# Neither Rescorla and Wagner nor SOCR Predict the Double Blocking Effect

A Senior Honors Thesis

Submitted in Partial Fulfillment of the Requirements
for Graduation in the Honors College

By

John Benicasco

Psychology, Philosophy, & Public Health Major

The College at Brockport
May 14, 2019

Thesis Director: Dr. James Witnauer, Associate Professor, Psychology

## Abstract

In a traditional blocking experiment, A+ trials precede AX+ trials, and responding to X at Test is weak relative to controls. Using the allergist task, the present study investigated the effect of an additional blocking cue presented in compound (i.e., ABX+, where B is also a previously trained excitor). The study found evidence of a *double blocking* effect, a proposed associative learning effect in which blocking with two blocking cues results in increased behavioral control by the target cue relative to blocking with a single blocking cue. In Phase 1 participants received A+/B+ training, in Phase 2 participants received ABX+ trials. Relative to controls, responding to X was greater in the double blocking group than in the traditional single blocking group. The Sometimes-Competing Retrieval (SOCR) model ostensibly should be able to explain double blocking through second-order comparator processes; however, a simulation analysis determined that SOCR was unable to satisfactorily explain the effect or fit the present data. Moreover, Rescorla-Wagner did no better than SOCR in fitting double blocking. The double blocking effect thus presents a challenge yet to be surmounted by a model.

Neither Rescorla and Wagner nor SOCR predict the double blocking effect.

## Introduction

When trained in compound, cues have the chance of enhancing the development of behavior control of each other (cue facilitation) or impeding it (cue competition). One of the most straightforward instances of cue competition is blocking. Originally described by Kamin (1968), blocking occurs when a previously trained excitatory cue (A) competes with another less-excitatory target cue (X) during reinforced compound training. After AX+ presentations, behavioral control by X is much weaker than behavioral control by A and relative to control conditions. In associative learning parlance, A is a blocking cue, and the term blocking comes from the idea that blocking cues block or prevent learning about other cues. While blocking was originally discovered in rats, subsequent research from several laboratories has documented blocking in humans (e.g. Shanks, 1985; Eippert, Gamer, & Buchel, 2012; Kruschke & Blair, 2000; Le Pelley, Beesley, & Suret, 2007; Mitchell, Lovibond, Minard, & Lavis, 2006; Prados et al., 2013). In general, associative learning research converges in the finding that basic learning effects are reproducible across species, though some differences do exist (see Rosas, Matías Gámez, León, Tirado, & Nelson, 2017).

One prominent model of associative learning, Rescorla & Wagner's (1972), posits that learning is proportional to the surprisingness of an outcome. Like other total error reduction models, R-W posits that learning happens only in situations where the animal had an incorrect prediction about the world that was subsequently challenged by reality. After this unexpected outcome, the associative strengths of presented cues have the opportunity to change. In a blocking situation, R-W argues that the previous A+ trials have established an

A-Outcome association. Thus, in AX+ trials, the outcome is expected based on A alone, and the animal is not sufficiently surprised to learn about X (Rescorla & Wagner, 1972).

Alternatively, the Sometimes-Competing Retrieval (SOCR) model holds that behavior at test is a function of the relative associative strengths of all cues presented over the entire course of training. Responding at test to any stimulus is, then, determined via a comparison between the directly and indirectly activated representations of the outcome elicited by the present cues. That is, the X-O association's strength is directly proportional to the strength of responding. Conversely, the strength of associations between X and any/all other comparator cues (*j*) as well as the associations between those cues and the outcome (i.e., the X-*j* and *j*-O associations) are inversely proportional to responding. The animal will only respond to X, then, when the X-O association is sufficiently strong enough to overpower the X-*j* and *j*-O associations. To increase responding to X, then, one could either strengthen the X-O association or weaken any of the X-*j* or *j*-O associations. In blocking, instead of a failure to learn about X, SOCR posits that *responding* to X is prevented by a stronger A-O association (which serves as one of the *j*-O associations) (Stout & Miller, 2007).

While both of the above models predict blocking, their predictions diverge in what the effect of an additional blocking cue (B) would be. That is, a situation in which A and B are independently previously trained excitors, and both are presented in compound with the target cue, X, in reinforced blocking trials (ABX+ trials as opposed to AX+ trials). R-W predicts that X would become inhibitory by the same logic that explains overexpectation and superconditioning (see Rescorla, 2003 for explanations of these effects). On an ABX+ trial, two outcomes are expected rather than one (one for each of the cues that signal one outcome), and the resulting lack of extra outcome is surprising. R-W's error-reduction mechanism

results in something akin to blaming X for the discrepancy. That is, X has been associated with the lack or reduction of the outcome (i.e., inhibition).

In contrast, SOCR posits responding to X blocked by two cues would increase relative to a single-blocked control. If that fails to make sense, it is because SOCR as described so far would never allot for cue facilitation. There are two important mechanisms within SOCR, though, that do allow for nontarget cues to facilitate X. One is direct, though a switching operator term that has a proportional effect on competition. The other, however, is more relevant and interesting to the present topic. SOCR proposes that not only do comparator cues modulate responding to X, they also have the ability to affect each other. When A is being compared to X, the A-O association *itself* is also modified by its own interaction with other comparator cues. As such, any training that would decrease responding to A would similarly weaken A's effectiveness as a comparator (Stout & Miller, 2007). In the proposed procedure, the B-O and A-O associations would hypothetically weaken each other as comparator cues, allowing for a stronger response to X.

Due to these differing predictions, an attempt to block with two blocking cues would serve as a useful catalyst to compare R-W and SOCR's relative validities. In a previous study, Witnauer, Urcelay, & Miller (2008) found that A+/B+ trials followed by ABX+ trials resulted in increased behavioral control by X relative to single (traditional) blocking; this *double blocking* effect was predicted by SOCR (Witnauer, Benicasco, & Kopunek, 2017).

