The Importance of Memory in Retrospective Revaluation Learning

Christine M. Chubala

Submitted to the Faculty of Graduate Studies of

The University of Manitoba

in partial fulfillment of the requirements for the degree of

MASTER OF ARTS

Department of Psychology

University of Manitoba

Winnipeg, Manitoba

Abstract

Retrospective revaluation—learning about implied but unpresented cues—poses one of the greatest challenges to classical learning theories. Whereas theorists have revised their models to accommodate revaluation, the empirical reliability of the phenomenon remains contentious. I present two sets of experiments that examine revaluative learning under different but analogous experimental protocols. Results provided mixed empirical evidence that is difficult to interpret in isolation. To address the issue, I apply two computational models to the experiments. An instance-based model of associative learning (Jamieson et al., 2012) predicts retrospective revaluation and anticipates participant behaviour in one set of experiments. An updated classical learning model (Ghirlanda, 2005) fails to predict retrospective revaluation, but anticipates participant behaviour in the other set of experiments. I argue that retrospective revaluation emerges as a corollary of basic memorial processes and discuss the empirical and theoretical implications.

Acknowledgements

I would like to thank my faculty advisor and friend, Dr. Randy Jamieson, for the infinite wisdom, guidance, and encouragement he provided over the course of this project. I would also like to thank my thesis committee members, Dr.'s Debbie Kelly and Jim Hare for their invaluable input. Additional gratitude must be given to my tirelessly supportive family and friends, and the many colleagues and lab mates who have helped along the way. You know who you are. And if you don't, you should probably see somebody about that.

Table of Contents

List of Tables

List of Figures

The Importance of Memory in Retrospective Revaluation Learning

In causal learning studies, participants observe situations in which a given cue or set of cues is followed or not followed by an outcome. Afterward, they are asked to judge the relationship between each cue and the outcome (e.g., how well does the cue predict the outcome?). While observing contingencies between presented cues and outcomes (Dickinson, Shanks, & Evenden, 1984), participants also learn about implied but unpresented cues—a phenomenon known as *retrospective revaluation* (Shanks, 1985).

A simple first-order revaluation protocol includes two training phases followed by a test. In phase one of training a compound cue *AB* is presented followed by an outcome (i.e., *AB+*). In phase two of training a single cue *A* is presented either followed by the outcome (i.e., *A+*) or not (i.e., *A-*). At test, participants rate how well cue *B* predicts the outcome. In general, participants (a) rate *B* as a *better* predictor of the outcome following *AB+, A-* training than following *AB+* training (i.e., upward revaluation of *B*) and (b) rate *B* as a *worse* predictor of the outcome following *AB+, A+* training than following *AB+* training (i.e., downward revaluation). These retrospective revaluation effects can be explained several ways. For example, the result can be explained by rational albeit uncertain inference: "Following *AB+* training it seemed that *A* and *B* were equally good predictors of the outcome. However, learning that *A* alone predicts the outcome suggests that perhaps *B* wasn't a good predictor after all." However, alternative explanations are necessary because retrospective revaluation has been documented in species for which symbolic reasoning is suspect, such as rats (i.e., Baker & Mercier, 1989; Miller & Matute, 1996) and honey bees (Giurfa & Bernard, 2006).

Retrospective revaluation was important historically because it contradicted predictions of the dominant learning theories (e.g., Rescorla & Wagner, 1972; Wagner, 1981) that took as first principle an inability to learn about unpresented cues. Indeed, prominent researchers have hailed retrospective revaluation as "the biggest challenge to traditional theories of associative learning" (Le Pelley, Cutler, & McLaren, 2000) and "the most daunting empirical challenge to the Rescorla and Wagner (1972) model" (Wasserman & Berglan, 1998).

According to the Rescorla-Wagner (RW; 1972) model, the predictive validity of a cue changes only on trials on which it is presented. This assertion is expressed in the model's rule for learning,

$$\Delta V_A = \alpha\beta(\lambda - \Sigma V_{tot}),$$

where ΔV_A is the change in predictive validity for cue *A* between the previous and current trial, α is the salience of the cue, β is the salience of the outcome, λ is the maximum amount of predictive strength supported by the outcome, and ΣV_{tot} is the sum of the associative strengths of all cues presented on the current trial. The challenge to the RW model is simple but absolute: *the model only updates the predictive validity of cues presented on a given trial*. Thus, in the second phase of retrospective revaluation (i.e., training with *A+* after training with *AB+*) new information about cue *A* will prompt an update to the predictive validity of cue *A* but not cue *B*. For a more detailed review of the Rescorla-Wagner model and its shortcomings, see Miller, Barnet, and Grahame (1995).

Van Hamme and Wasserman (1994) modified the RW model to accommodate retrospective revaluation by allowing learning of implied but unpresented stimuli. For example, presenting cues *A* and *B* in compound establishes a within-compound association between the two. Later, when cue *A* is presented alone, the within-compound association causes the model to expect cue *B*. The discrepancy between what the model expects and what is actually presented forces a revaluation of cue *B.*

Dickinson and Burke (1996) applied the same reasoning in a modification of Wagner's (1981) Sometimes Opponent Processes model (MSOP). In MSOP, stimuli (i.e., cues and outcomes) are composed of elements in one of three states: an inactive state (I), a primary active state (A1), or a secondary active state (A2). Positive associations develop between elements in the same active state (both A1, or both A2), and negative associations develop between elements in different states. No association occurs with a stimulus whose elements are inactive. Thus, when cue *A* is presented alone its elements enter state A1 and elements of cue *B* are retrieved into state A2. If the outcome occurs on the trial, it also enters state A1, strengthening a positive association between *A* and the outcome and strengthening a negative association between *B* and the outcome (i.e., prompting a downward revaluation of *B*). Researchers considered the problem solved until a new result was published.

In 2002, De Houwer and Beckers renewed the challenge initially posed by first-order retrospective revaluation when they extended the result to the second-order of association. Their experiment included three learning phases. In phase one, a compound cue *AB* was presented followed by the outcome (i.e., *AB+*). In phase two, a compound cue *BC* was presented followed by the outcome (i.e., *BC+*). In phase three, participants were presented with a single cue *A* either followed by the outcome (i.e., *A+*) or not (i.e., *A-*). Following training, participants rated how well cues *A*, *B*, and *C* predicted the outcome. The experimental results were clear. Participants in the *A+* training group rated cue *A* as a good predictor, cue *B* (the first-order associate to cue *A*) as a poor predictor, and cue *C* (the second-order associate to cue *A*) as a good predictor. In contrast, participants in the *A*-training group rated cue *A* as a poor predictor, cue *B* as a good predictor, and cue *C* as a poor predictor. The *seesaw pattern* of ratings in both conditions—where the first-order associate of the presented cue (i.e., *B* in this protocol) is revalued in the opposite direction of both the presented cue (i.e., *A* in this protocol) and the second-order associate of the presented cue (i.e., *C* in this protocol)—is the hallmark of second-order retrospective revaluation.

Just as Shanks' (1985) demonstration of first-order revaluation frustrated the classical learning theories, De Houwer and Beckers' (2002) demonstration of second-order revaluation frustrated the neoclassical theories that accommodated first-order revaluation but failed to anticipate the seesaw pattern that identifies the second-order effect (Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994). Theorists once again redesigned their models to handle second-order revaluation. According to the post-neoclassical models, learning *AB+* and then *BC+* establishes a first-order within-compound association between cues *A* and *B* and a second-order within-compound association between cues *A* and *C*, mediated by *B*. Thus, presenting cue *A* in the third training phase retrieves both cue *B* and cue *C* and opens both up to learning. To accommodate the seesaw pattern, a *positive learning parameter* is applied to the presented cue (i.e., *A*), a *negative learning parameter* is applied to the first-order associate of the presented cue (i.e., *B*), and a *positive learning parameter* is applied to the second-order associate of the presented cue (i.e., *C*). The elaborated scheme, if built into the Rescorla-Wagner model (Van Hamme & Wasserman, 1994), accommodates first- and second-order retrospective revaluation (see Witnauer &

Miller, 2011). Given the updates, theorists have argued that the original Rescorla and Wagner (1972) model ought remain the dominant and standard account of associative learning (see Miller, Barnet, & Grahame, 1995, for strong disagreement with that position).

