences observed by Stout et al. In addition, SOCR's BIC and SSE were lower than those of the TER models. In conjunction with the results of Simulation 1a, the results of Simulation 1b challenge the TER approach.

3.2 Simulation 2 – Reinforcing an inhibitory CS in the presence of an excitatory CS

One of the most widely cited failures of the TER approach is the observation that an inhibitor gains more behavioral control than an excitor when they are reinforced in compound (Rescorla, 2000). At least superficially, this is problematic for the TER approach because the Rescorla-Wagner model predicts that stimuli of equal salience will gain the same amount of associative strength when they are conditioned in compound because a common error term drives associative changes for all stimuli. Rescorla used a sophisticated compound testing approach to reveal the differences in associative strength gained by excitors and inhibitors. The critical groups from the magazine approach studies from Rescorla's report are summarized in Table 6. During Phase 1, all subjects received training that established two equivalent excitors (A and C) and two equivalent conditioned inhibitors (B and D), with the physical cues counterbalanced within groups which effectively equated saliences of A, B, C, and D. In Phase 2, subjects received reinforced presentations of A and B in compound in Experiments 1 and 2 and nonreinforced presentations of A and B in compound in Experiment 3. Subjects were tested on both the AD and BC compounds. Critically, any difference in behavioral control by these compounds is attributable to differences between A and B with respect to changes in associative strength during Phase 2 because the compounds are otherwise matched (i.e., Phase 1 training was equated between the two test compounds). This avoids problems in interpretation that would be inherent in testing the critical elements (A and B alone), the most notable of which would be reduced sensitivity to seeing an increase in responding to A after Phase 2 reinforcement because of a ceiling effect. The aim of Simulation 2 was to simulate Rescorla's experiments with SOCR (i.e., a LER model) and the TER models.

The best-fitting predictions and critical observations are reproduced in Table 7 and the models' fits are summarized in Table 3. In the present simulations, we simulated Rescorla's (2000) Experiments 1–3 and Experiment 5 (Experiment 4 was omitted because it was

conducted in a different experimental preparation). In the interest of simplifying presentation of the results, Table 7 only includes data from the critical groups across these experiments. Consistent with the intuition that Rescorla's data are problematic for the TER approach, the Rescorla-Wagner (1972), backpropagation, and Pineno (2007) models performed poorly in terms of both statistical fit and explaining the ordinal relationships among groups. Those models performed poorly because they were unable to predict that B gained more behavioral control than A during AB+ trials. Surprisingly, the other TER models (i.e., Van Hamme & Wasserman and Hybrid) were able to explain all of the critical ordinal differences. The Hybrid model performed better than all other models, despite its being fundamentally TER. The Hybrid model assumes that during AB-US trials, establishment of a within-compound B-A association allows for stronger activation of the US when testing B than does the A-B association that could allow for activation of the US during testing of A. Importantly, Rescorla controlled for within-compound associations in Experiment 5. Despite this, the Hybrid model still manages to explain the results of Experiment 5 by assuming that the inhibitor has a weaker association with the context than does the excitor (due to overshadowing of the inhibitor-context within-compound association by the training excitor). Extinction of the inhibitory context during the intertrial interval reduces responding to the excitor and, to a lesser extent, the inhibitor. That is, retrospective revaluation of the excitor during the intertrial interval is stronger than retrospective revaluation of the inhibitor based on its having stronger association with the context. Of course, this prediction assumes that the context is inhibitory based on it being nonreinforced (extinguished) during the ITI, which is tenuous at best. SOCR also provided a reasonably good qualitative explanation of Rescorla's results based on its rejection of the TER approach. The statistical fits of the models varied widely. The Rescorla-Wagner, backpropagation, and Pineno models provided a poor fit to the data. SOCR and Van Hamme and Wasserman provided a modest fit to the data, and the Hybrid model provided the best overall fit to the data. Thus, Rescorla's data most strongly highlight a need for the TER approach, if retained, to be combined with a nontrivial performance rule (e.g., Pineno, 2007) and an additional mechanism to explain retrospective revaluation (e.g., Van Hamme & Wasserman, 1994). It is also noteworthy that Rescorla proposed a TER alternative to the simple Rescorla-Wagner model that seemingly explains his results. This model (called the constrained model), in which learning is driven by both TER and LER (also see Brandon, Vogel, & Wagner, 2003; LePelley, 2004), anticipates less cue interaction than the Rescorla-Wagner model because learning is determined only in part by a common error term. The model succeeds in explaining situations in which cue competition is less than complete (e.g., Rescorla, 2000). However, other models explain the same phenomenon (and many others) by appealing to performance factors. In SOCR, the strength of predicted competition among cues is controlled by k_2 (the comparator weight). In the hybrid and Pineno models, positive mediation reduces the amount of competition that occurs. Thus, the constrained model predicts Rescorla's results, but we decided not to explore this approach computationally because it is less parsimonious than simple TER and LER models and it does not provide as much explanatory power as the performance rules in SOCR and Pineno. Moreover, early simulations in this project revealed that the constrained rule by itself did not fit much more than the results of Rescorla (2000).

3.3 Simulation 3 - Effects of blocking treatment on the blocking cue

Blocking occurs when a well-established signal for the US disrupts acquisition of behavioral control by the target stimulus when it is present during target CS-US pairings. This phenomenon inspired the TER approach because it suggests that animals fail to learn when the US is not surprising based on all cues present on a given trial. Considerable research has focused on the determinants of blocking, but decidedly less research has focused on the effect of blocking treatment on the blocking cue. According to the TER approach, animals

should not learn about the target cue and no further learning should occur to the blocking cue during the compound trials in blocking situations because the US is fully predicted on the basis of the blocking cue (assuming A-US training in Phase 1 was asymptotic). In contrast, SOCR predicts that the behavioral control by the blocking cue should be reduced as a result of compound training because the target cue should be established as a first-order comparator stimulus to the blocking cue, thereby allowing the two cues to compete. Arcediano, Escobar, and Miller (2004) conducted the experiment summarized in Table 8 in order to test these divergent predictions. Subjects received either blocking treatment (condition Block) in which X was trained in the presence of a well-established signal for the US (A) or a control treatment in which X was trained in the presence of B. Orthogonally, X was either of high or low salience. Based on the TER approach, one would expect that suppression to the blocking cue (A) would be the same across all groups because A should have been asymptotic after Phase 1. Suppression to X should have been weaker when it trained in the presence of A relative to the control groups. In contrast, SOCR predicts that responding to both X and A should be weaker in the Block condition than in the Control condition. Moreover, when X is of relatively high salience, behavioral control by A should be weaker than when X is of low salience. In Simulation 3, we compared SOCR and the TER models with respect to their fits to the data of Arcediano et al.