The present experiment sought to replicate Witnauer et al.'s (2008) double blocking effect, but utilizing human participants, and to test how well SOCR fit such data. In addition, following the arguments of Ghirlanda & Ibadullayev (2015), we attempted to simulate both

training and test data using SOCR, theoretically increasing the validity of our simulation results. Furthermore, the present study included a more robust single-blocking control group, one that received compound training with three cues to match the three (A, B, and X) used in the double blocking group, as opposed to Witnauer et al.'s (2008) AX+ single blocking control group.

## Methods

### Participants

132 students enrolled in an introductory psychology class at the College at Brockport served as participants in this study. Each participant received class credit for participation in the experiment.

### Materials

The present experiment utilized the allergist task (Wasserman, 1990) in an electronic form in the style of Larkin, Aitken, & Dickinson (1998) programmed and carried out with Psychopy (version 1.85.3). Notably, the allergist task has been shown in previous research to be capable of reproducing the blocking effect (e.g., Beckers, De Houwer, Pineño, & Miller, 2005; Le Pelley et al., 2007; Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003; Mitchell et al., 2006), though there have been failures to do so (e.g., Vandorpe & de Houwer, 2005).

During the allergist task, participants are asked to imagine themselves in the role of an allergist trying to figure out the allergic triggers for a hypothetical patient. Instructions and

information in the present study were all presented on screen (over the course of four screens), requiring the participant to press the spacebar to continue on. Instructions were as follows:

> *This is an investigation of how people learn. It is not a test of your personal abilities or skills. Your name will not be linked with any of the data. In this experiment, you will assess the probability of different foods to cause an allergic reaction. Please press the space bar to continue.*
>
> *Now imagine that you are an allergist who tries to discover the causes of allergic reactions. An allergic reaction can potentially be caused by many different foods that people come into contact with on a daily basis. In order to discover what foods are causing an allergic reaction, you monitor potentially allergenic food. During some meals, you will analyze many potential allergens. During other meals, you will analyze a single potential allergen. Please press the spacebar to continue.*
>
> *Each screen displayed will indicate the potentially allergenic foods consumed in a given meal. After each presentation, please predict the PROBABILITY of an allergic reaction by choosing a point on a scale from -10 to 10. +10 means that a food will always cause an allergic reaction, –10 means that a food will always prevent an allergic reaction that may otherwise be caused by another food. 0 means that a food has no influence on the occurrence or the non-occurrence of an allergic response. The computer will then tell you whether or not an allergic reaction actually occurred. Please press the space bar when ready.*
>
> *Speed of response is not important, so take as much time as you need before responding. At first, you will have to guess, but in time you will begin to learn which foods cause an allergic reaction. Please press the spacebar when ready.*

Participants are shown labeled pictures of food(s) their patient ate, followed by the *noted* presence or absence of an allergic reaction, signaled by “There Was A Severe Allergic Reaction!” framed in a yellow and black zig-zagging line or “There Was No Allergic Reaction.” in a smaller font within a blue-bordered text box.

While the previously cited studies used a variety of foods, the present study chose to use only fruits and vegetables to reduce chances of pre-established biases of allergic potential as well as judgments based on food category as opposed to specific cue. Underneath each set of cue(s), participants indicated how likely they felt the food(s) shown would cause an allergic reaction on a scale of -10 to 10, with -10 indicating protection from an allergic reaction, 0 indicating no correlation, and 10 indicating that the food will always cause a reaction. Intertrial intervals lasted 3 seconds, during which the participant viewed an on-screen message, "Please wait while information about the food is retrieved." Participants were thanked and debriefed with an on-screen message following the test phase.

To control for possible order effects or idiosyncratic acceleration of learning based on cue identity (e.g., bananas always being rated higher), four subgroups were created for each group (Zero Block, One Block, and Two Block) and participants were randomly assigned to one of the twelve subgroups (i.e., one of the twelve psychopy programs, as discussed below). Within each group (each set of four subgroups), all participants received training as per Table 1 (all figures and tables are at the end); however, the identity of cues was counterbalanced both in groups and between groups. Cues A, B, C, and D, the excitors, were banana, apple, cherries, and pear (counterbalanced). The inhibitory cues, E and F, were pineapple and raspberries for all participants, respectively. Finally, all participants were shown kiwi as the target cue, X.

**Procedure**

Participants were run through the experiment in groups of up to five at a time in a research lab with a bank of five computers. Before each run, computers were set up with

different psychopy programs, following a previously generated list of programs to run at each computer. While not entirely random, the list ensured that each program was evenly distributed across the five available computers, and that each program was utilized relatively equally. After obtaining informed consent, participants were collectively asked to sit at any of the prepared computers and self-selected which one they used. This procedure was used to approximate random assignment.

In phase 1, all participants received 24 elemental training trials during which they learned about four food cues. Two of these cues were always paired with the outcome, and the other two were never paired with the outcome. In phase two, half of the trials were reinforced compound trials with three cues, including the novel target cue, X. The remaining three trials were elemental, in which never-paired cues from Phase 1 again failed to cause an allergic reaction.

Manipulation occurred in phase two: the novelty of the two non-target cues differed between groups (see Table 1). For participants in the Zero Block control group, all phase two cues were novel. In the One Block group, one cue out of three was a previously trained excitor (i.e., a blocking cue). In the Two Block group, both of the non-target cues were previously trained excitors. X was then tested alone.

Data from all participants were used to generate csv files, written by the psychopy test programs during data collection. Data were compiled into one Microsoft Excel file (xlsx) utilizing a Matlab script. These procedures reduced chances of human error in copying data.