In difference to theorists who have worked to salvage the Rescorla-Wagner (1972) model in the face of the challenges posed by retrospective revaluation, a number of other researchers have proposed new and alternative theories of associative learning. De Houwer's (2009) propositional model of associative learning is one example. According to his model, learning involves the formation and evaluation of symbolic propositional structures that describe the relationships between cues or events. Thus, the propositional model explains retrospective revaluation as a process of deductive reasoning, whereby participants who learn *AB+* and then *A-,* for example, infer that cue *B* causes the outcome because cue *A* does not. Miller and colleagues' comparator model (e.g., Miller & Schactman, 1985; Stout & Miller, 2007) provides a second example. It assumes that a cue elicits both direct and indirect retrieval from memory, and that retrospective revaluation follows from a comparison of the directly and indirectly retrieved cues. Melchers, Lachnit, and Shanks (2004) provide a third explanation based on memory rehearsal. In their account, a presented cue retrieves memory of previous learning trials and a process of rehearsal enables learning about the unpresented cues retrieved. Kruschke's (2006; see also Daw & Courville, 2008) Bayesian account provides yet another explanation for revaluation. In that account, learning is represented by a probability distribution of causal knowledge about cues. Retrospective revaluation is explained by a summing-to-one constraint so that learning about *A* requires equal unlearning about *B*. Ghirlanda's (2005; see also Ghirlanda & Enquist, 1999, 2007) neural network model of associative learning provides yet another explanation. It is a post-neoclassical approach but it solves retrospective revaluation by using a vector-based rather than scalar representation scheme for cues and outcomes, instead of the more orthodox approach of introducing additional parameters to the learning rule. Jamieson, Crump, and Hannah (2012; see also Jamieson, Hannah & Crump, 2010) provide a final explanation based in principles of human memory. In their account, events from each learning trial are stored in memory and associative learning follows from the storage and deployment of those instances. Whereas these unorthodox models provide solutions to the problem of retrospective revaluation, a deeper problem underlies their application and relevance.

Despite theorists' enthusiasm to update and revise learning theory to accommodate revaluation, empiricists have warned that those efforts are premature. In short, retrospective revaluation is notoriously difficult to elicit in the lab, even at the first order (e.g., Williams, Sagness, & McPhee, 1994). Indeed, my discussions with researchers suggest that a *file drawer effect*, where researchers have published their successful but not their unsuccessful attempts at demonstrating retrospective revaluation, has left a false impression of a robust phenomenon. The problem is severe. If retrospective revaluation is a false learning phenomenon, we ought to be more reserved in rejecting the classical theory of learning and we ought to be less enthusiastic about developing new models of associative learning that predict the phenomenon.

Part 1 of this thesis provides a clear empirical demonstration of retrospective revaluation at the first-, second-, and third-orders of association. My experiments will be based my experiments on De Houwer and Beckers' (2002) influential procedures. However, I will improve their experiments by extending their protocols and adding necessary control conditions that they did not report. I will also present a second set of experiments based on a different experimental procedure. Part 2 reports computer simulations of my experiments with two models of associative learning: Jamieson et al.'s (2010, 2012) Minerva-AL memory model of associative learning, and Ghirlanda's (2005) post-neoclassical neural network model of associative learning (see also Dawson, 2008; Ghirlanda & Enquist, 1999, 2007). I will use these two models so that I can contrast an unorthodox memory-based account of learning against a more orthodox post-neoclassical approach.

Part 1: Empirical

Since 1994, theorists have been eagerly updating and revising classical learning theory to accommodate retrospective revaluation. However, empiricists have warned that the behavioural data remain suspect.

I address the discrepancy in six experiments. Experiments 1A, 1B, and 1C document first-, second-, and third-order retrospective revaluation using the tank task (De Houwer & Beckers, 2002; Shanks, 1985). Experiments 2A, 2B, and 2C reevaluate those results using the allergist task (Larkin, Aitken, & Dickinson, 1998; Mutter, 2012). To anticipate the results, I report clear evidence for retrospective revaluation with the tank task but no evidence of retrospective revaluation with the allergist task. I will discuss reasons for the difference later.

Experiment 1A: First-Order Tank Task

Theorists have reengineered learning theory to accommodate retrospective revaluation but empiricists have warned that it is an unreliable phenomenon. In light of the uncertainty, Experiment 1A was conducted to evaluate first-order retrospective revaluation in a standard task using our own procedure and instruments.

The experiment included a training phase followed by a test. On each trial in the training phase, the participant watched a cartoon tank drive across a computer screen. Depending on the trial, one or two of five visible weapons fired on the tank. On some trials the tank exploded and on other trials it did not. Following the training phase, participants rated the efficacy of each individual weapon on a scale from 0 to 100. The task follows Shanks (1985) and more closely De Houwer and Beckers (2002a).

Table 1. *Full design of Experiments 1A, 1B, and 1C.*

Conditions	Phase 1	Phase 2	Phase 3
Experiment 1A			
A+ Condition	**AB+** C- D+	**A+** C- D+	----
A- Condition	**AB+** C- D+	**A-** C- D+	----
Control Condition	**AB+** C- D+	----	----
Experiment 1B			
A+ Condition	**AB+** D- E+	**BC+** D- E+	**A+** D- E+
A- Condition	**AB+** D- E+	**BC+** D- E+	**A-** D- E+
Control Condition	**AB+** D- E+	**BC+** D- E+	----
Experiment 1C			
A+ Condition	**AB+** E- F+	**BC+ CD+** E- F+	**A+** E- F+
A- Condition	**AB+** E- F+	**BC+ CD+** E- F+	**A-** E- F+
Control Condition	**AB+** E- F+	**BC+ CD+** E- F+	----

Note. Letters indicate weapons that participants learned about in each phase, in an order randomized per participant. Bolding specifies the critical cues for revaluation. The presence of a tank explosion is indicated by a "+" symbol, whereas the absence of tank explosion is indicated by a "-" symbol. In a subsequent test phase, participants rated the effectiveness of all weapons they had seen in all learning phases.

Participants were assigned to one of three treatments (see the top Panel of Table 1). Participants assigned to a control condition learned *AB+*, *C-*, *D+* before going to test (where letters are weapons and the +/- notation signals explosion/no explosion). Participants assigned to an *A+* condition learned *AB+*, *C-*, *D+* and then *A+*, *C-*, *D+* before going to test. Participants assigned to an *A*- condition learned *AB+*, *C-*, *D+* and then *A-*, *C-*, *D+* before going to test.

Retrospective revaluation will be observed (a) if cue *B* (the first-order associate to cue *A*) is rated to be a worse predictor by participants in the *A+* condition than by participants in the control condition (i.e., downward revaluation of cue *B* given training on

A+ after training on *AB+*) and (b) if cue *B* is rated to be a better predictor by participants in the *A-* condition than by participants in the control condition (i.e., upward revaluation of cue *B* given training on *A-* after training on *AB+*).

Methods

Participants

Eighty-seven undergraduate psychology students from the University of Manitoba participated in this experiment in exchange for course credit. Participants were randomly assigned in equal numbers to the control, the *A+*, and the *A-* treatment conditions.

Apparatus

The experiment was presented on seven different Dell Optiplex PCs, all of which were equipped with identical keyboards, mice, and 22" LCD monitors. The experimental program was written and implemented in Turbo Pascal 7.0 and was obtained with permission from Jan De Houwer. I adapted the software to the relevant experimental protocols. However, I did not alter the timing or graphical presentation of events. All data were collected via keyboard input.

Procedure

The procedure was a close match to the one reported in De Houwer and Beckers (2002a). Participants were informed about the task by typed instructions. The instructions were presented in white font on a black background:

> *In this task, army tanks will ride across the computer screen. A tank can be destroyed by certain weapons. Each weapon is represented as a square on the bottom of the screen. There are five weapons and thus five squares. When a weapon fires, you will see a white light appearing in the square that represents the weapon that fires. On the*

basis of what you observe during these situations, you will afterwards have to judge FOR EACH WEAPON SEPARATELY how likely it is that a tank will be destroyed when that weapon fires. Your task will be complicated by the fact that two weapons can fire simultaneously, that is, at the exact same moment in time. If that happens and the tank is destroyed, you do not know which of the two weapons was responsible for destroying the tank. But, each time that weapons fire, you will receive information about the combined impact that all fired weapons had on the tank. The larger the impact, the larger the chance that the tank will be destroyed. Note that the maximal impact that can be measured corresponds to an impact score of 20. So try to determine FOR EACH WEAPON SEPARATELY how likely it is that that weapon will destroy a tank if it fires on its own.

Press the spacebar when you are ready to begin.

When the participant pressed the space bar, the screen was cleared. 500 ms later a white horizontal line was presented, extending from the left to the right side of the computer monitor, and five 3.0 x 2.5 cm white rectangular borders numbered 1 through 5 from left to right were presented at the bottom of the screen. Panel A in Figure 1 shows a schematic of the display. For each individual participant, weapons (hereafter denoted as weapons *A*, *B*, *C*, *D*, and *E*) were randomly assigned to rectangular borders at the outset of the experiment.

Participants were instructed to press the “B” key to begin. Approximately 3 s later, a white cartoon tank approximately 4 cm wide by 2 cm high moved smoothly from the left to the right side of the computer screen. The bottom of the tank’s wheels touched the horizontal white line.

As a tank moved across the screen, one or two of the five weapons fired, as indicated by the appearance of a 2.7 x 2.3 cm solid white square for 300 ms inside of the relevant rectangular borders arrayed at the bottom of the screen. On "+" trials the tank exploded. This was conveyed by removing the tank from the screen approximately 2 seconds after it first appeared and replacing it with a fan shaped splay of 10 lines that gradually grew in length from 1 cm to 7 cm before contracting back to 0 cm. The message "IMPACT 10/20" was presented in white block letters at the top of the screen to remind learners that only one weapon was responsible for the explosion (see Panel B in Figure 1). On "-" trials, the tank did not explode but instead continued to the right side of the screen. The message "IMPACT 0/20" was presented in white block letters at the top of the screen (see Panel C in Figure 1). It took approximately 6 s for a tank to move all the way from the left to the right side of the screen.