The best-fitting predictions of Simulation 3 are reproduced in Table 9 and the fit of each model is summarized in Table 3. Inspection of Table 9 reveals that SOCR was the only model simulated that could explain all of the critical observations of Arcediano et al. SOCR anticipated weaker responding to A in the Block than in the Control condition. This difference was stronger after blocking a highly salient stimulus than a stimulus of low salience. Moreover, the model anticipated the basic blocking effect. In contrast, the TER models failed to explain the effect of blocking on the blocking cue. Some of the TER models even failed to anticipate any blocking of X because doing so would force the models to predict the opposite of what was observed with respect to A. In addition, SOCR provided a robust quantitative fit to the data relative to the TER models. Thus, the results of Arcediano et al. challenge the TER view.

3.4 Simulation 4 – Within-compound associations in overshadowing

When a target stimulus is conditioned in compound, animals have the opportunity to learn both direct CS-US associations and within-compound associations. For example, overshadowing refers to a situation in which behavioral control by a target stimulus (X) is attenuated following compound AX-US trials relative to elemental X-US trials. With the exception of attentional models, the majority of associative models assume that variations in conditioned responding to X due to the presence of A are produced by the interaction of X-US and A-US associations. In TER accounts, these two associations conjointly decrease the extent to which the US is surprising during AX-US trials (relative to X-US trials) and decrease learning about A and X compared to when error is based solely on the X-US association. In addition to using information provided by the direct X-US and A-US associations, animals might access the US representation and respond based, at least in part, on the X-A or A-X associations. This possibility has gained considerable theoretical (Dickinson & Burke, 1996; Miller & Matzel, 1988; Pineno, 2007; Van Hamme & Wasserman, 1994) and empirical (Amundson, Witnauer, Pineno, & Miller, 2008) attention. The TER models simulated in the present research differ in important and systematic ways from each other and from SOCR with respect to the emphasis placed on within-compound associations. The Rescorla-Wagner model assumes that within-compound associations are learned but are irrelevant to performance. The backpropagation model does not consider within-compound associations (at least in its current formulation). The Van Hamme and Wasserman model assume that within-compound associations allow for retrieval of and

learning about absent stimuli (e.g., retrospective revaluation). Pineno's model assumes that higher-order activation of the US representation through within-compound associations functions to increase excitatory behavioral control by the target stimulus (e.g., positive mediation), predominantly with a few training trials. Thus, the Hybrid model assumes that both positive mediation and retrospective revaluation are driven, at least partially, by within-compound associations and that simple cue competition and conditioned inhibition phenomena are driven by variations in total error. In contrast to the TER models, SOCR assumes that within-compound associations are critical for all cue interactions, including positive mediation, retrospective revaluation, and cue competition. In Simulation 4, we compared these models' fits to experiments that aimed to illuminate the role of within-compound associations in cue competition (specifically, overshadowing).

3.4.1 Simulation 4a—Studies in human causal learning paradigms sought to study the role of within-compound associations in cue competition and retrospective revaluation (e.g., Aitken, Larkin, & Dickinson, 2001; Melchers, Lachnit, & Shanks, 2004). Their results suggested that within-compound associations contribute to retrospective revaluation but not cue competition. A general weakness of these studies is that the correlation between memory for within-compound associations and the degree of cue competition was used as an index of the extent to which within-compound associations drive cue competition. Aitken et al. manipulated (in addition to measuring) within-compound memory by presenting a distractor between serial compound presentations of cues. Unfortunately, the serial compound trials preclude simulation because the models described here are trial-wise (instead of real time). Amundson et al. (2008) sought to further study the role of within-compound associations by directly manipulating within-compound associations and measuring the effect of doing so on cue competition. Specifically, Amundson et al. conducted the experiment summarized in Table 10. Rats received either overshadowing (AX-US) or elemental (X-US) training. Critically, prior to target training, subjects received a treatment designed to establish an inhibitory X-A association (AC- / XC-) or irrelevant training designed to control for latent inhibition (AD- / XC-). Amundson et al. observed less overshadowing after inhibitory pretraining than after the control treatment presumably because the X-A association is important for overshadowing. In Simulation 4a, we compared the models' fit to the results of Amundson et al.'s experiment.

Table 11 reproduces the best fitting predictions and group means of Amundson et al.'s (2008) Experiment 1 and Table 3 summarizes the fit of the models to Amundson et al.'s results. Notably, while Tables 10 and 11 summarize only data pertaining to Experiment 1, the predictions, SSE, and BIC scores are based on applying the models to both Experiments 1 and 2. We simulated both experiments because increasing the amount of data simulated should decrease (improve) BIC scores and we report only the results of the critical experiment (Experiment 1). The only models that explained the attenuated overshadowing achieved by inhibitory pretraining were those that emphasize within-compound associations (SOCR and the Hybrid model). The Rescorla-Wagner, Van Hamme and Wasserman, Pineno, and backpropagation models were unable to explain the attenuated overshadowing caused by pretraining an inhibitory X-A association, presumably because of their deemphasizing the role of within-compound associations in cue competition. Interestingly, the Hybrid model provided a better fit to the data than SOCR despite its assuming that within-compound associations are irrelevant to cue competition per se. The Hybrid model explained the results by assuming that establishing an X-C-A-US associative chain prior to overshadowing increased behavioral control by X because it allowed for stronger higher-order activation of the US representation. It is unclear why Pineno's original model did not provide a similar explanation of Amundson et al.'s results. Perhaps the hillclimbing procedure discovered local minima that differed markedly from the global minima.

3.4.2 Simulation 4b—Towards determining the role of within-compound associations in overshadowing, Blaisdell, Bristol, Gunther, and Miller (1998) conducted the experiment summarized in Table 12. They tested the prediction that latent inhibition of the putative target stimulus (X)-overshadowing cue (A) within-compound association would attenuate overshadowing. That is, if the X-A association is critical for overshadowing, latent inhibition of X, which should retard acquisition of the X-A association, should attenuate overshadowing. Similarly, if the X-training context association is critical for latent inhibition as is predicted by several models including SOCR, overshadowing of the X-training context association achieved through compound training should attenuate the latent inhibition effect. To test these predictions, subjects received either a CS-preexposure treatment or exposure to an irrelevant control cue and, orthogonally, subjects received either elemental or overshadowing training (i.e., in compound) of the target stimulus. Blaisdell et al. observed that both the CS preexposure and the overshadowing training treatments attenuated behavioral control by the target cue when they were administered alone. However, behavioral control by the target stimulus was greater when the two treatments were administered conjointly than when either treatment was administered alone. Note that this finding is not ubiquitous (e.g., Nakajima, Ka, & Imada, 1999). This suggests that selective attenuation of the expression of the X-A and X-training context associations attenuates the reduction in behavioral control caused by compound training and CS preexposure. The purpose of Simulation 4b was to compare the fit of TER models with that of SOCR as applied to the results of Blaisdell et al.