## Results

**Elimination Criteria**

*N* = 132 students were sorted into the three groups: Two Block (*n* = 44), One Block (*n*=43), and Zero Block (*n*=45). However, data from 33 students were dropped from final analyses as explained below (resulting in Two Block's *n* = 29, One Block's *n* = 34, and Zero Block's *n* = 36). Participants were dropped due to a failure to learn about the cues as defined by incorrect discrimination between cues at the end of phase one. Ratings on the sixth presentation of each of the always-paired cues was compared to ratings on the sixth presentation of the never-paired foods. Participants were assigned one elimination point for each wrong discrimination (defined as rating the never-paired cue as *more likely* to cause a reaction than the always-paired cue), and those with three or four points were dropped from final analyses. Since all cues in phase one were either always paired with the outcome (contingency 1.0) or never paired with the outcome (contingency 0.0), a failure to discriminate between these cues is not likely the result of task difficulty, confusion, or ambiguity in the task, especially since the allergist task has been thoroughly vetted and verified as valid for use in associative learning experiments.

**Control Procedures**

ANOVAs were run to determine the effectiveness of control procedures. Pre-elimination ANOVAs found no main effect on response to X at test of Psychopy program, $F(9, 132) = 1.81$, $p = 0.168$, session, $F(15, 132) = 0.93$, $p = 0.560$, or which computer participants sat at, $F(4, 132) = 2.32$, $p = 0.116$. All subsequent analyses refer to only those data that survived elimination. Control procedures were reassessed to find no main effect of program, $F(9, 99) = 1.81$, $p = 0.168$, session, $F(14, 99) = 0.63$, $p = 0.788$, or computer, $F(4, 99) = 1.50$, $p = 0.290$. Assignment procedures ensured a fair assortment into program; almost every program included 8 or 9

participants, with one exception: one of the Two Block programs only included 5 participants. Participants were also fairly sorted between computers, with 16-23 participants using each computer.

**Acquisition/Training**

Acquisition occurred smoothly, as seen in Figure 1. Moreover, a three way ANOVA between Group (3) x Phase 1 Trial (6) x Cue (4) found main effects of cue, $F(3, 288) = 246.05$, $p < 0.001$ and trial, $F(5, 480) = 7.64$, $p < 0.001$, as well as an interaction between cue and trial, $F(15, 1440) = 68.47$, $p < 0.001$, but no main effect of group $F(2, 96) = 0.41$, $p = 0.668$. Across a range of ANOVA assumptions, no other significant interaction effects were detected, all $F$s $\leq$ 0.89, $p$s $\geq$ 0.415. These results indicate that only cue identity and trial (i.e. how much experience had with each cue) affected phase 1 learning, while group did not (ensuring consistency in phase 1 training across groups).

ANOVAs (3 trial x 3 group) conducted on phase 2 data found that, on ABX+ trials, there was a main effect of trial, $F(2, 192) = 42.00$, $p < 0.001$ and group, $F(2, 96) = 11.42$, $p < 0.001$, as well as an interaction between trial and group, $F(4, 192) = 5.23$, $p = 0.001$. These ANOVAs indicate that phase 1 training did, in fact, affect phase 2 responding. Additionally, participants continued to learn during phase 2, as revealed by the effect of trial.

**Test**

A simple ANOVA discovered a main effect of group on ratings to X at test, $F(2, 96) = 7.43$, $p = 0.001$, denoting an effect of phase 1 training. Planned comparisons were chosen before

data collection, in line with Witnauer et al.'s (2008) analyses. ANOVA significance was achieved, making use of planned comparisons valid. A planned comparison revealed that ratings to X at test were greater in group Zero Block ($M$ = 7.22, $SD$ 0.71) than One Block ($M$ = 3.29, $SD$ = 0.73), $t(96) = 3.85$, $p < 0.001$. The procedure successfully reproduced the basic blocking effect. Finally, a planned comparison between ratings to X at test in groups One Block and Two Block ($M$ = 5.44, $SD$ = 0.79) found, as predicted, an increase in behavioral control by X in the Two Block group, $t(96) = 2.00$, $p = 0.049$ (see Figure 2 for comparison of means). The double blocking effect had been validated by the present study.

## Simulations

So which model, if either, actually predicts the above results? While it is easy to speculate, the nature of science demands a more systematic and unbiased approach, especially since the true implications of a complicated model may be tricky to disentangle without a proper test (Ghirlanda & Ibadullayev, 2015). To do this, we utilized the formal mathematical models of R-W and SOCR and tested them against each other using Matlab simulations. For more in depth coverage of simulation methods used herein (as well as arguments in favor of those methods), see Witnauer, Hutchings, & Miller (2017).

Both R-W and SOCR are defined by sets of equations and assumptions, first explicated by Rescorla & Wagner (1972) and Stout & Miller (2007), respectively. Importantly, R-W functions primarily as a model of *learning*. That is, it predicts how much animals learn about any given cue in a given trial. Responding to that cue (the responses actually measured in any

experiment) is assumed to be directly proportional to what is known about that cue. For example, if an animal knows twice as much about an A-O association as it does a B-O association, R-W predicts that responding to A will be roughly twice that of responding to B. Alternatively, SOCR is primarily a *response rule* and is mostly concerned with how associative strengths and other factors influence the directly measurable responding. The consequent assumption in SOCR is that learning (the strength of mental associations) is roughly contiguity-based. The practical difference in a basic simulation is negligible, though the theoretical difference is important.

**Models**

R-W is formulated as follows:

$$\Delta V_A = \alpha_A * \beta_O * (\lambda_O - \Sigma V_{i-O}) \quad (1)$$

$V_A$ refers to the associativity of A with the outcome (i.e., the strength of the A-O association) and Δ refers to change, such that this equation models learning a series of successive changes in associative strength. α and β represent stimulus properties relevant to speed and strength of learning (e.g., stimulus salience), λ is the maximum possible associative strength of the outcome (i.e., $V_A$ max, as a function of physical properties of the outcome) and the final term refers to the sum of associative strengths of all cues presented on the given trial (e.g., on ABX+ trials, this value would be equal to $V_{A\text{-}O} + V_{B\text{-}O} + V_{X\text{-}O}$).