Figure 1. *Schematic drawings of screens seen by participants in the tank task.*

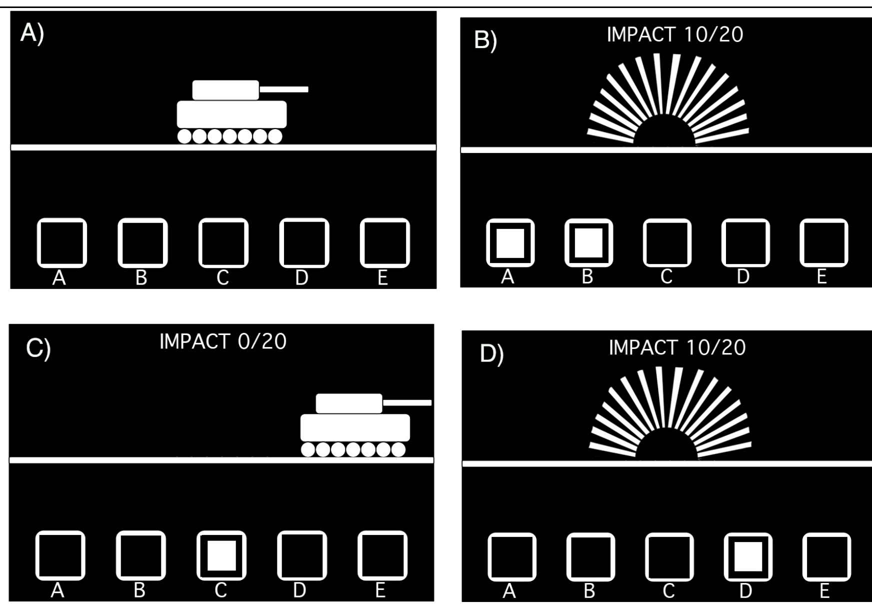

Panel A: A general view of the screen, with no weapons firing. ***Panel B:*** An example of an *AB+* trial, on which weapons *A* and *B* fire and the tank explodes. ***Panel C:*** An example of a *C-* trial, on which weapon *C* fires and the tank does not explode. ***Panel D:*** An example of a *D+* trial, on which weapon *D* fires and the tank explodes.

Approximately 3000 ms after a trial was completed, the next trial began. This cycle continued uniformly until all of the learning trials had been presented.

Following training, participants rated the efficacy of each weapon, moving from left to right on the screen. On each trial, a question was displayed on the screen (e.g., “How effective is weapon number 1?”) and a 10 cm rating scale was displayed beneath it. The rating scale ran from 0 to 100 from left to right. The phrases *very ineffective* and *very effective* were presented at the extreme left and right of the line, respectively. Participants recorded their rating by typing a number between 0 and 100 on the computer keyboard and pressing the enter key.

When the participant pressed the enter key the question on the screen was changed to “How confident are you?” and the labels on the rating scale were changed from *very ineffective* and *very effective* to *very unsure* and *very sure*, respectively. Participants recorded their confidence by typing a number between 0 and 100 on the computer keyboard and pressing the enter key. Confidence ratings were collected by De Houwer and Beckers (2002a) as part of an ongoing project, but they did not report on it. We maintained this part of the experimental procedure to minimize differences between our procedure and that of De Houwer and Beckers (2002a), but we did not analyze confidence data.

To ensure that participants remembered the weapons correctly, the five rectangles that represented the weapons remained on the screen throughout the test phase, just as they had appeared during the training phase. The entire experiment took participants approximately 15 minutes to complete.

Experimental design

Procedures in all three conditions involved weapons firing at a tank. The weapons were labeled *A, B, C, D,* and *E*, and trials are described in those terms. For example, on an *AB+* trial, weapons *A* and *B* fire together and the tank explodes. On a *C-* trial, weapon *C* fires and the tank does not explode. On a *D+* trial, weapon *D* fires and the tank explodes; etcetera (see Panels B through D in Figure 1).

The top Panel in Table 1 shows the three different training regimes that participants were assigned to. As shown, participants assigned to the *Control* condition observed *10 AB+*, *5 C-*, *5 D+* trials before going to test. Participants assigned to the *A+* condition observed *10 AB+*, *5 C-*, *5 D+* trials and then *10 A+, 5 C-, 5 D+* trials before going to test. Participants in the *A-* condition observed *10 AB+*, *5 C-*, *5 D+* trials and then *10 A-, 5 C-, 5 D+* trials before going to test. The order of trial types was randomized within each phase and for each subject.

Results

The top Panel of Figure 2 shows the mean ratings of effectiveness that participants assigned to weapons *A, B, C, D,* and the irrelevant cue *E* that participants did not learn about during training.

As shown, participants revalued weapon *B* (the first-order associate to weapon *A*) in a manner consistent with first-order retrospective revaluation. Firstly, participants in the *A+* condition (M = 21.97, SEM = 5.24) rated weapon *B* to be a <u>less</u> effective weapon compared to participants in the control condition (M = 41.34, SEM = 4.55), $F(1, 56) = 7.78$, $p < .01$. Secondly, participants in the *A-* condition (M = 66.21, SEM = 5.93) rated weapon *B* as a <u>more</u> effective weapon compared to participants in the control condition (M = 41.34, SEM = 4.55), $F(1, 56) = 11.07$, $p < .01$.

Figure 2. *Behavioural and simulated data from Experiment 1A.*

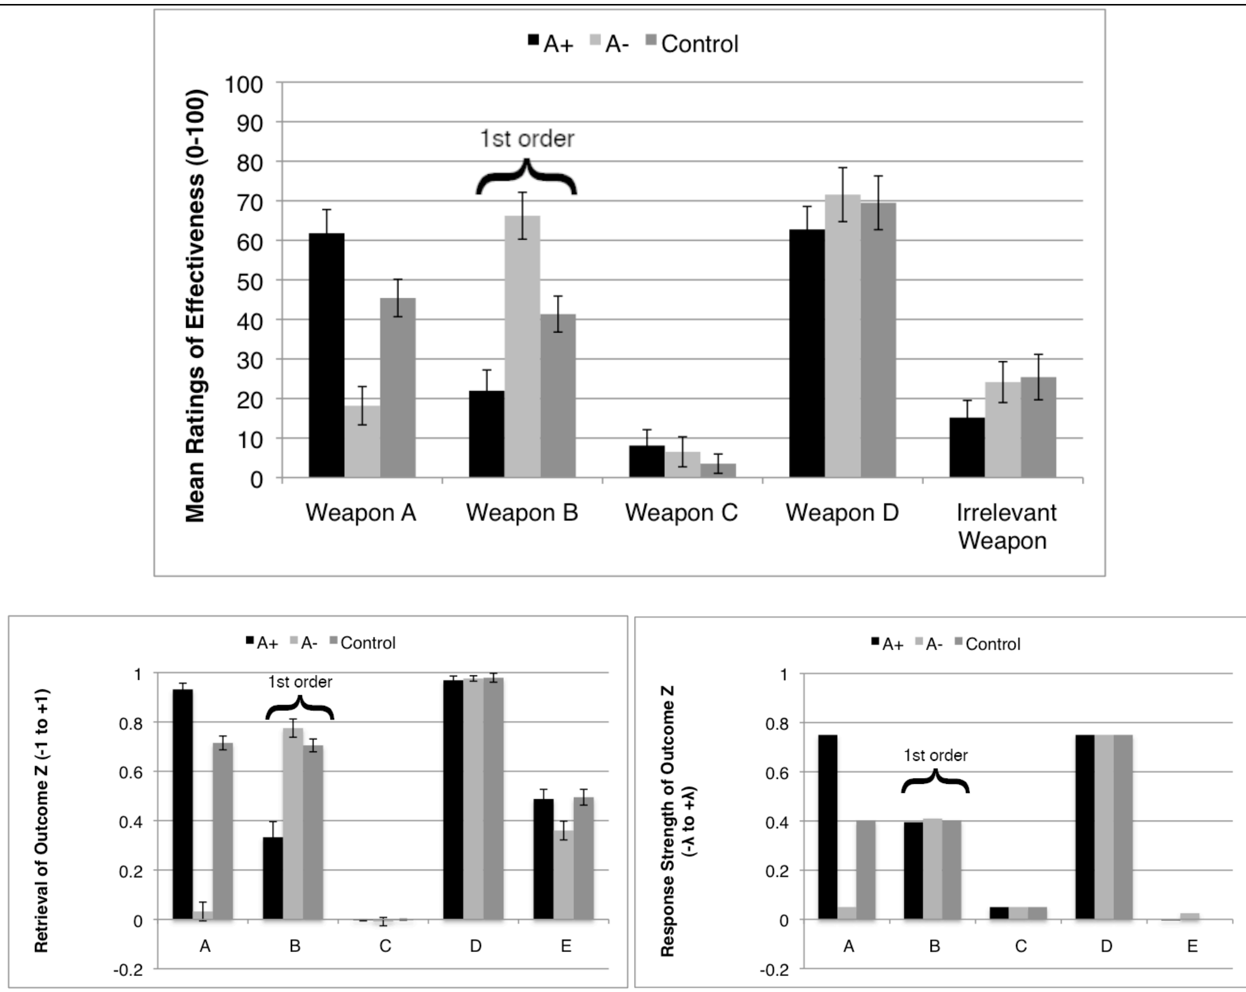

Note. Black bars correspond to the *A+* condition; light grey bars correspond to the *A-* condition; and dark grey bars correspond to the control condition. Error bars indicate 1 standard error. ***Top Panel:*** Mean ratings of effectiveness as a function of experimental condition and weapon. ***Bottom left Panel:*** Mean echo intensity of target outcome, *Z*, as a function of condition and cue used as a probe to memory, using Minerva-AL. ***Bottom right Panel:*** Strength of responding to the outcome, *Z*, as a function of condition and cue in the Ghirlanda model.