The best-fitting predictions and group means of Blaisdell et al.'s (1998) Experiment 1 are reproduced in Table 13 and the fits of the models to the data are summarized in Table 3. Similar to Simulation 4a, SOCR explained all of the critical differences observed by Blaisdell et al., including: overshadowing, latent inhibition, less overshadowing after CS preexposure, and less latent inhibition after compound training. Also similar to Simulation 4a, the Rescorla and Wagner, Van Hamme and Wasserman, backpropagation, and Pineno models proved unable to explain the critical results, presumably because of their insufficient emphasis on within-compound associations in forward cue competition. Finally, the Hybrid model was unable to explain the critical results of Blaisdell et al. These conclusions are evidenced by the differences among the models with respect to their fits to the data (Table 3). The critical difference between the Hybrid and SOCR models is that the Hybrid model assumes that within-compound associations are necessary only for retrospective revaluation and positive mediation, whereas SOCR assumes that within-compound associations are also important for cue competition (e.g., overshadowing). Thus, the results of Simulation 4b (in conjunction with those of Simulation 4a) suggest that within-compound associations contribute to cue competition and that data speaking to the role of within-compound associations are problematic for the TER approach.

3.5 Simulation 5

Recent empirical research in our laboratory revealed a surprising and potentially theoretically challenging group of phenomena called counteraction effects (Wheeler & Miller, 2008). Counteraction effects occur when the administration of two cue interaction treatments results in less cue interaction than when either treatment is administered alone. The Blaisdell et al. (1998) study described above demonstrates a counteraction between latent inhibition and overshadowing treatments. Urushihara and Miller (2006) provide a very different example of counteraction. They discovered that US-preexposure counteracts with overshadowing. In their studies, subjects received either simple elemental or compound (overshadowing) training. Orthogonally, subjects received either massive preexposure to the US in the training context or a control treatment. Urushihara and Miller observed a response decrement after either US preexposure or overshadowing relative to simple elemental

training. However, when subjects received both overshadowing and US preexposure, behavioral control by the target cue was augmented relative to either treatment alone. That is, the US preexposure and overshadowing treatments mutually counteracted in their potentials to diminish responding to the target stimulus. Counteraction effects are inconsistent with simple TER models because such models assume that when multiple cues are trained in compound they summate (rather than counteract) in their potentials to influence total error. However, some recent theoretical work suggests that positive mediation (e.g., second-order conditioning) might contribute to counteraction effects (Witnauer & Miller, 2010). That is, models based on the Pineno's (2007) rule of positive mediation can explain some aspects of counteractions.

Importantly, counteraction effects have been observed in both negative mediation (e.g., counteraction between overshadowing and US preexposure; Urushihara & Miller, 2006) and positive mediation (e.g., counteraction between second-order conditioning mediated by the context and second-order conditioning mediated by a discrete training excitator [Witnauer & Miller, 2011a]) situations. While preliminary simulations and those reported in previous papers (Witnauer & Miller, 2010) suggest that counteractions can be explained by assuming that positive mediation is stronger in counteraction situations than in situations where a single cue interaction treatment is administered, it remains unclear whether this account could explain counteraction in both negative and positive mediation situations if the same set of parameters was used. Witnauer and Miller (2012) conducted a series of experiments in which they observed counteraction between two positive mediation treatments and between two negative mediation treatments in the same experiment. Table 14 summarizes the design and results of the critical groups from Witnauer and Miller's (2012) experiments. Within the Extinction condition, subjects received X-US pairings followed by nonreinforced presentations of X either elemental or in compound with an independently trained excitator (Y). All extinction treatments occurred in a context that was physically distinct from the acquisition context. Orthogonal to the manipulation of the extinction stimulus (X vs. XY), subjects either received or did not receive unsignalled US presentations during the intertrial interval that were intended to associatively inflate the context of extinction. When extinction occurred either in the presence of a concurrent, independently trained excitator or in the presence of an excitatory context, behavioral control by X was attenuated relative to simple elemental extinction, which is predicted by TER models. However, when both the excitatory context and compound (XY) extinction treatments were administered, behavioral control was greater than when either treatment alone was administered. Thus, the excitatory context and compound extinction treatments counteracted. Within the Second-Order Conditioning condition, subjects were treated similar to the Extinction condition except that the target stimulus was associatively neutral prior to Phase 2. Thus, excitatory behavioral control by X at test was entirely driven by second-order conditioning to X that was mediated by the excitatory context or the concomitant excitator (Y). In the SOC-Compound-Neutral group, responding was elevated relative to the SOC Control group, indicating that Stimulus Y could indeed mediate second-order conditioning. In the SOC-Elemental-Excitatory group, responding was also elevated relative to the control group, which suggests that the excitatory context could also mediate second-order conditioning. However, when the target stimulus was nonreinforced in the presence of both the excitatory context and Stimulus Y, responding to the target stimulus was attenuated relative to when X was trained with either treatment alone. Thus, counteraction occurs in both extinction (negative mediation) and second-order conditioning (positive mediation) situation. As mentioned, previous simulations have indicated that the TER approach can explain either of these findings in isolation (e.g., Witnauer & Miller, 2010). The present simulations tested whether TER models could explain both.

The best fitting predictions for the critical groups from Witnauer and Miller's (2012) experiments are summarized in Table 15 and the fit of each model is summarized in Table 3. Notably, SOCR was the only simulated model that explained all of the critical ordinal results. SOCR predicted counteraction in both the Extinction and SOC conditions. That is, SOCR predicts that either context inflation or compound extinction results in less behavioral control by the target stimulus than when the two treatments are administered conjointly. Similarly, SOCR predicts that either context inflation or compound training will augment responding to the target stimulus in second-order conditioning; however, both treatments combined are less effective than when either is administered alone. None of the TER models anticipate this pattern of data. Correspondingly, the quantitative fits of the TER models are decidedly worse than SOCR's fit. These simulations indicate that the observation of counteraction in both positive and negative mediation situations is problematic for the TER approach.

4. General Discussion

Several critical things were revealed by the present simulations. First, all of the TER models failed to explain the effect of amount of training in elemental and compound training situations (Simulation 1). Second, the TER approach failed to predict the response reduction to a blocking cue observed after compound training with a blocked cue of high salience (Simulation 3). Third, the TER approach failed to explain the important role of within-compound associations in cue competition situations (Simulation 4). Fourth, the TER approach failed to explain counteraction in both positive and negative mediation situations (Simulation 5). In contrast to the simple response rule used by the TER models simulated here (i.e., $R_x = V_x - US$), SOCR, which is a model based on a nontrivial performance rule and noncompetitive (i.e., LER) acquisition rule, explained all of the results that were problematic for the TER approach. Moreover, among the TER models, the Hybrid model performed better in terms of uncontrolled fit (SSE) than the other TER models as applied to Rescorla's (2000, see Simulation 2) and Amundson et al.'s (2008, see Simulation 4a) results, presumably because, similar to SOCR, it assumes a critical role for within-compound associations in cue interactions in general rather than uniquely during retrospective reevaluation. Notably, simple TER models that did not use within-compound associations to explain cue interactions (the backpropagation and the Rescorla and Wagner models), failed to explain any of the experiments simulated in the present paper. These simulations strongly support the claim that learning is not driven by total error and does not function to reduce total error. Moreover, these simulations support SOCR's interpretation of cue interactions.