SOCR, in contrast, is not as easily presented as the model is multilayered and defined by a series of equations. For simplicity, the overall response rule is discussed below, leaving the gory details to Stout & Miller (2007). That response rule is as follows:

$$R_x = V_{X,O} - k_2 * f[\Sigma Op_{X,j,O} * r(V_{X,j}) * r(V_{j,O})] \quad (2)$$

$R_x$ is the predicted level of response to X at test, and $V_{X,O}$, $V_{X,j}$, and $V_{j,O}$ are associative strengths of the connections listed in their subscripts. The *r*s and *f* refer to specific manipulations accomplished by secondary equations using the parenthetical values as inputs; $k_2$ represents a learning-rate parameter (parameter in this sense defined below), and $\sum Op_{X,j,O}$ is the switching operator (the term responsible for basic facilitation and competition), also calculated separately. Important to this project is the effect of the *r*-manipulations, which reflect how SOCR handles meta-interference by comparator cues.

**Simulation Methods**

Matlab (with access to optimization toolbox) accomplished all calculations, modeling, and simulations discussed in this section. The goal of a model simulation is to see how well the model accounts for, or predicts, the *shape* of the data rather than the exact values obtained. For example, the practical value of pinpointing exactly how many fewer leverpresses a rat performs is minimal. Instead, the search for a best-fitting model is a search for the *kinds* of variables that matter, for those factors in reality that are relevant to predicting and understanding behavior and for predictions about relative responding across groups of different conditions. This allows general models to be applied across situations (even across species), while still retaining useful information about behavior. Practically, this is accomplished through use of a scaling factor, a variable multiplied by a model's general predictions to make those predictions directly comparable to real data.

The scaling factor is a free parameter in a model, meaning it is a variable that can take any value necessary to bring the model's predictions closest to reality (again, we are interested in the shape of our model's predictions, not necessarily the exact values obtained). The scaling factor is common to all associative learning models, but each individual model has other free parameters, albeit ones that may be constrained in line with theoretical considerations. For example, in R-W, $\beta_O$ must be less than or equal to one ($\beta_O \leq 1$). These variables are only free to take any value within their specified domain. Importantly, the process of modeling involves programs like Matlab being used to find optimal values for these parameters, such that they fit the given data maximally well as determined by a hill climbing algorithm.

Hill Climbing, in this context, is a process of searching for variables that bring the model's predictions as close to reality as the model will allow. Starting with random values for parameters, a hill climbing algorithm (or, hypothetically, a human with a vast amount of free time), will slightly modify each parameter's value until further modifications continuously lead to a worse fit to data. The best fitting parameters are saved. Hill climbing algorithms repeat this process a specified number of times and will save the best parameters from each repetition. Hill climbing is repeated as certain characteristics of data can lead the hill climbing algorithm to lose its way (though also because Matlab and other programs can perform several thousand calculations in a number of seconds, dramatically reducing the practical cost of double-checking).

The final model's predictions are then statistically compared to data through use of the sum of squared error ($SSE = \Sigma(Predicted - Observed)^2$), and model fit is assessed with the Bayesian Information Criterion (BIC), and the Akaike Information Criterion (AIC) (see Witnauer, et al., 2017). BIC and AIC are measures of how poorly the model fits data, taking into account absolute differences (*SSE*), sample size, and number of free parameters. The model with

the lowest BIC and AIC scores is preferred. Moreover, despite optimization, models may still fail to generate accurate hypotheses.

Models can also be assessed by comparing the optimized free parameters to theoretical considerations. For example, an experimental situation might involve making one cue, A, more salient than another, B (perhaps by magnifying intensity). If hill climbing identifies that the best fitting salience parameter for B is greater than that of A, then the model shows a fundamental disconnect with reality. These considerations are equally, if not more, important than the fit obtained.

**Simulation Results**

The *SSE* calculated for SOCR for our data was 30.72, BIC = 29.88, AIC = 23.62. The SSE for R-W was 47.70, BIC = 34.99, AIC = 29.54. While SOCR's fit was better than R-W's, SOCR's predictions were nonetheless unsatisfactory. SOCR did predict Phase 1 data (see Figure 1), but had difficulty with Phase 2 (see Figure 3). Furthermore, SOCR failed to predict even the rank order of ratings to test, i.e., SOCR predicted that responding in Two Block would be less than responding in One Block, the opposite of what was observed. Since our simulations required that SOCR fit both test and learning data, upon SOCR's failure to do so, we re-ran simulations that did not require SOCR to predict training data. SOCR was able to predict greater responding in Two Block only by giving undue weight to the context and was unable to fit the data without context. The best fitting SOCR simulations of Test are plotted in Figure 4. Moreover, the parameters obtained that allowed SOCR to predict test data generated incredibly

odd predictions when used to simulate training data (e.g., SOCR predicted responding north of 80 on a scale of -10 to 10).

## General Discussion

First and foremost, the double blocking effect was reproduced by the present study. Participants performed greater responding to X at test following blocking with two blocking cues relative to blocking with a single blocking cue. The effect is not new (Witnauer et al., 2008) but replication with the use of an ABX control condition and a different species adds validity to the previous finding. However, the present study does little to explain the blocking effect. To the contrary, simulations provide evidence that SOCR cannot explain or fit double blocking, despite theoretically being the better contender for the job. SOCR and R-W both failed to fit the double blocking effect, and as such, double blocking does not serve to bolster either model's relative merits. The search for a model, paradigm, or correction that can fit double blocking could thus prove useful in improving models of associative learning.

Double blocking, though, has implications for more than just models, it fundamentally undermines a dominant understanding of learning: total error reduction (TER). Learning, as TER theories understand it, is analogous to a primitive version of the scientific method. Animals have hypotheses about their world in the form of predictions or assumptions of what is *about to happen* given what is currently happening. These predictions are routinely tested by reality, and those data slowly update the animal's predictions until they sufficiently approximate reality. That is, TER argues animals learn in response to prediction errors in a way that functions to reduce future prediction errors. R-W is an example of a TER model.