The data in Figure 2 also show that the magnitude of the *A*- group's upward revaluation of *B* relative to the control group was equal to the magnitude of the *A+* group's downward revaluation of *B* relative to the control group. The equal size of upward and downward revaluation is unusual. De Houwer and Beckers (2002b), for example, found stronger downward than upward first-order revaluation, albeit nested within a second-order procedure.

In summary, the data provide clear evidence of both upward and downward first-order retrospective revaluation. Therefore the data recommend changes to the classical theories.

In addition to the differences in participants' ratings for weapon *B*, the pattern of ratings over weapons *A*, *C* and *D* were rational. Participants in the *A+* treatment condition (M = 61.79, SEM = 5.99) rated weapon *A* as <u>more</u> effective than those in the control condition (M = 45.41, SEM = 4.73), $F(1, 56) = 4.60$, $p < .05$. Participants in the *A*- treatment condition (M = 18.17, SEM = 4.84) rated weapon *A* as <u>less</u> effective than those in the control condition, $F(1, 56) = 16.17$, $p < .001$. The pattern of ratings for cue *A* is wholly consistent with differences in the frequency structure in the three treatment conditions. Participants' ratings for cues *C* and *D* were also rationally consistent with the presented protocols (see Table 1).

The results reinforce and improve De Houwer and Beckers' (2002a) demonstration of first-order retrospective revaluation on two counts. First, I have evaluated upward and downward retrospective revaluation independent of one another by the addition of a control group. De Houwer and Beckers (2002a), by contrast, evaluated retrospective revaluation far more liberally by comparing ratings for weapon *B* in the *A+* and *A*-

treatment conditions. Whereas De Houwer and Beckers' (200a) analysis is not incorrect, it conflates the measurement of upward and downward revaluation. For example, if ratings for cue *B* in the control condition were equal to ratings for cue *B* in the *A+* condition, I would have detected upward revaluation without downward revaluation. De Houwer and Beckers' (2002a) method cannot make that distinction. Second, De Houwer and Beckers' (2002a) evidence for first-order retrospective revaluation was obtained within a more complicated second-order procedure. This introduces potential confounds. Our procedure simplifies and thereby strengthens their evidence.

Experiment 1A provides a simplest possible demonstration of retrospective revaluation and justifies the need for the neoclassical models. In Experiment 1B, I expand on these results to show retrospective revaluation at the second order of association. Successful demonstration of second-order retrospective revaluation will confirm the need for the complications of the post-neoclassical models.

Experiment 1B: Second-Order Tank Task

Experiment 1B was conducted to establish second-order revaluation, using the same task and basic procedure from Experiment 1A.

Participants were assigned to one of three groups (see the middle Panel of Table 1). Participants assigned to a control condition learned *AB+, D-, E+* and then *BC+, D-, E+* before going to test. Participants assigned to an *A+* treatment condition learned *AB+, D-, E+*, then *BC+, D-, E+,* and then *A+, D-, E+* before going to test. Participants assigned to an *A-* treatment condition learned *AB+, D-, E+*, then *BC+, D-, E+,* and then learned *A-, D-, E+* before going to test.

Second-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate weapon *C* (the second-order associate to weapon *A*) higher than participants in the control condition and (b) if participants in the *A-* treatment condition rate weapon *C* lower than participants in the control condition. First-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate weapon *B* (the first-order associate to weapon *A*) lower than participants in the control condition and (b) if participants in the *A-* treatment condition rate weapon *B* higher than participants in the control condition.

Methods

Participants

One hundred and eleven undergraduate psychology students from the University of Manitoba participated in this experiment in exchange for course credit. Participants were randomly assigned in equal numbers to the control, *A+*, and *A-* treatment conditions.

Apparatus and procedure

The apparatus and procedure were the same as in Experiment 1A.

Experimental design

The experimental design is shown in the middle Panel of Table 1. As shown, participants in the control condition observed *10 AB+*, *5 D-*, *5 E+* trials and then *10 BC+, 5 D-, 5 E+* trials before going to test. Participants assigned to the *A+* treatment condition observed *10 AB+*, *5 D-*, *5 E+* trials, then *10 BC+, 5 D-, 5 E+* trials, and then *10 A+, 5 D-, 5 E+* trials before going to test. Participants in the *A-* treatment condition observed *10 AB+*, *5 D-*, *5 E+* trials, then *10 BC+, 5 D-, 5 E+* trials, and then *10 A-, 5 D-, 5 E+* trials before going to test. As in Experiment 1A, the presentation of trial types was randomized within each phase and for each participant.

Results

The top Panel of Figure 3 shows the mean ratings of effectiveness that participants assigned to weapons *A, B, C, D*, and *E*.