The present simulations, in conjunction with the behavioral data reviewed, suggest that TER fails to provide an accurate account of appreciable associative learning and performance. Moreover, the successes of SOCR and, to some extent the Hybrid model, suggest that nontrivial response rules and within-compound associations both contribute to behavioral control by a target stimulus. But one might ask how can these results be reconciled with the neurobiological data supporting the TER view reviewed in the introduction? First, preliminary simulations indicate that SOCR provides a reasonably good fit to some of these data. Specifically, we have simulated the data of McNally et al. (2004) and Waelti et al. (2001). For example, SOCR explains the effect of naloxone on blocking by assuming that naloxone modulates the salience of the US (McNally et al.). Thus, SOCR might be applied to the neurobiology of fear conditioning. Moreover, SOCR provided a good fit to the pattern of firing among dopamine neurons in a blocking design by assuming that dopamine firing reflects establishment of a bias in later retrieval (Waelti et al.). In general, a survey of the empirical literature concerning TER failed to convince the authors that neuroscience has produced data that are categorically incompatible with SOCR. Perhaps a yet to be developed biologically informed alternative to the TER approach (e.g., a biological SOCR) could

explain both the neurobiological data supportive of the TER approach and the behavioral data that are problematic for the TER approach. Indeed, Vinogradova (2001) has argued that some neural circuits in the hippocampus function like a comparator wherein incoming information can be inhibited through a comparison process that considers the expectation of the outcome.

In addition, some researchers (e.g., Berridge, 2012; Berridge & Robinson, 1998) have argued against the conventional view that dopamine neurons encode a total error signal that drives learning. Instead they argue that the observed dopamine signals reflect changes in motivation. Additionally, neurobiological data (e.g., Panlilio, Thorndike, & Schindler, 2007) have occasionally failed to support some predictions of a strict TER view of dopamine function. According to the TER view of dopamine function, dopamine agonists (e.g., many illicit drugs) have the potential to produce abnormally strong habits because they produce a strong total error signal that drives large synaptic changes in forebrain targets (Redish, 2004). Panlilio et al. tested this view by replicating a blocking procedure with cocaine as a US. In Phase 1, rats in a blocking group received presentations of a tone CS accompanied by a self-administered cocaine US. A control group involved similar training without the tone. In Phase 2, rats received compound tone and light presentations followed by the cocaine US. During a subsequent test phase, the light was used as a conditioned reinforcer for a new leverpress response. Importantly, despite the use of a cocaine US (which should not allow for blocking based on its consistent potential to produce a strong total error signal), blocking was observed. Moreover, rats that are genetically modified to be unable to synthesize dopamine still have the ability to learn to prefer a spout that delivers sucrose solution over a spout that delivers water, which reflects simple reward learning without dopamine (Cannon & Palmiter, 2003). Similarly, rats that are genetically modified to lack the potential to deactivate synaptic dopamine (i.e., they lack the dopamine reuptake transporter) have vastly more extracellular dopamine and a commensurate increase in preference for sweet substances. However, such mice do not show faster or stronger reinforcement learning, as would be predicted by the TER approach (Cagniard, Balsam, Brunner, & Zhuang, 2005). Such data, in conjunction with the behavioral data simulated in the present paper, could constitute the foundation of a non-TER neuroscientific model of associative learning. It is our hope that neuroscientists and researchers in human cognition consider the numerous behavioral refutations of the basic TER approach (e.g., Rescorla & Wagner, 1972) and keep an open mind to alternative explanations of the neurobiological data.

One limitation of the present analysis is that we could not fully explore the range of different TER models that are available. Because of the widespread acceptance of the TER view and consequently the many TER models that have been proposed, it is impossible to consider and simulate all versions of it. In an effort to address this limitation, we simulated models that we considered to be the most likely to offer a viable account of the experiments that were simulated (based on other simulations and previous theorizing) with the constraint that they be reasonably well-matched to SOCR (e.g., they are trial-wise implementations of TER and employ a modest number of parameters). We also carefully selected the experiments to be simulated in an effort to address this limitation. Specifically, experiments were chosen that test the TER approach in general and not just specific examples of TER models. Indeed, there are many phenomena that are easily explained by the TER approach; however, these phenomena are also compatible with LER models such as SOCR (reviewed by Stout & Miller, 2007). Our conclusions must also be qualified because many of the present simulations fit data from conditioned suppression preparations. Hence, one might be concerned about generalization to preparations other than conditioned suppression. However, there is good agreement between aversive and appetitive preparations with respect to cue interactions in most situations. Nevertheless, there are well-established differences between the preparations as well (e.g., Jozefowicz, Witnauer, & Miller, 2011; 2012). It

should be noted that we are not particularly wedded to SOCR or the comparator hypothesis. Indeed, some data are problematic for SOCR (e.g., Aitken et al., 2001; Esber et al., 2009) and other simulations have revealed limitations of SOCR (Hannah & Allen, 2011). Certainly, we are not centrally claiming that SOCR provides the best account of associative learning. Rather, we are claiming that the TER approach does not seem to provide an adequate general account of all associative learning and performance.

The TER view asserts that learning functions to reduce the discrepancy between events which actually occur and event predictions based on all cues present on that trial. Many behavioral, neuroscientific, and cognitive models have been proposed that learn by TER. These models are fundamentally similar both in terms of their formal assumptions and in their psychological implications. The merit of this view is seriously weakened in light of the present simulations and in conjunction with empirical results reviewed by others (e.g., Berridge, 2007, 2012; Urcelay & Miller, 2008; Wheeler & Miller, 2008). The intent of the current analysis was to encourage researchers across levels of analysis to develop and test models that do not subscribe to TER. Importantly, just because the Rescorla-Wagner model is consistent with much (not all [see Berridge, 2007]) of that data, does not mean that the data are inconsistent with alternative models. For example, the neural activity observed at training that has been interpreted as reflecting TER could actually reflect preparation for subsequent retrieval or some other process that is compatible with LER.

Acknowledgments

NIMH Grant 33881 provided support for this research. We thank Patricia Di Lorenzo, Ken Kurtz, Mario Laborda, Lisa Mash, Bridget McConnell, Gonzalo Miguez, and Cody Polack for comments on an earlier version of the manuscript.