The TER approach to learning is widely accepted within the field, but also has spilled into neuroscience (for review see den Ouden, Kok, & de Lange, 2012; Holland & Schiffino, 2016). Several specific dopaminergic neural circuits have been identified as sources of TER-like learning mechanisms (see Montague, Hyman, & Cohen, 2004; Schultz, 2016). Suggesting an intimate connection between dopamine and the ability to learn, one study found that the dopamine precursor L-DOPA improved reward learning in humans (Chowdhury et al., 2013). Furthermore, L-DOPA proved helpful *only* for those with age-related dopaminergic neuron degeneration and may inhibit learning in younger individuals.

Moreover, dopamine has been implicated in learning about appetitive *and* aversive cues (Talmi, Atkinson, & El-Deredy, 2013). Deficits in dopaminergic signaling in learning pathways have been implicated in symptoms of various psychopathologies, as well (Balsters et al., 2017; Gradin et al., 2011; Griffiths, Langdon, Le Pelley, & Coltheart, 2014). Prediction error is a constant theme in learned motor adaptation research (e.g., Izawa & Shadmehr, 2011; Tseng, Diedrichsen, Krakauer, Shadmehr, & Bastian, 2007), and pupillary correlates of prediction errors in humans have even been identified (Braem, Coenen, Bombeke, van Bochove, & Notebaert, 2015). Suffice it to say, the TER approach is influential and popular.

However useful TER hypotheses may be, though, they fail to explain the present study. In the double blocking procedure, a TER framework hypothesizes that X would become inhibitory or remain neutral. Given that A and B are excitors, the presence of the outcome in ABX+ trials could be interpreted in one of two ways by a TER-based learner. 1) The learner expects two outcomes based on the presence of two excitors. Upon only receiving one, a TER mechanism kicks in which reduces the contingency between at least one of the cues presented and the outcome. In this case X is likely to become inhibitory, though nothing in TER precludes

X from remaining neutral while the associativity of A and/or B is reduced instead. Or, 2) The learner expects an outcome and summarily receives one. While this is in conflict with overexpectation, it is possible that participants assumed a “severe allergic reaction” already represented the maximum possible outcome and thus were not surprised when only one occurred. In this situation, the lack of a prediction error means TER-based learning does not occur. In either case, TER predicts that X has no chance of gaining excitatory properties.

At best for TER, the present study suggests that learning is not solely or principally driven by error reduction mechanisms. At worst, the TER hypothesis is only approximating some other learning mechanism. Either way, further experimentation is needed to sort out the true role of TER-based learning mechanisms.

## References

Balsters, J. H., Apps, M. A. J., Bolis, D., Lehner, R., Gallagher, L., & Wenderoth, N. (2017). Disrupted prediction errors index social deficits in autism spectrum disorder. *Brain*, *140*(1), 235–246. https://doi.org/10.1093/brain/aww287

Beckers, T., De Houwer, J., Pineño, O., & Miller, R. R. (2005). Outcome additivity and outcome maximality influence cue competition in human causal learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(2), 238–249. https://doi.org/10.1037/0278-7393.31.2.238

Braem, S., Coenen, E., Bombeke, K., van Bochove, M. E., & Notebaert, W. (2015). Open your eyes for prediction errors. *Cognitive, Affective, & Behavioral Neuroscience*, *15*(2), 374–380. https://doi.org/10.3758/s13415-014-0333-4

Chowdhury, R., Guitart-Masip, M., Lambert, C., Dayan, P., Huys, Q., Düzel, E., & Dolan, R. J. (2013). Dopamine restores reward prediction errors in old age. *Nature Neuroscience*, *16*(5), 648–653. https://doi.org/10.1038/nn.3364

den Ouden, H. E. M., Kok, P., & de Lange, F. P. (2012). How prediction errors shape perception, attention, and motivation. *Frontiers in Psychology*, *3*. https://doi.org/10.3389/fpsyg.2012.00548

Eippert, F., Gamer, M., & Buchel, C. (2012). Neurobiological mechanisms underlying the blocking effect in aversive learning. *Journal of Neuroscience*, *32*(38), 13164–13176. https://doi.org/10.1523/JNEUROSCI.1210-12.2012

Ghirlanda, S., & Ibadullayev, I. (2015). Solution of the comparator theory of associative learning. *Psychological Review*, *122*(2), 242.

Gradin, V. B., Kumar, P., Waiter, G., Ahearn, T., Stickle, C., Milders, M., … Steele, J. D. (2011). Expected value and prediction error abnormalities in depression and schizophrenia. *Brain*, *134*(6), 1751–1764. https://doi.org/10.1093/brain/awr059

Griffiths, O., Langdon, R., Le Pelley, M. E., & Coltheart, M. (2014). Delusions and prediction error: re-examining the behavioural evidence for disrupted error signalling in delusion formation. *Cognitive Neuropsychiatry*, *19*(5), 439–467. https://doi.org/10.1080/13546805.2014.897601

Holland, P. C., & Schiffino, F. L. (2016). Mini-review: Prediction errors, attention and associative learning. *Neurobiology of Learning and Memory*, *131*, 207–215. https://doi.org/10.1016/j.nlm.2016.02.014

Izawa, J., & Shadmehr, R. (2011). Learning from sensory and reward prediction errors during motor adaptation. *PLoS Computational Biology*, *7*(3), e1002012. https://doi.org/10.1371/journal.pcbi.1002012

Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Maimi Symposium on the Prediction of Behvaior, 1967: Aversive Stimulation* (pp. 9–31). Coral Gables, Florida: University of Miami Press.

Kruschke, J. K., & Blair, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin & Review*, *7*(4), 636–645.

Larkin, M. J., Aitken, M. R., & Dickinson, A. (1998). Retrospective revaluation of causal judgments under positive and negative contingencies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(6), 1331.