Figure 3. *Behavioural and simulated data from Experiment 1B.*

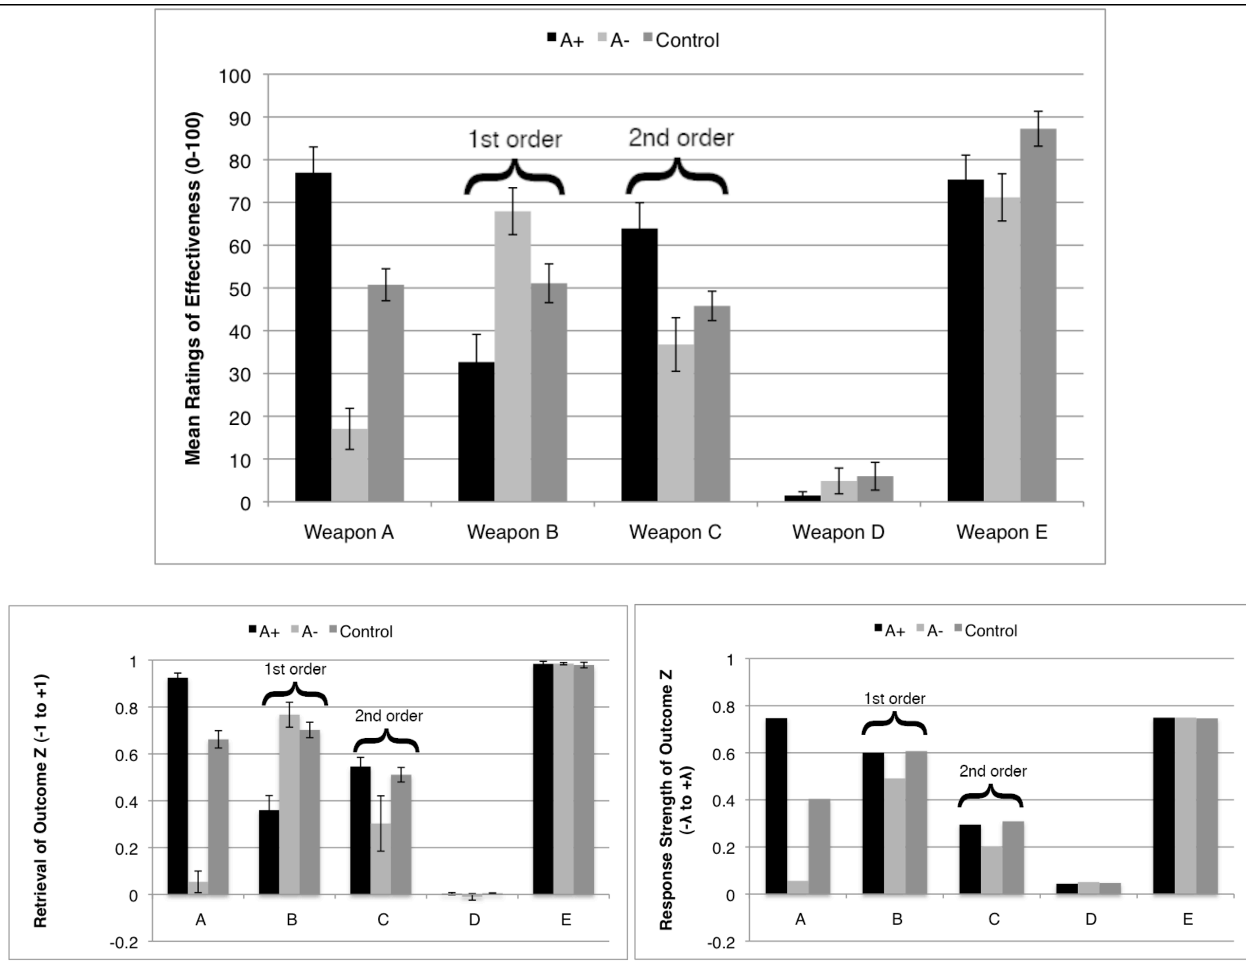

Note. Black bars correspond to the *A+* condition; light grey bars correspond to the *A-* condition; and dark grey bars correspond to the control condition. Error bars indicate 1 standard error. ***Top Panel:*** Mean ratings of effectiveness as a function of experimental condition and weapon. ***Bottom left Panel:*** Mean echo intensity of target outcome, *Z*, as a function of condition and cue used as a probe to memory, using Minerva-AL. ***Bottom right Panel:*** Strength of responding to the outcome, *Z*, as a function of condition and cue in the Ghirlanda model.

As shown, participants revalued weapon *C* (the second-order associate to weapon *A*) in a manner moderately consistent with second-order retrospective revaluation. Participants in the *A+* condition rated weapon *C* as more effective (M = 63.92, SEM = 6.00) compared to participants in the control condition (M = 45.81, SEM = 3.42), $F(1, 72) = 6.86$, $p < .05$. In contrast, whereas participants in the *A-* condition (M = 36.78, SEM = 6.27) rated weapon *C* as less effective than participants in the control condition (M = 45.81, SEM = 3.42), that difference was not statistically reliable, $F(1, 72) = 1.60$, $p = .21$. In summary, I observed both upward and downward second-order retrospective revaluation, however only the upward revaluation was statistically reliable.

Whereas the results of second-order revaluation were somewhat uncertain, the data once again provide evidence of first-order retrospective revaluation. Participants in the *A+* condition (M = 32.67, SEM = 6.47) rated weapon *B* as less effective compared to participants in the control condition (M = 51.11, SEM = 4.54), $F(1, 72) = 5.44$, $p < .05$, and participants in the *A-* condition (M = 67.95, SEM = 5.47) rated weapon *B* as more effective compared to participants in the control condition, $F(1, 72) = 5.61$, $p < .05$.

A comparison of first- and second-order revaluation reveals an asymmetry. In the first-order case, upward and downward revaluation were of equal magnitude. In the second-order case, upward revaluation was stronger than downward revaluation. By comparison, De Houwer and Beckers' (2002b) found stronger downward than upward first-order revaluation, and stronger upward than downward second-order revaluation. Our data suggest a potentially increasing instability of retrospective revaluation as the order of association is increased. De Houwer and Beckers' data show evidence of a vacillation in the magnitude of upward and downward revaluation as the order of association is increased. I will revisit this difference in Experiment 1C.

In addition to the differences in participants' ratings for weapons *C* and *B*, the pattern of ratings over weapons *A*, *D* and *E* was rational. Participants in the *A*+ condition (M = 76.97, SEM = 6.01) rated weapon *A* as <u>more</u> effective than participants in the control condition (M = 50.76, SEM = 3.73), $F(1, 72) = 13.72$, $p < .001$, who in turn rated weapon *A* as <u>more</u> effective than did those in the *A*- condition (M = 17.05, SEM = 4.80), $F(1, 72) = 30.72$, $p < .0001$. The pattern of ratings for weapon *A* is wholly consistent with differences in the frequency structure in the three treatment conditions. Participants' ratings for weapons *D* and *E* were also rationally consistent with the presented protocols (see Table 1).

In summary, the data provide evidence for both first- and second-order retrospective revaluation. However, first-order retrospective revaluation was stronger than second-order revaluation.

As with Experiment 1A, the data reinforce but improve upon De Houwer and Beckers (2002a). Namely, the use of a control condition allowed me to evaluate upward and downward revaluation independent of one another. In contrast, De Houwer and Beckers evaluated retrospective revaluation by comparing ratings for weapon *C* in the *A*+ and *A*- treatment conditions. Notably, if I evaluated the evidence for second-order revaluation as De Houwer and Beckers (2002a) did I would conclude in favor of clear second-order retrospective revaluation, $F(1, 72) = 9.77$, $p < .01$, but fail to appreciate that upward revaluation was stronger than downward revaluation. The discrepancy between our results and those of De Houwer and Beckers (2002a) illustrates the importance of running appropriate controls. More specifically, however, the discrepancy suggests that second-order retrospective revaluation is not as strong an effect as first-order revaluation, or as strong as De Houwer and Beckers (2002a) have claimed it is. In retrospect, the diminished nature of second-order revaluation is unsurprising, as we have known for almost a century that the conditioned response in second-order excitatory conditioning is reliably and measurably weaker than the conditioned response in first-order excitatory conditioning (Pavlov, 1927). Placed in this light, the present results are comforting.

Experiment 1B provides a demonstration of second-order retrospective revaluation and justifies the need for the post-neoclassical models. There are, however, some outstanding questions. Does the effect extend to the third order of association or beyond?

Experiment 1C: Third-Order Tank Task

Experiment 1A established first-order retrospective revaluation and Experiment 1B established second-order retrospective revaluation. The evidence suggests that the magnitude of revaluative learning diminishes as the order of association increases. In Experiment 1C, I evaluate that horizon on revaluation by extending the second-order protocol from Experiment 1B to a third-order protocol. The results of the present experiment ought to clarify whether the magnitude of revaluation does indeed diminish with increasing order of association.

Participants were assigned to one of three treatment conditions (see the bottom Panel of Table 1). Participants assigned to a control condition learned *AB+, E-, F+* and then *BC+, CD+ E-, F+* before going to test. Participants assigned to an *A+* treatment condition learned *AB+, E-, F+*, then *BC+, CD+, E-, F+,* and then *A+, E-, F+* before going to test. Participants assigned to an *A-* treatment condition learned *AB+, E-, F+*, then *BC+, CD+, E-, F+,* and then *A-, E-, F+* before going to test.

Third-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate weapon *D* (the third-order associate to weapon *A*) lower than participants in the control condition and (b) if participants in the *A-* treatment condition rate weapon *D* higher than participants in the control condition. Second-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate weapon *C* (the second-order associate to weapon *A*) higher than participants in the control condition and (b) if participants in the *A-* treatment condition rate weapon *C* lower than participants in the control condition. First-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate weapon *B* (the first-order associate to weapon *A*) lower than participants in the control condition and (b) if participants in the *A-* treatment condition rate weapon *B* higher than participants in the control condition.

Methods

Participants

Seventy-two undergraduate psychology students from the University of Manitoba participated in this experiment in exchange for course credit. Participants were randomly assigned in equal numbers to the control, *A+*, and *A-* treatment conditions.

Apparatus and procedure

The apparatus and procedure were the same as in Experiments 1A and 1B, except that a sixth numbered box was added to the bottom of the screen to accommodate third-order revaluation learning. The screen was slightly rearranged to handle a sixth weapon. Six numbered boxes measuring 3.0 x 2.5 cm were arrayed at the bottom of the screen. Although they were closer together than in Experiments 1A and 1B, they were easily distinguished. Moreover, participants were tested on six weapons, rather than five.

Experimental design

The experimental design is shown in the bottom Panel of Table 1. As shown, participants in the control condition observed *10 AB+*, *5 E-, 5 F+* trials and then *10 BC+, 10 CD+, 5 E-, 5 F+* trials before going to test. Participants assigned to the *A+* treatment condition observed *10 AB+*, *5 E-*, *5 F+* trials, then *10 BC+, 10 CD+, 5 E-, 5 F+* trials, and then *10 A+, 5 E-, 5 F+* trials before going to test. Participants in the *A-* treatment condition observed *10 AB+*, *5 E-*, *5 F+* trials, then *10 BC+, 10 CD+, 5 E-, 5 F+* trials, and then *10 A-, 5 E-, 5 F+* trials before going to test. As in Experiments 1A and 1B, the presentation of trial types was randomized within each phase and for each participant.

Results

The top Panel of Figure 4 shows the mean ratings of effectiveness that participants assigned to weapons *A, B, C, D*, *E,* and *F*.

As shown, participants revalued weapon *D* (the third-order associate to weapon *A*) in a manner only partly consistent with third-order retrospective revaluation. Participants in the *A+* condition (M = 33.38, SEM = 6.23) rated weapon *D* as less effective compared to participants in the control condition (M = 52.71, SEM = 4.88), $F(1, 46) = 5.97$, $p < .05$. However, participants in the *A-* condition (M = 66.25, SEM = 5.48) did not rate weapon *D* as more effective than participants in the control condition (M = 52.71, SEM = 4.88), at least not reliably so, $F(1, 46) = 2.61$, $p = .11$.