References

- Aitken MR, Larkin MJ, Dickinson A. Re-examination of the role of within-compound associations in retrospective revaluation of causal judgements. *Quarterly Journal of Experimental Psychology*. 2001; 54B:27–51. [PubMed: 11216299]
- Amundson JC, Witnauer JE, Pineno O, Miller RR. An inhibitory within-compound association attenuates overshadowing. *Journal of Experimental Psychology: Animal Behavior Processes*. 2008; 34:133–143. [PubMed: 18248120]
- Annau Z, Kamin LT. The conditioned emotional response as a function of intensity of the US. *Journal of Comparative and Physiological Psychology*. 1961; 51:128–132.
- Arcediano F, Escobar M, Miller RR. Is stimulus competition an acquisition deficit or a performance deficit? *Psychonomic Bulletin & Review*. 2004; 11:1105–1110. [PubMed: 15875983]
- Azorlosa JL, Cicala GA. Increased conditioning in rats to a blocked CS after the first compound trial. *Bulletin of the Psychonomic Society*. 1988; 26:254–275.
- Berridge KC. The debate over dopamine's role in reward: The case for incentive salience. *Psychopharmacology*. 2007; 191:391–431. [PubMed: 17072591]
- Berridge KC. From prediction error to incentive salience: Mesolimbic computation of reward motivation. *European Journal of Neuroscience*. 2012; 35:1124–1143. [PubMed: 22487042]
- Berridge KC, Robinson TE. What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Research Reviews*. 1998; 28:306–369.
- Blaisdell AP, Bristol AS, Gunther LM, Miller RR. Overshadowing and latent inhibition counteract each other: Support for the comparator hypothesis. *Journal of Experimental Psychology: Animal Behavior Processes*. 1998; 24:335–351. [PubMed: 9679309]
- Bouton ME, Frohardt RJ, Sunsay C, Waddeil J, Morris RW. Contextual control of inhibition with reinforcement: Adaptation and timing mechanisms. *Journal of Experimental Psychology: Animal Behavior Processes*. 2008; 34:223–236. [PubMed: 18426305]

- Brandon SE, Vogel EH, Wagner AR. Stimulus representations in SOP: I theoretical rationalization and some implications. *Behavioural Processes*. 2003; 62:5–25. [PubMed: 12729966]
- Bush, R.; Mosteller, F. Stochastic models of learning. New York: John Wiley and Sons; 1955.
- Cannon CM, Palmiter RD. Reward without dopamine. *Journal of Neuroscience*. 2004; 23:10827–10831. [PubMed: 14645475]
- Cagniard B, Balsam PD, Brunner D, Zhuang X. Mice with chronically elevated dopamine exhibit enhanced motivation, but not learning, for a food reward. *Neuropsychopharmacology*. 2005; 31:1362–1370. [PubMed: 16319913]
- Cole S, McNally GP. Temporal-difference prediction errors and pavlovian fear conditioning: The role of NMDA and opioid receptors. *Behavioral Neuroscience*. 2007; 5:1043–1052. [PubMed: 17907835]
- Corlett PR, et al. Prediction error during retrospective revaluation of causal associations in humans: fMRI evidence in favor of an associative model of learning. *Neuron*. 2004; 44:877–888. [PubMed: 15572117]
- Denniston, JC.; Savastano, HI.; Miller, RR. The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In: Mowrer, RR.; Klein, SB., editors. *Handbook of contemporary learning theories*. Hillsdale, NJ: Erlbaum; 2001. p. 65–117.
- Dickinson A, Burke J. Within-compound associations mediate the retrospective revaluation of causality judgments. *Quarterly Journal of Experimental Psychology*. 1996; 49B:60–80. [PubMed: 8901386]
- Esber GR, Pearce JM, Haselgrove M. Enhancement of responding to A following A+ / AX+ training: Challenges for a comparator-theory of learning. *Journal of Experimental Psychology: Animal Behavior Processes*. 2009; 35:485–497. [PubMed: 19839701]
- Fanselow MS. Pavlovian conditioning, negative feedback, and blocking: Mechanisms that regulate association formation. *Neuron*. 1998; 20:625–627. [PubMed: 9581755]
- Furlong TM, Cole S, Hamlin AS, McNally GP. The role of the prefrontal cortex in predictive fear learning. *Behavioral Neuroscience*. 2010; 124:574–586. [PubMed: 20939658]
- Gluck MA, Bower GH. From fear conditioning to category learning: An adaptive network model. *Journal of experimental Psychology: General*. 1988; 117:227–247. [PubMed: 2971760]
- Hannah SD, Allan LG. The criterion-calibration model of cue interaction in contingency judgements. *Learning & Behavior*. 2011; 39:171–190. [PubMed: 21308428]
- Jozefowicz J, Witnauer JE, Miller RR. Two components of responding in Pavlovian lick suppression. *Learning & Behavior*. 2011; 39:138–145. [PubMed: 21264570]
- Jozefowicz J, Witnauer J, Miller RR. The temporal pattern of responding in conditioned bar-press suppression: The role of the context switch and training mode. *Behavioural Processes*. 2012; 89:239–243. [PubMed: 22178451]
- Kamin, LJ. Attention-like” processes in classical conditioning. In: Jones, MR., editor. *Symposium on the prediction of behavior: Aversive stimulation*. Miami, FL: University of Miami Press; 1968. p. 9–31.
- Kaufman MA, Bolles RC. A nonassociative aspect of overshadowing. *Bulletin of the Psychonomic Society*. 1981; 18:318–320.
- Kruschke JK. ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*. 1992; 99:22–44. [PubMed: 1546117]
- Larrauri JA, Schmajuk NA. Attentional, associative, and configural mechanisms in extinction. *Psychological Review*. 2008; 115:640–676. [PubMed: 18729595]
- Le Pelley ME. The role of associative history in models of associative learning: A selective review and a hybrid model. *Quarterly Journal of Experimental Psychology*. 2004; 57B:193–243. [PubMed: 15204108]
- Matzel LD, Schachtman TR, Miller RR. Recovery of an overshadowed association achieved by extinction of the overshadowing stimulus. *Learning and Motivation*. 1985; 16:398–412.
- McClelland JL, Rumelhart DE. Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*. 1985; 114:159–188. [PubMed: 3159828]