Le Pelley, M. E., Beesley, T., & Suret, M. B. (2007). Blocking of human causal learning involves learned changes in stimulus processing. *The Quarterly Journal of Experimental Psychology*, *60*(11), 1468–1476. https://doi.org/10.1080/17470210701515645

Lovibond, P. F., Been, S.-L., Mitchell, C. J., Bouton, M. E., & Frohardt, R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory & Cognition*, *31*(1), 133–142. https://doi.org/10.3758/BF03196088

Mitchell, C. J., Lovibond, P. F., Minard, E., & Lavis, Y. (2006). Forward blocking in human learning sometimes reflects the failure to encode a cue–outcome relationship. *The Quarterly Journal of Experimental Psychology*, *59*(05), 830–844. https://doi.org/10.1080/17470210500242847

Montague, P. R., Hyman, S. E., & Cohen, J. D. (2004). Computational roles for dopamine in behavioural control. *Nature*, *431*(7010), 760.

Prados, J., Alvarez, B., Acebes, F., Loy, I., Sansa, J., & Moreno-Fernández, M. M. (2013). Blocking in rats, humans and snails using a within-subjects design. *Behavioural Processes*, *100*, 23–31. https://doi.org/10.1016/j.beproc.2013.07.014

Rescorla, R. A. (2003). Contemporary study of Pavlovian conditioning. *The Spanish Journal of Psychology*, *6*(2), 185–195.

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

Rosas, J. M., Matías Gámez, A., León, S. P., Tirado, G. G., & Nelson, J. B. (2017). Of rats and people: A select comparative analysis of cue competition, the contents of learning, and

retrieval. *International Journal of Psychology & Psychological Therapy*, *17*(2). Retrieved from https://www.researchgate.net/profile/James_Nelson9/publication/317380625_Of_Rats_and_People_A_Select_Comparative_Analysis_of_Cue_Competition_the_Contents_of_Learning_and_Retrieval/links/5937cc65a6fdcca6588cdd91/Of-Rats-and-People-A-Select-Comparative-Analysis-of-Cue-Competition-the-Contents-of-Learning-and-Retrieval.pdf

Schultz, W. (2016). Dopamine reward prediction error coding. *Dialogues in Clinical Neuroscience*, *18*(1), 23.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *The Quarterly Journal of Experimental Psychology Section B*, *37*(1), 1–21. https://doi.org/10.1080/14640748508402082

Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review*, *114*(3), 759–783. https://doi.org/10.1037/0033-295X.114.3.759

Talmi, D., Atkinson, R., & El-Deredy, W. (2013). The feedback-related negativity signals salience prediction errors, not reward prediction errors. *Journal of Neuroscience*, *33*(19), 8264–8269. https://doi.org/10.1523/JNEUROSCI.5695-12.2013

Tseng, Y., Diedrichsen, J., Krakauer, J. W., Shadmehr, R., & Bastian, A. J. (2007). Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *Journal of Neurophysiology*, *98*(1), 54–62. https://doi.org/10.1152/jn.00266.2007

Vandorpe, S., & de Houwer, J. (2005). A comparison of forward blocking and reduced overshadowing in human causal learning. *Psychonomic Bulletin & Review*, *12*(5), 945–949. https://doi.org/10.3758/BF03196790

Wasserman, E. A. (1990). Attribution of causality to common and distinctive elements of compound stimuli. *Psychological Science, 1*(5), 298–302.

Witnauer, J. E., Benicasco, J. W., & Kopunek, K. (2017, April). *Two blocking cues counteract in human contingency learning*. Poster presented at the Conference on Comparative Cognition, Melbourne, FL.

Witnauer, J. E., Hutchings, R., & Miller, R. R. (2017). Methods of comparing associative models and an application to retrospective revaluation. *Behavioural Processes, 144*, 20–32. https://doi.org/10.1016/j.beproc.2017.08.004

Witnauer, J. E., Urcelay, G. P., & Miller, R. R. (2008). Reduced blocking as a result of increasing the number of blocking cues. *Psychonomic Bulletin & Review, 15*(3), 651–655.

## Tables and Figures

Table 1

*Design of present experiment*

| Group | Phase 1 | Phase 2 | Test |
|---|---|---|---|
| Zero Block (Control) | 6C+ / 6D+ / 6E- / 6F- | | X |
| One Block | 6A+ / 6D+ / 6E- / 6F- | 3ABX+ / 2E- / 1F- | X |
| Two Block | 6A+ / 6B+ / 6E- / 6F- | | X |

Note: Slashes denote interspersed pseudorandom presentations of cues in one phase. Cues A, B, C, and D were apple, banana, cherries, and pear (counterbalanced). A, B, C, and D were always paired with the allergic reaction outcome (signified by +). Cues E and F were pineapple and raspberries, respectively, and were never paired with the outcome (signified by - ). X was kiwi for all participants.