Also shown in Figure 4, participants revalued weapon *C* (the second-order associate to weapon *A*) in a manner only partially consistent with second-order retrospective revaluation.

Figure 4. *Behavioural and simulated data from Experiment 1C.*

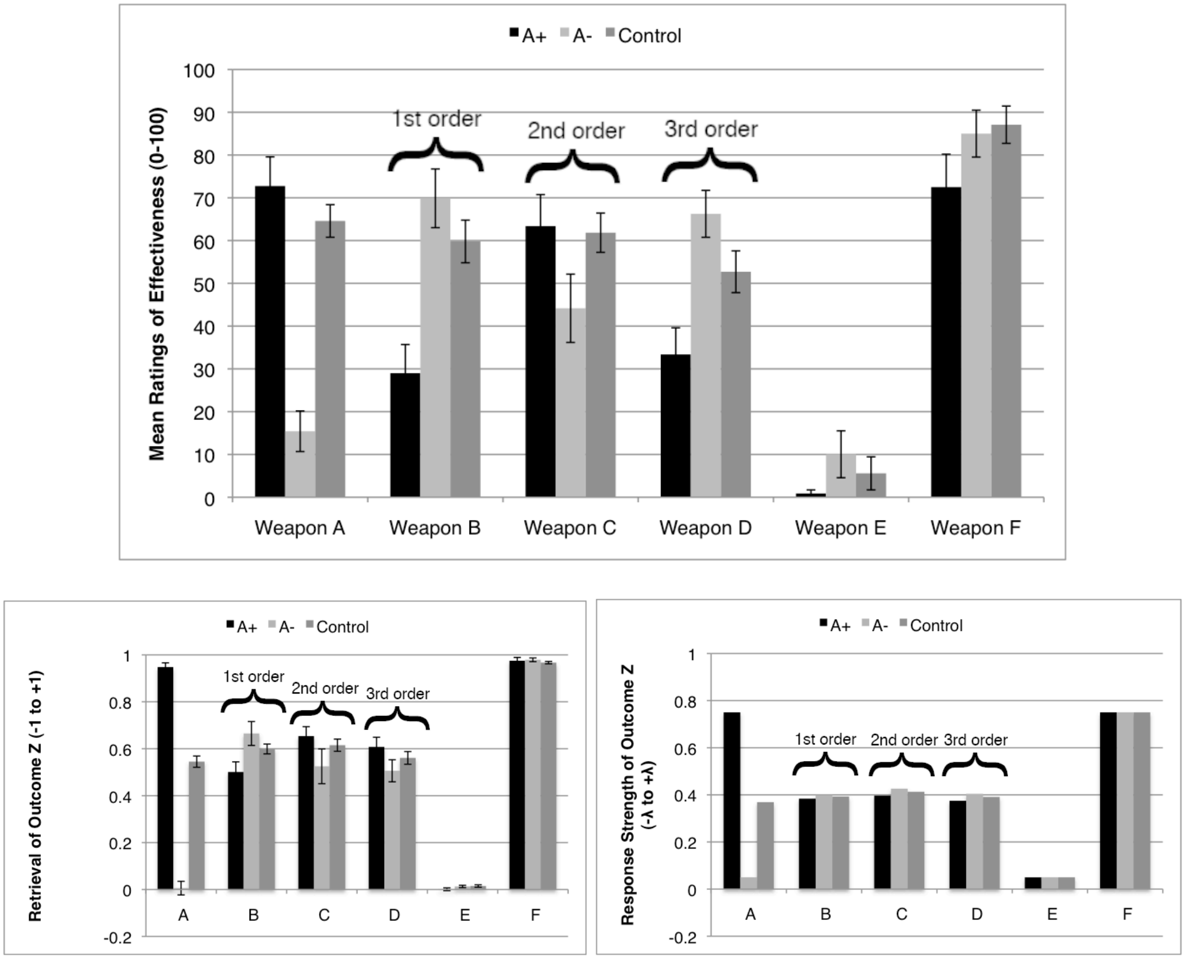

Note. Black bars correspond to the *A+* condition; light grey bars correspond to the *A-* condition; and dark grey bars correspond to the control condition. Error bars indicate 1 standard error. ***Top Panel:*** Mean ratings of effectiveness as a function of experimental condition and weapon. ***Bottom left Panel:*** Mean echo intensity of target outcome, *Z*, as a function of condition and cue used as a probe to memory, using Minerva-AL. ***Bottom right Panel:*** Strength of responding to the outcome, *Z*, as a function of condition and cue in the Ghirlanda model.

Whereas participants in the *A+* condition (M = 63.38, SEM = 7.37) did not rate weapon *C* as <u>more</u> effective than participants in the control condition (M = 61.83, SEM = 4.59), $F(1, 46) = 0.03$, $p = .86$, participants in the *A-* condition rated weapon *C* as <u>less</u> effective (M = 44.17, SEM = 7.98) compared to participants in the control condition (M = 61.83, SEM = 4.59), $F(1, 46) = 3.68$, $p = .06$.

Finally, and also shown in Figure 4, participants revalued weapon *B* (the first-order associate to weapon *A*) in a manner only partially consistent with first-order retrospective revaluation. Whereas participants in the *A+* condition (M = 29.00, SEM = 6.69) rated weapon *B* as <u>less</u> effective compared to participants in the control condition (M = 59.79, SEM = 4.99), $F(1, 46) = 13.61$, $p < .001$, participants in the *A-* condition (M = 69.88, SEM = 6.86) did not rate weapon *B* as <u>more</u> effective compared to participants in the control condition, $F(1,46) = 1.41$, $p = .24$.

Interestingly, the third-order learning protocol seems likewise to have impacted learning about the basic contingency structure of the experiment. That is, although participants in the *A-* condition rated weapon *A* as <u>less</u> effective (M = 15.42, SEM = 4.74) than did those in the control condition (M = 64.58, SEM = 3.81), $F(1, 46) = 65.44$, $p < .0001$, the difference in ratings between the *A+* condition and the control condition did not reach significance, $F(1, 46) = 1.09$, $p = .30$. This difference was nevertheless in the correct direction, with higher ratings in the *A+* condition (M = 72.75, SEM = 6.81) than in the control condition (M = 64.58, SEM = 3.81).

Participants were at least partially sensitive to differences in the frequency structure of the experimental training protocols. Participants in all three conditions rated cue *E* as an ineffective weapon and weapon *F* as an effective one. These ratings map clearly onto the frequency structure for those cues in the experimental protocols.

In summary, the third-order learning protocol caused learning to deteriorate overall. However, participants in the *A+* condition showed evidence of strong first- and third-order retrospective revaluation and participants in the *A-* condition showed evidence of strong second-order and weak first-order retrospective revaluation. We conclude that introducing a more complex learning design frustrates the processes that underlie learning in the retrospective revaluation task.

As I have previously discussed, De Houwer and Beckers (2002a) evaluated retrospective revaluation by comparison of ratings between the *A+* and *A-* treatment conditions rather than against a control condition, as I have done. By their more liberal criterion, I would conclude strong evidence for first- and third-order retrospective revaluation, $F(1, 46) = 18.19$ and 12.70 respectively, $p < .0001$, and less compelling evidence for second-order retrospective revaluation, $F(1, 46) = 3.13$, $p = .08$. By this count, my conclusions are entirely consistent with those of De Houwer and Beckers (2002a). The difference underscores the importance of appropriate controls in the evaluation of retrospective revaluation. Whereas the appropriate analysis can be debated, an evaluation of learning against a control condition cannot.

Based on my analysis, there is a horizon on learning. Experiment 1A showed clear evidence of first-order retrospective revaluation. Experiment 1B also showed evidence of first-order revaluation, but less certain evidence of second-order revaluation. Experiment 1C showed a complicated pattern of ratings with unclear evidence of first-, second-, and third-order revaluation. Although this horizon was clear in my data, it was not in De

Houwer and Beckers' (2002a, b) data. By my estimation, the horizon is a feature of retrospective revaluation that a competent account of learning ought to accommodate.

There are, however, two possible alternative explanations for the deterioration of revaluation learning in Experiment 1C. First, Experiment 1C examines revaluation in a smaller sample of participants than either Experiment 1A or 1B. It could be argued that a larger sample size would have resolved the statistically unreliable findings. I counter this contention with three pieces of evidence. First, the standard errors in the current experiment were consistent with those of the previous experiments, suggesting that the addition of more participants would only have strengthened the observed results. Second, sample size cannot be the critical factor, given that I doubled the sample size of De Houwer and Beckers' (2002a) analogous experiment. If anything, it is more likely that De Houwer and Beckers reported a Type I error. Third, De Houwer and Beckers (2002a) reported evidence of second- and third-order revaluation. In contrast, De Houwer and Beckers (2002b) reported a replication of the second-order result with neither replication nor mention of the third-order result. The omission reinforces the speculation that the evidence for third-order revaluation in the 2002a paper is a Type I error.