- McNally GP, Pigg M, Weidemann G. Blocking, unblocking, and overexpectation of fear: A role for opioid receptors in the regulation of Pavlovian association formation. *Behavioral Neuroscience*. 2004; 118:111–120. [PubMed: 14979787]
- Medin DL, Shchaffer MM. Context theory of classification learning. *Psychological Review*. 1978; 85:207–238.
- Melchers KG, Lachnit H, Shanks DR. Within-compound associations in retrospective revaluation and in direct learning: A challenge for comparator theory. *Quarterly Journal of Experimental Psychology*. 2004; 57B:25–53. [PubMed: 14690848]
- Miller RR, Barnet RC, Grahame NJ. Assessment of the Rescorla-Wagner model. *Psychological Bulletin*. 1995; 117:363–386. [PubMed: 7777644]
- Miller, RR.; Matzel, LD. The comparator hypothesis: A response rule for the expression of associations. In: Bower, GH., editor. *The psychology of learning and motivation*. Vol. 22. San Diego, CA: Academic Press; 1988. p. 51-92.
- Montague PR, Hyman SE, Cohen JD. Computational roles for dopamine in behavioral control. *Nature*. 2004; 431:760–767. [PubMed: 15483596]
- Nakajima S, Ka H, Imada H. Summation of overshadowing and latent inhibition in rats' conditioned taste aversion: Scapegoat technique works for familiar meals. *Appetite*. 1999; 33:299–307. [PubMed: 10625523]
- Niv Y, Schoenbaum G. Dialogues on prediction errors. *Cell*. 2008; 12:265–272.
- Nosofsky RM, Kruschke JK, McKinley SC. Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1992; 18:211–233.
- Panlilio LV, Thorndike EB, Schindler CW. Blocking of conditioning to a cocaine-paired stimulus: Testing the hypothesis that cocaine perpetually produces a signal of larger-than-expected reward. *Pharmacology, Biochemistry and Behavior*. 2007; 86:774–777.
- Pavlov, IP. *Conditioned Reflexes*. London UK: Oxford University Press; 1927.
- Pearce JM. Similarity and discrimination: A selective review and a connectionist model. *Psychological Review*. 1994; 101:587–607. [PubMed: 7984708]
- Pineno O. A response rule for positive and negative stimulus interaction in associative learning and performance. *Psychonomic Bulletin & Review*. 2007; 14:1115–1124. [PubMed: 18229484]
- Redish AD. Addiction as a computational process gone awry. *Science*. 2004; 306:1944–1947. [PubMed: 15591205]
- Rescorla RA. Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*. 1968; 66:1–5. [PubMed: 5672628]
- Rescorla RA. Reduction in the effectiveness of reinforcement after prior excitatory conditioning. *Learning and Motivation*. 1970; 1:372–381.
- Rescorla RA. Variation in the effectiveness of reinforcement and nonreinforcement following inhibitory conditioning. *Learning and Motivation*. 1971; 2:113–123.
- Rescorla RA. Simultaneous second-order conditioning produces S-S learning in conditioned suppression. *Journal of Experimental Psychology: Animal Behavior Processes*. 1982; 8:23–32. [PubMed: 7057142]
- Rescorla RA. Associative changes in excitors and inhibitors differ when they are conditioned in compound. *Journal of Experimental Psychology: Animal Behavior Processes*. 2000; 26:428–438. [PubMed: 11056883]
- Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH.; Prokasy, WF., editors. *Classical Conditioning: II. Current theory and research*. New York, NY: Appleton-Century Crofts; 1972. p. 64-99.
- Rumelhart, DE.; Hinton, GE.; Williams, GE. Learning internal representations by error propagation. In: Rumelhart, DE.; McClelland, JL., editors. *Parallel distributed processing: Explorations in the microstructure of cognition*. Vol. 1. Cambridge, MA: MIT Press; 1986.
- Schultz W. Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*. 1998; 80:1–27. [PubMed: 9658025]

- Schultz W. Getting formal with reward. *Neuron*. 2002; 36:241–263. [PubMed: 12383780]
- Schultz W, Dayan P, Montague RR. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
- Shanks DR. Forward and backward blocking in human contingency judgement. *Quarterly Journal of Experimental Psychology*. 1985; 37B:1–21.
- Stout SC, Arcediano F, Escobar M, Miller RR. Overshadowing as a function of trial number: Dynamics of first- and second-order comparator effects. *Learning & Behavior*. 2003; 31:85–97. [PubMed: 18450071]
- Stout S, Escobar M, Miller R. Trial number and compound stimuli temporal relationship as joint determinants of second-order conditioning and conditioned inhibition. *Learning & Behavior*. 2004; 32:230–239. [PubMed: 15281395]
- Stout SC, Miller RR. Sometimes competing retrieval (SOCR): A formalization of the extended comparator hypothesis. *Psychological Review*. 2007; 114:759–783. [PubMed: 17638505]
- Sutton RS. Learning to predict by the methods of temporal differences. *Machine Learning*. 1988; 3:9–44.
- Tobler PN, Dickinson A, Schultz W. Coding of predicted reward omission by dopamine neurons in a conditioned inhibition paradigm. *Journal of Neuroscience*. 23:10402–10410. [PubMed: 14614099]
- Urcelay GP, Miller RR. Retrieval from memory. In: Menzel, R., editor. *Learning and memory: A comprehensive reference*. Oxford, UK: Elsevier; 2008.
- Urcelay GP, Witnauer JE, Miller RR. The dual role of the context in postpeak performance deficits resulting from extended training. *Learning & Behavior*. 2012; 40:476–493. [PubMed: 22359182]
- Urushihara K, Miller RR. Overshadowing and outcome-alone exposure counteract each other. *Journal of Experimental Psychology: Animal Behavior Processes*. 2006; 32:253–270. [PubMed: 16834493]
- Van Hamme LJ, Wasserman EA. Cue competition in causality judgements: The role of nonpresentation of compound stimulus elements. *Learning and Motivation*. 1994; 25:127–151.
- Vingradova OS. Hippocampus as comparator: Role of two input and two output systems of the hippocampus in selection and registration of information. *Hippocampus*. 2001; 11:578–598. [PubMed: 11732710]
- Waelti P, Dickinson A, Schultz W. Dopamine responses comply with basic assumptions of formal learning theory. *Nature*. 2001; 412:43–48. [PubMed: 11452299]
- Waekliem DL. Introduction to the special issue on model selection. *Sociological Research Methods*. 2004; 33:167.
- Wasserman EA, Castro L. Surprise and change: Variations in the strength of present and absent cues in causal learning. *Learning & Behavior*. 2005; 33:131–146. [PubMed: 16075834]
- Wheeler DS, Miller RR. Determinants of cue interactions. *Behavioral Processes*. 2008; 78:191–203.
- Witnauer JE, Miller RR. Within-compound associations: Models and data. In: Schmajuk, NA., editor. *Computational Models of Pavlovian Conditioning*. Cambridge University Press; 2010. p. 108–149.
- Witnauer JE, Miller RR. Associative status of the context determines the effectiveness of compound extinction. *Journal of Experimental Psychology: Animal Behavior Processes*. 2012; 40:476–493.
- Witnauer JE, Miller RR. Conditioned suppression is an inverted-U function of US intensity. *Learning & Behavior*. 2013; 41:94–106. [PubMed: 22926983]
- Witnauer JE, Miller RR. Some determinants of second-order conditioning. *Learning & Behavior*. 2011a; 39:12–26. [PubMed: 21264564]
- Witnauer JE, Miller RR. The role of within-compound associations in learning about absent cues. *Learning & Behavior*. 2011b; 39:146–162. [PubMed: 21264569]
- Young SL, Fanselow MS. Associative regulation of Pavlovian fear conditioning: Unconditional stimulus intensity, incentive shifts, and latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*. 1992; 18:400–413. [PubMed: 1402696]

Highlights

- Behavioral and neurological phenomena consistent with total error reduction driving learning are reviewed.
- A model based on local error reduction is seen to also account for these phenomena.
- Several different learning models based on total error reduction and one based on local error reduction are fit to a number of data sets, with local error reduction usually providing the best fit.
- Caution is recommended in embracing the principle of total error reduction.