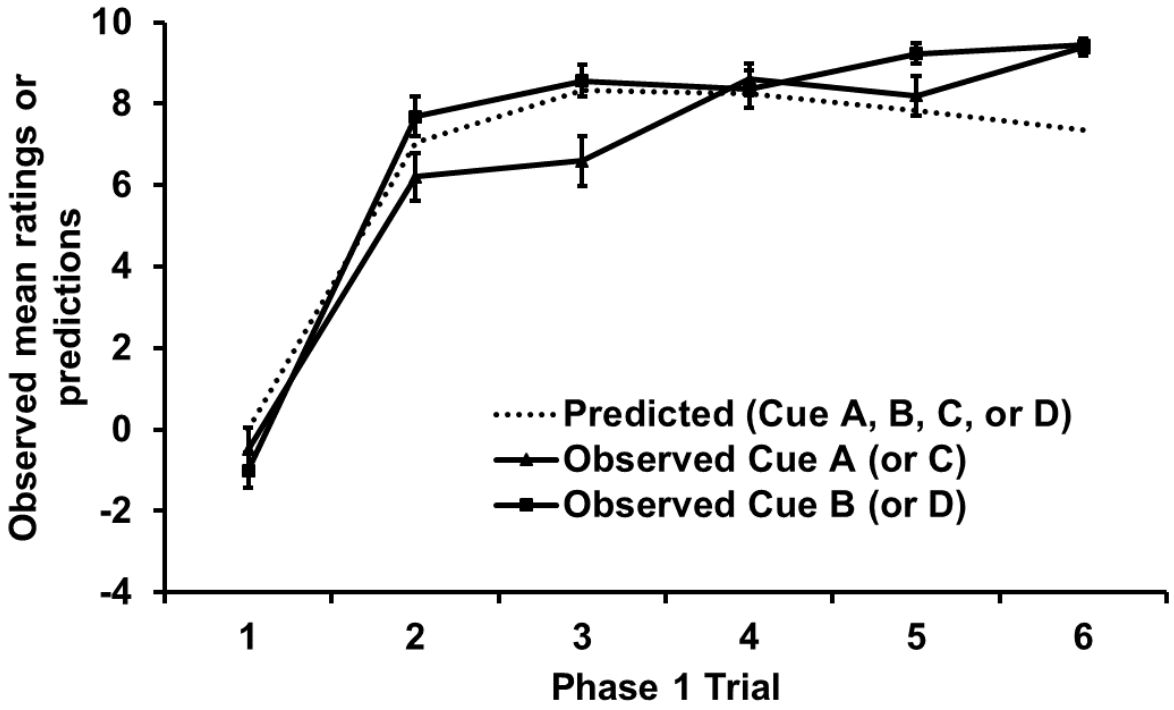


*Figure 1*. Acquisition data (means) from Phase 1 for all groups. Observed A (or C) refers to the first of the two blocking cues, and Observed B (or D) refers to the second blocking cue as ordered in the design table (though not necessarily during the experiment). Trial numbers refer to how many times participant has seen the specified cue (e.g. 1 represents the first trial with that cue and 6 represents the final Phase 1 trial that *contained the specified cue*). SOCR's best fitting predictions for Phase 1 are also plotted and fit data nicely.

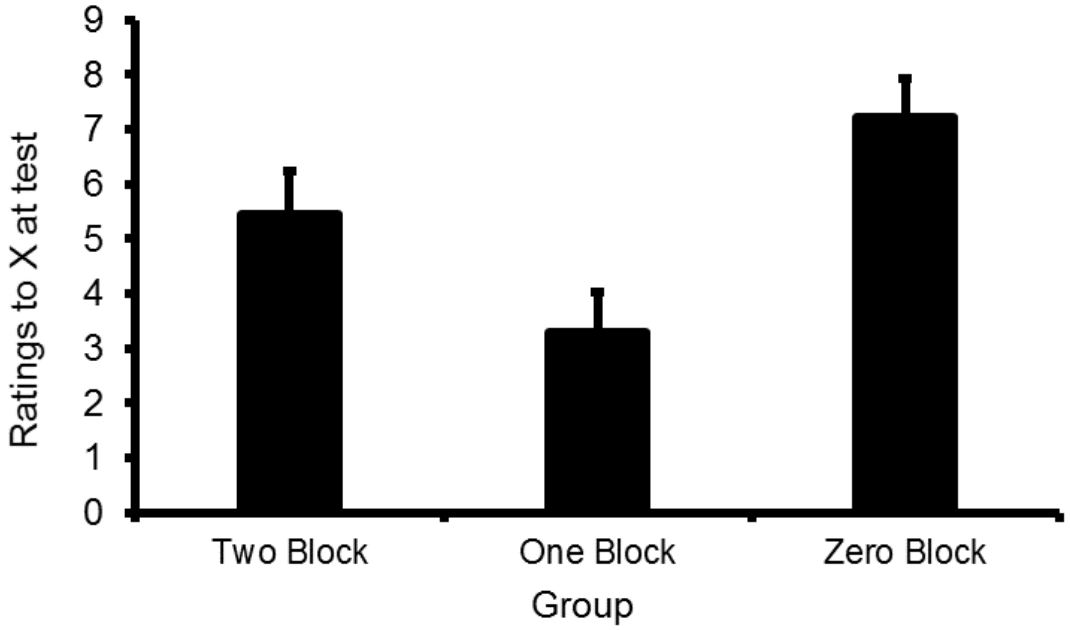


*Figure 2.* Mean responding to X at Test Phase by group. Bars represent standard error.