A second alternative explanation for the deterioration may be the experimental procedure itself. Namely, the procedure adapted from De Houwer and Beckers' (2002a) Experiment 2 is not a strict example of third-order retrospective revaluation. Because *BC+* and *CD+* trials both occur in the second phase of learning, participants hardly need to recall *BC+* from memory in order to make inferences about *CD+*. In fact, it is equally possible that participants will observe a *CD+* trial before a *BC+* as it is for them to observe the trials in an order conducive to revaluation, in a sense allowing cues *C* and *D* to form unclear and complicated relationships with cue *A*. Because De Houwer and Beckers (2002a) provided the first example of higher-order revaluation in humans, it is as yet unclear exactly how the order of trials would impact the results, but it is not controversial to argue that it would.

Discussion of Experiments 1A, 1B, and 1C

Experiment 1A, 1B and 1C provided evidence for first-, second-, and third-order retrospective revaluation. A comparison of results over those three experiments, however, suggests that retrospective revaluation is weakened as the order of association increases. Thus, whereas the results support and strengthen De Houwer and Beckers' (2002a, b) conclusions, they also suggest a horizon for revaluation learning that is consistent with Pavlov's (1927) finding that second-order excitatory conditioning is weaker than first-order excitatory conditioning. I conclude that retrospective revaluation dilutes with increasing orders of association. This notion casts doubt on De Houwer and Beckers' (2002a, b) conviction that higher-order revaluation is as strong as analogous first-order effects, and I think such reservations makes good sense. If retrospective revaluation did not dilute, then every new experience would force a strong revaluation of every other experience in memory. The idea that revaluation occurs for proximate but not distant cues is a comforting rather than surprising one.

In addition to the empirical importance of such findings, the current experiments provide important constraints for computational models seeking to explain human learning. Any such model must be able to account for higher-order revaluation, as well as the dilution of the effect with increasing order. It is not clear that the neoclassical and post-neoclassical models appreciate the dilution. However, the modern and unorthodox models are well situated.

Indeed, De Houwer (2009) submitted a novel account of human learning to explain revaluation at all orders. His model suggests that learning involves the formation and evaluation of symbolic propositional structures about the relationships between cues or events. Although it can theoretically predict higher-order revaluation, De Houwer's propositional approach suffers three shortcomings. First, the model is conceptual, a sort of thought experiment that De Houwer (2009) conducts to provide a post hoc account of revaluation. Because the model has no systematic formalization, it is limited in the extent to which it can predict new results and more importantly the extent to which it can be falsified. Second, De Houwer commits to no cognitive or neurological mechanism to define propositional units. Without mechanism, there can be no explanation of how similar demonstrations of retrospective revaluation can emerge in human and nonhuman animals. Although the reader can use introspection to uncover his or her own propositional thought processes, the symbolic and propositional abilities of nonhuman animals, particularly as they pertain to syllogistic reasoning, are unlikely to be identical (see, for example, Beckers et al., 2006; Haselgrove, 2010). Finally, there is nothing to suggest that the model would accommodate a horizon on revaluative learning.

A second viable explanation of retrospective revaluation is Miller's comparator hypothesis (Miller & Matzel, 1988) that has recently been formalized as a mathematical model called the Sometimes-Competing Retrieval model (SOCR; Stout & Miller, 2007). De Houwer and Beckers (2002a, b) acknowledge the model's success with their findings and explain how it accomplishes higher-order revaluation.

In the model, judgements are based on a comparison of associative strengths between the outcome and each of the two compounded cues (A and B). Consider second-order revaluation. In the first phase of learning, *AB+* trials establish cue *A* as a potential competitor for cue *B*, so that judgements of cue *B* will depend on the comparison between the outcome given cue *B* and the outcome given cue *A*. In the second phase of learning, *BC+* trials establish cues *B* and *C* as competitors, so that judgements of cue *B* now depend on two comparisons: the outcome given cue *B* as compared to the outcome given cue *A* and the outcome given cue *C*. Judgements of cue *C* will depend on the comparison between the outcome given cue *B* and the outcome given cue *C*. In the third phase of learning, *A+* trials strengthen the association between cue *A* and the outcome, relative to cue *B* and the outcome, resulting in a downward revaluation of cue *B*. When cue *C* is now judged, the outcome given cue *C* will be compared against the weaker association between the outcome and cue *B*, and will be upwardly revalued. In this way, the comparator hypothesis and its mathematical formalization can account for the seesaw pattern observed in higher-order revaluation. The predictions of the comparator hypothesis are not limited to the second order, but it is not clear whether the model can explain the dilution of the effect at higher orders.

Whereas De Houwer and Beckers' (2002a, b) data and my own give the appearance that retrospective revaluation is a solved problem, numerous reports of difficulties in obtaining the results suggest otherwise (e.g., Batson & Batsell, 2000; Mitchell, Killedar & Lovibond, 2005; Vadillo & Matute, 2010; Williams, Sagness, & McPhee, 1994). Any number of factors may contribute to the failures, reversals, and contradictory results: details of the task and cover story (Tangen & Allan, 2004), participants' level of engagement with the task (Le Pelley, Oakeshott, & McLaren, 2005; Wasserman & Berglan, 1998), and—as underscored by the above data—the presence or absence of appropriate control groups.

Rather than attempt to unravel the myriad potential contributions to the mystery of revaluation by means of hypothesis, I turn next to further empirical exploration.

Whereas De Houwer and Beckers' (2002a, b) tank task was informed by Shanks' (1985) early examination of revaluation in humans, another procedure known as the allergist task has since emerged as a dominant one in the field. Introduced by Wasserman (1990), the allergist task invites participants to learn which foods do and do not lead to an allergic reaction in a hypothetical patient.

In Experiments 2A, 2B, and 2C, I examine first-, second-, and third-order revaluation, respectively, using a version of the allergist task popularized by Van Hamme and Wasserman (1994). The purpose of these experiments was to examine the impact of task details on retrospective revaluation. Critically, I modeled the contingencies in these experiments after those used in the previous experiments.

More importantly, the results of the second set of experiments will provide further constraints for theories of learning seeking to accommodate retrospective revaluation.

Experiment 2A: First-Order Allergist Task

In the allergist task (e.g., Van Hamme & Wasserman, 1994), participants make predictions about whether foods will (+) or will not (-) cause an allergic reaction in a fictional patient. To illustrate, they may see a screen that reads, "*Mr. X eats a meal consisting of CARROTS and BANANA. Do you think this will cause an allergic reaction?*" On each trial, participants submit their predictions and are then given feedback (i.e., *"WRONG!!! Eating CARROTS and BANANA DID NOT cause an allergic reaction in Mr. X"*). Participants thereby learn about their patient's set of allergies over multiple trials. Typically, in a revaluation experiment participants first learn about pairs of foods (i.e., *AB+*, *CD+*) and later learn about individual foods (e.g., *A+*, *C-*). In a test phase, participants rate how likely each individual food is to cause an allergic reaction in their patient.

The allergist task has a long and varied history of use in human contingency judgement. In contrast to the tank task, in which participants passively observe contingencies, the allergist task involves active prediction. As a result, whereas ratings for the tank task are unanimously collected on a subsequent test phase, researchers making use of the allergist task have been divided in their methods of collecting ratings: some experiments have made use of end-test ratings, whereas others have examined changes in predictions over trials, and still others have used both. Combined with the inclusion or exclusion of control cues for comparison, the results have been rather mixed.

For example, Dickinson and Burke (1996) designed their experiments to evaluate upward and downward revaluation simultaneously within subjects (i.e., *AB+, CD+* training followed by *A+, C-* training). They assessed revaluation by comparing changes over trials in ratings for the target cues between the *A+* (cue *C*) and the *C-* (cue *D*) contingencies. Like De Houwer and Beckers (2002a), they did not collect any data on control groups or control cues, and therefore they could not independently confirm the presence of revaluation in either contingency.

In contrast, Larkin, Aitkin, and Dickinson (1998) used a between-subjects design that split the procedure over two groups. Although they found evidence of upward revaluation, they failed to find evidence of downward revaluation.

Like Dickinson and Burke (1996), but unlike Larkin et al. (1998), Wasserman and Berglan (1998) assessed upward and downward revaluation simultaneously in a single group of participants, but with the important addition of control cues. They uncovered revaluation in both groups independently, by comparison against control cues.

Of note is that none of these experiments made use of independent control conditions. If comparisons against controls were made at all, they were made against within-subject control cues. That is, participants would learn *AB+* and *CD+* in the first phase of learning and *A+* in the second phase, and would be asked about cues *B* and *D* at test. If ratings for cue *B* were lower than ratings for cue *D*, it was taken as evidence for revaluation. Although this is undoubtedly better than no control cues, there is no evidence that participants treat cues independently such that learning about cue *A* in a second phase would not impact cue *D* in some way. Likewise, there is no evidence that participants learning both contingencies simultaneously (i.e., Dickinson & Burke, 1996; Wasserman & Berglan, 1998) treat the two contingencies independent of one another.