Table 1

Summary of Urcelay et al.'s (2012) critical groups (Simulation 1a)

Group	Phase 1	Phase 2	Test X
5	5 X+	NA	1.6
10	10 X+	NA	1.6
20	20 X+	NA	1.6
50	50 X+	NA	1.1
5-Ctx Ext	5 X+	480 min in Ctx	1.7
50-Ctx Ext	50 X+	480 min in Ctx	1.7

Note: Ctx = context, Ext = extinction, + = reinforcement with a footshock US. Test scores are mean log latencies to resume drinking in the presence of X.

Table 2

Urceley et al. (2012) group means and best-fitting predictions.

	5	10	20	50	5-Ctx Ext	50-Ctx Ext
Observed	1.64	1.64	1.57	1.06	1.72	1.70
SOCR	1.58	1.69	1.58	1.13	1.77	1.67
Backprop	1.53	1.53	1.53	1.53	1.53	1.53
Hybrid	1.61	1.42	1.45	1.45	1.61	1.45
Pineno	1.60	1.66	1.61	1.46	1.60	1.46
VH&W	1.52	1.52	1.52	1.52	1.52	1.52
RW	1.50	1.50	1.50	1.50	1.50	1.50

Note: SOCR = sometimes competing retrieval. Backprop = backpropagation. Hybrid is a hybrid of the Pineno and Van Hamme and Wasserman models (se text). VH&W = Van Hamme and Wasserman. RW = Rescorla and Wagner. Scores are predicted or observed lick suppression scores in log seconds.

Table 3

Comparison of the fit of models

Model	Sim 1a	Sim 1b	Sim 2	Sim 3	Sim 4a	Sim 4b	Sim 5
SOCR	-65.7 (0.13)	-46.5 (0.51)	-2.7 (3.21)	-221.8 (2 * 10 ⁻⁵)	-105.6 (0.02)	-90.2 (0.02)	-44.2 (0.54)
Backprop	-38.4 (0.83)	-30.7 (1.81)	37.7 (39.75)	-56.6 (0.21)	-91.5 (0.05)	-36.1 (0.49)	-12.1 (3.96)
Hybrid	-29.8 (0.68)	-16.2 (1.88)	-34.5 (0.46)	-56.6 (0.11)	-115.5 (0.01)	-22.7 (0.65)	-9.0 (2.51)
Pineno	-31.0 (0.72)	-18.9 (1.88)	21.1 (9.84)	-50.8 (0.17)	-80.5 (0.05)	-23.2 (0.69)	-11.8 (2.53)
VH&W	-33.5 (0.82)	-24.5 (1.88)	-5.3 (3.19)	-34.5 (0.54)	-93.1 (0.04)	-27.5 (0.65)	-9.7 (3.86)
RW	-37.6 (0.76)	-25.8 (2.06)	40.1 (39.62)	-35.9 (0.56)	-87.0 (0.06)	-28.0 (0.69)	-12.6 (3.86)

Note: Sim = simulation. SOCR = sometimes competing retrieval. Backprop = backpropagation. The Hybrid model was a hybrid of the Pineno and Van Hamme and Wasserman models. VH&W = the Van Hamme and Wasserman model. RW = the Rescorla and Wagner model. Each cell contains the best fitting Bayesian information criterion (BIC) and sum of squared error (SSE, in parentheses) for that model. Lower numbers for BIC (including negative) and SSE indicate a better fit.

Table 4

Summary of Stout et al.'s (2003) critical conditions (Simulation 1b)

Group	Phase 1	Phase 2	Test X	Test Y
Compound	36 AX+ / 4 BY+	NA	2.2	1.6
Elemental	36 X+ / 4 Y+	NA	2.1	2.1
Compound-Ctx Ext	36 AX+ / 4 BY+	480 min in Ctx	0.9	0.8

Note: Ctx = context, Ext = extinction, + = reinforcement with a footshock US. Test scores are mean log latencies (s) to resume drinking in the presence of X or Y. A and B were counterbalanced, as were X and Y. A and B were of higher salience than X and Y. Slashes denote interspersed trials within sessions.

Stout et al. (2003) group means and best-fitting predictions.

Table 5

	Compound-Test X	Compound-Test Y	Elemental-Test X	Elemental-Test Y	Compound Ctx Ext-Test X	Compound Ctx Ext-Test Y
Observed	2.2	1.6	2.1	2.1	0.9	0.8
SOCR	2.1	1.4	1.8	2.0	1.0	1.0
Backprop	2.0	1.8	2.0	1.8	1.4	1.3
Hybrid	1.7	1.5	2.1	1.8	1.7	1.5
Pineno	1.7	1.4	2.1	1.8	1.7	1.4
VH&W	1.7	1.4	2.0	1.9	1.7	1.4
RW	1.7	1.4	2.0	1.9	1.7	1.4

Note: SOCR = sometimes competing retrieval. Backprop = backpropagation. Hybrid is a hybrid of the Pineno and Van Hamme and Wasserman models (see text). VH&W = Van Hamme and Wasserman. RW = Rescorla and Wagner. Scores are predicted or observed lick suppression scores in log s.

Table 6

Design and summary of Rescorla's (2000) critical groups (Simulation 2)

Group	Phase 1	Phase 2	Test AD	Test BC
Reinforce (Exp 1–3)	A+ / C+ / X+ / XB– / XD–	AB+	5.83	9.09
Nonreinforce (Exp 5)	A+ / C+ / X+ / XB– / XD–	AB–	1.52	3.94

Note: A, B, C, and D were counterbalanced in the reinforced condition (Experiments 1 and 2). A was counterbalanced with C and B was counterbalanced with D in the nonreinforced condition (Experiment 3). + was the delivery of a food pellet. – was nonreinforcement. Slashes denote interspersed trials within sessions. Test scores indicate mean magazine entries per minute.

Table 7

Rescorla (2000) group means and best-fitting predictions.

	Reinforce-Test AD	Reinforce-Test BC	Nonreinforced-Test AD	Nonreinforce-Test BC
Observed	5.8	9.1	1.5	3.9
SOCR	5.9	9.3	2.5	4.6
Backprop	7.5	7.5	2.7	2.7
Hybrid	5.8	9.0	1.6	4.0
Pineno	5.6	9.3	1.7	4.1
VH&W	5.7	9.6	0.8	4.6
RW	7.7	7.7	2.9	2.8

Note: SOCR = sometimes competing retrieval. Backprop = backpropagation. Hybrid is a hybrid of the Pineno and Van Hamme and Wasserman models (see text). VH&W = Van Hamme and Wasserman. RW = Rescorla and Wagner. Numbers are predicted or observed mean magazine entries per minute.