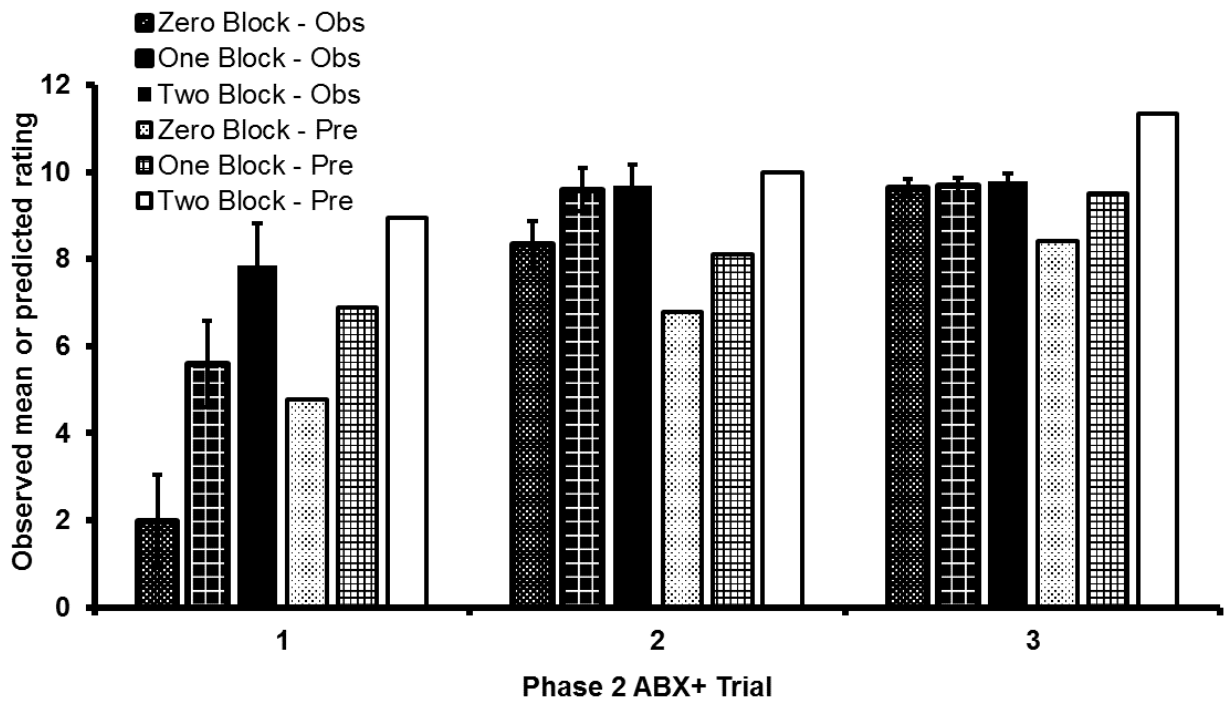


*Figure 3.* Comparison of Phase 2 data to SOCR predictions. "Obs" refers to mean scores from data (i.e., observed scores). "Pre" refers to SOCR predictions. By trial 3, SOCR's predictions no longer matched obtained data.

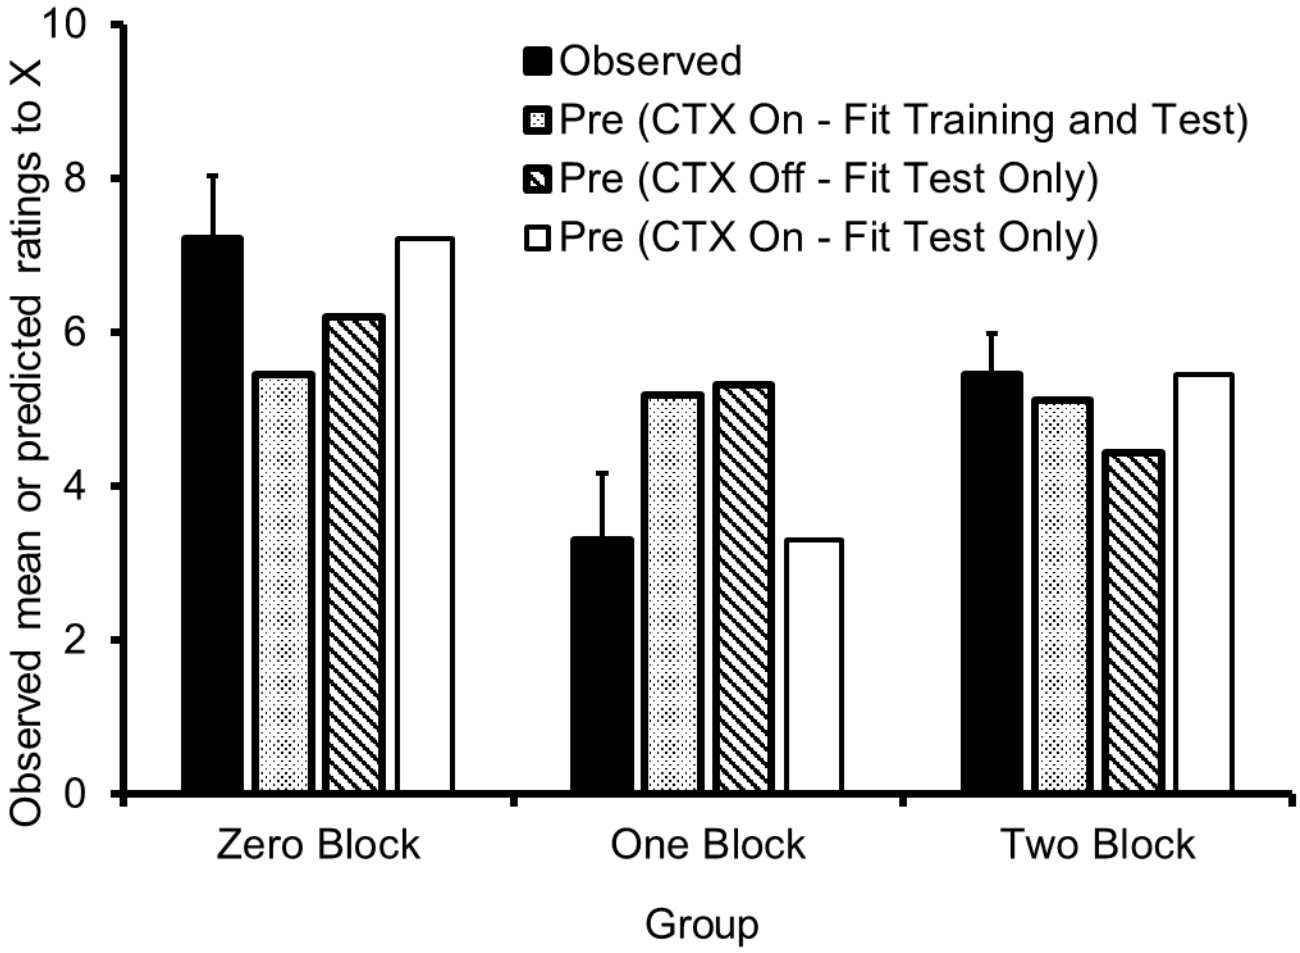


*Figure 4.* Various SOCR predictions compared to mean ratings to X at Test. "Pre" refers to predictions, "CTX" refers to the context, On/Off indicates whether or not context was utilized as a free parameter for each prediction simulation, and "Fit" indicates which experimental phases that simulation was forced to model (either all phases, as in the original simulation, or just the Test Phase). Notice SOCR's poor fit to Test data.