Table 8

Design and summary of Arcediano et al. (2003; Simulation 3)

Group	Phase 1	Phase 2	Test A	Test X
Block-High	A+	AX +	1.4	2.1
Control-High	A+	BX +	1.8	2.1
Block-Low	A+	AX+	1.8	1.7
Control-Low	A+	BX+	2.0	2.0

Note: A and B were counterbalanced. Bold indicates that the stimulus was of high salience. Subjects in the High condition were exposed to compound presentations of A or B with a high salience stimulus (**X**). Low indicates that subjects received blocking or control treatment with a low salience stimulus. + indicates reinforcement. Scores indicate mean latencies in log s to resume drinking in the presence of the test stimulus.

Table 9

Arcediano et al. (2003) group means and best-fitting predictions.

	Block-High-Test A	Ctrl-High-Test A	Block-Low-Test A	Ctrl-Low-Test A	Block-High-Test X	Ctrl-High-Test X	Block-Low-Test X	Ctrl-Low-Test X
Observed	1.4	1.8	1.8	2.0	2.1	2.1	1.7	2.0
SOCR	1.3	1.8	1.8	2.0	2.0	2.1	1.7	1.9
Backprop	1.7	1.7	1.8	1.8	2.1	2.1	1.7	1.7
Hybrid	1.6	1.6	1.8	1.9	2.0	2.2	1.8	1.9
Pineno	1.6	1.6	1.8	1.8	2.0	2.1	1.8	1.9
VH&W	1.6	1.6	1.7	1.7	1.7	1.7	1.7	1.7
RW	1.7	1.7	1.7	1.7	1.7	1.7	1.7	1.7

Note: SOCR = sometimes competing retrieval. Backprop = backpropagation. Hybrid is a hybrid of the Pineno and Van Hamme and Wasserman models (see text). VH&W = Van Hamme and Wasserman. RW = Rescorla and Wagner. Scores are predicted or observed lick suppression scores in log s.

Table 10

Design and summary of Amundson et al. (2008, Experiment 1; Simulation 4a).

Group	Phase 1	Phase 2	Test X
Inhib-Compound	AC– / XC–	AX+	.06
Ctrl-Compound	AD– / XC–	AX+	.18
Inhib-Elemental	AC– / XC–	X+	.06
Ctrl-Elemental	AD– / XC–	X+	.09

Note: C and D were counterbalanced. + indicates reinforcement. – represents nonreinforcement. Slashes denote interspersed trials within sessions. Inhib = Inhibitory and Ctrl = Control. Scores indicate mean suppression ratios.

Table 11

Amundson et al. (2008, Experiment 1) group means and best-fitting predictions.

	Inhib-Compound	Ctrl-Compound	Inhib-Elemental	Ctrl-Elemental
Observed	0.06	0.18	0.06	0.09
SOCR	0.05	0.18	0.05	0.07
Backprop	0.09	0.11	0.09	0.11
Hybrid	0.04	0.17	0.06	0.09
Pineno	0.11	0.11	0.07	0.07
VH&W	0.12	0.11	0.07	0.07
RW	0.15	0.15	0.07	0.07

Note: SOCR = sometimes competing retrieval. Backprop = backpropagation. Hybrid is a hybrid of the Pineno and Van Hamme and Wasserman models (see text). VH&W = Van Hamme and Wasserman. RW = Rescorla and Wagner. Scores are predicted or observed mean suppression ratios.

Table 12

Design and summary of Blaisdell et al. (1998, Experiment 1; Simulation 4b).

Group	Phase 1	Phase 2	Test X
Pre-Compound	X–	AX+	1.65
No Pre-Compound	Y–	AX+	0.75
Pre-Elemental	X–	X+	1.12
No Pre-Elemental	Y–	X+	1.64

Note: X and Y were counterbalanced. + indicates reinforcement. – represents nonreinforcement. Pre = CS Preexposure. Scores indicate mean latencies to resume drinking in the presence of X in log s.

Table 13

Blaisdell et al. (1998, Experiment 1) group means and best-fitting predictions.

	Pre-Compound	No Pre-Compound	Pre-Elemental	No Pre-Elemental
Observed	1.65	0.75	1.12	1.64
SOCR	1.65	0.75	1.12	1.75
Backprop	1.25	1.18	1.39	1.38
Hybrid	1.21	1.16	1.55	1.55
Pineno	1.20	1.20	1.54	1.54
VH&W	1.21	1.16	1.55	1.55
RW	1.20	1.20	1.54	1.54

Note: SOCR = sometimes competing retrieval. Backprop = backpropagation. Hybrid is a hybrid of the Pineno and Van Hamme and Wasserman models (see text). VH&W = Van Hamme and Wasserman. RW = Rescorla and Wagner. Scores are predicted or observed lick suppression scores in log s.

Table 14

Design and summary of Witnauer and Miller's (2011a) critical groups from Experiments 2 and 3.

Group	Phase 1	Phase 2	Test X
Ext-Compound-Inflate Ctx	X+ / Y+	XY- / +	2.13
Ext-Compound-Neutral Ctx	X+ / Y+	XY-	1.35
Ext-Elemental-Inflate Ctx	X+ / Y+	X- / +	1.44
Ext-Elemental-Neutral Ctx	X+ / Y+	X-	1.76
SOC-Compound-Inflate Ctx	Z+ / Y+	XY- / +	1.28
SOC-Compound-Neutral Ctx	Z+ / Y+	XY-	1.74
SOC-Elemental-Inflate Ctx	Z+ / Y+	X- / +	1.71
SOC-Ctrl	Z+ / Y+	Y-	1.22

Note: Z and X were counterbalanced. + indicates reinforcement. - represents nonreinforcement. Slashes denote interspersed trials within sessions. Ext = extinction. SOC = second-order conditioning. Ctrl = control. Ctx = context. Test scores indicate group mean log latencies (s) to drink in the presence of X.

Witnauer and Miller (2011a) group means and best-fitting predictions (Simulation 5).

Table 15

	Ext-Compound-Inflate Ctx	Ext-Compound-Neutral Ctx	Ext-Elemental-Inflate Ctx	Ext-Elemental-NeutralCtx	SOC-Compound-Inflate Ctx	SOC-Compound-Neutral Ctx	SOC-Elemental-Inflate Ctx	SOC-Ctrl
Obs	2.13	1.35	1.44	1.76	1.28	1.74	1.71	1.22
SOCR	2.11	1.42	1.23	1.98	1.13	1.84	1.45	0.70
Backprop	2.12	1.40	2.16	1.42	1.63	1.32	1.68	1.32
Hybrid	1.78	1.56	1.78	1.56	1.48	0.70	1.48	0.70
Pineno	1.83	1.54	1.83	1.54	1.44	0.70	1.44	0.70
VH&W	1.68	1.68	1.68	1.68	0.70	0.70	0.70	0.70
RW	1.68	1.68	1.68	1.68	0.70	0.70	0.70	0.70

Note: SOCR = sometimes competing retrieval. Backprop = backpropagation. Hybrid is a hybrid of the Pineno and Van Hamme and Wasserman models (see text). VH&W = Van Hamme and Wasserman. RW = Rescorla and Wagner. Scores are predicted or observed lick suppression scores in log s.