In summary, different researchers have used the allergist task in different ways, and the patterns of results conflict. In the work that follows, I evaluate retrospective revaluation using the allergist task with our own procedure and instruments.

Experiment 2A included two training phases followed by a test. On each trial of the first phase, participants predicted whether a pair of foods would or would not cause an allergic reaction in a fictional patient. After making a prediction, participants were given feedback; on some trials the presented food pair caused an allergic reaction and on other trials it did not. On each trial of the second phase, participants predicted whether a single food would or would not cause an allergic reaction. Participants once again received feedback. Following training, participants rated the likelihood of allergic reaction given the consumption of each individual food using a scale from 0 to 100.

Table 2. *Full design of Experiments 2A, 2B, and 2C*

Conditions	Phase 1	Phase 2	Phase 3	Phase 4
Experiment 2A				
A+ Condition	**AB+** CD+ FG- IJ-	**A+** E+ H- K-	----	----
A- Condition	**AB+** CD+ FG- IJ-	**A-** E- H+ K+	----	----
Control Condition	**AB+** CD+ FG- IJ-	----	----	----
Experiment 2B				
A+ Condition	**AB+** EF+ IJ- KL-	**BC+** FG+ MN- QR-	**A+** D+ H- O-	----
A- Condition	**AB+** EF+ IJ- KL-	**BC+** FG+ MN- QR-	**A-** D- H+ O+	----
Control Condition	**AB+** EF+ IJ- KL-	**BC+** FG+ MN- QR-	----	----
Experiment 2C				
A+ Condition	**AB+** EF+ JK- LM-	**BC+** FG+ JK- NO-	**CD+** GH+ LM- QR-	**A+** I+ P- S-
A- Condition	**AB+** EF+ JK- LM-	**BC+** FG+ JK- NO-	**CD+** GH+ LM- QR-	**A-** I- P+ S+
Control Condition	**AB+** EF+ JK- LM-	**BC+** FG+ JK- NO-	**CD+** GH+ LM- QR-	----

Note. Letters indicate foods that participants learned about in each phase, in an order randomized per block, per participant. Bolding specifies the critical cues for revaluation. The presence of allergic reaction is indicated by a "+" symbol, whereas the absence of an allergic reaction is indicated by a "-" symbol. In a subsequent test phase, participants rated the likelihood of allergic reaction following all foods they had seen in all training phases.

Participants were assigned to one of three treatment conditions. Participants assigned to a control condition learned *AB+, CD+* before going to test (where letters represent individual foods and the *+/-* notation signals allergic reaction/no allergic reaction). Participants assigned to an *A+* condition learned *AB+, CD+* and then *A+, E+* before going to test. Participants assigned to an *A-* condition learned *AB+, CD+* and then learned *A-, E-* before going to test. All training phases included a number of noncritical foods and food pairs to equate the number of foods that did and did not cause an allergic reaction (see top Panel of Table 2).

First-order retrospective revaluation will be observed (a) if food *B* (the first-order associate to food *A*) is rated to be a <u>worse</u> predictor by participants in the *A+* condition than by participants in the control condition (i.e., downward revaluation of *B* given training on *A+*) and (b) if food *B* is rated to be a <u>better</u> predictor by participants in the *A-* condition than by participants in the control condition (i.e., upward revaluation of *B* given training on *A-*).

Methods

Participants

One hundred and seventy-four undergraduate psychology students from the University of Manitoba participated in this experiment in exchange for course credit. Participants were randomly assigned in equal numbers to the three treatment conditions (i.e., control, *A+*, and *A-*).

Apparatus

The experiment was computerized and was presented on seven different Dell Optiplex PCs, all of which were equipped with identical keyboards, mice, and 22" LCD monitors. I wrote the experimental program and implemented it in Unity3D (Unity Technologies, San Francisco, CA). All data were collected via keyboard input.

Stimuli

The stimulus set for Experiments 2A, 2B, and 2C consisted of 33 common fruits and vegetables: apple, asparagus, avocado, banana, beans, beets, berries, broccoli, cabbage, carrot, celery, coconut, corn, cucumber, figs, garlic, grapes, kiwi, lime, mango, melon, mushroom, olives, onion, papaya, peach, pear, pineapple, potato, pumpkin, radish, spinach, and tomato. All stimuli were presented visually in capitalized 14 pt. Arial font. Each participant in the *A+* and *A-* conditions learned about a randomly selected subset of 11 foods. Each participant in the control condition learned about a randomly selected subset of 8 foods.

Procedure

Instructions were presented on the computer screen:

In this experiment, you are asked to imagine that you are an allergist (someone who tries to discover the causes of allergic reactions in people). Please note that this is entirely fictional! The prevalence of real world allergies does not apply and your patients may be allergic to many foods.

You have just been presented with a new patient, 'Mr. X.', who suffers from allergic reactions following some meals, but not others. In an attempt to discover which foods cause Mr. X to have allergic reactions, you arrange for him to eat a number of different meals and observe whether or not he has an allergic reaction.

On the following screens, you will be shown the contents of meals eaten by Mr. X and will be asked to predict whether he will suffer an allergic reaction or not. To make your predictions: Press the YES button if you think Mr. X will suffer an allergic reaction; Press the NO button if you do not think Mr. X will suffer an allergic reaction; Then press the OK button to finalize your prediction. You will then be told by the computer whether an allergic reaction actually occurred.

You will have to guess at first, but with the aid of feedback, your predictions should soon start to become more accurate. Your reactions times are not important in this experiment; you may take as long as you like on each trial. Please do not write anything down, it is all to be done in your head. Press SPACE to begin the experiment.

Instructions were presented one screen at a time, with a new paragraph indicating a new screen. Instruction screens were advanced when the participant pressed the spacebar.

The entire experiment was conducted in black 14 pt. Arial font on a grey background unless otherwise indicated, and everything appearing on the screen was centered horizontally.

After the participant had read the instructions and pressed the spacebar, the first training trial was displayed. Trial instructions were displayed approximately 8 cm from the top of the screen: "*Mr. X eats a meal consisting of FOOD and FOOD. Do you think this will cause an allergic reaction? Select YES or NO, then press "OKAY.*" The *FOOD* labels here apply to two different randomly assigned foods. Approximately 2 cm below these instructions were two 1.5 cm x 3.5 cm dark grey buttons, labeled "*( )NO*" and "*( )YES*" in white 12 pt. Arial font. Participants made predictions by clicking on one of the buttons with the mouse, at which point a letter *X* would appear within the brackets of the selected button. Participants were free to change their selection as many times as they liked before confirming their response by clicking a dark grey "*OK*" button located approximately 2 cm below the Yes/No buttons and measuring approximately 1.5 cm x 5.0 cm.

After clicking "*OK*," the trial screen was immediately replaced by a feedback screen. Approximately 10 cm from the top of the screen, the word "*CORRECT!!!*" or "*WRONG!!!*" appeared, depending on whether participants' response matched the actual outcome defined by the trial contingency. Approximately 3 cm below the feedback, a reiteration of the facts was displayed: "*Eating FOOD and FOOD did cause an ALLERGIC REACTION in Mr. X*" on trials for which the assigned outcome was "+", or "*Eating FOOD and FOOD DID NOT cause an allergic reaction in Mr. X*" on trials for which the assigned outcome was "–". Participants were instructed to press the spacebar to progress to the next trial. Upon pressing the spacebar on the keyboard, the screen was replaced with an asterisk at the very center of an otherwise blank grey screen and one asterisk was added per 0.5 seconds of wait time, until

1.5 seconds had elapsed. At this point, the screen was cleared and the next study trial was displayed.

Training was presented in blocks. In each block, four food-allergy contingencies were presented in a randomized order. Feedback was provided following each prediction. Foods were presented either in pairs or individually, depending on the particular phase of the experiment (see Table 2 for trial contingencies) and training repeated for a minimum of four and a maximum of six blocks. If a participant correctly predicted all four trials on block 4 or block 5 within a learning phase, the next phase of learning began. There was no indication to participants that a phase had ended or begun.[1]

Following training, the computer screen was cleared, and 1.5 sec later test instructions were displayed.

You will now be asked to rate your opinions on the effect that eating various foods has on Mr. X. You will be asked to rate the effects of these foods by moving a sliding bar from 0 to 100.

A rating of 100 means that eating the food is very likely to CAUSE an allergic reaction in Mr. X. A rating of 0 means that eating the food is NOT very likely to cause an allergic reaction in Mr. X.

Press SPACE to begin.

As in the training phase, each paragraph indicates a new screen and participants advanced the instructions by pressing the spacebar.

---

[1] Unfortunately, due to an oversight on the part of the experimenter, participant predictions were not collected and there is therefore no data on whether the number of training blocks (4, 5, or 6) experienced by a participant had any impact on their ratings at test.

After the participant had read the instructions and pressed the spacebar, the first test trial was displayed. Trial instructions were displayed approximately 11 cm from the top of the screen: "*Will eating FOOD cause an allergic reaction in Mr. X? Slide the horizontal bar to indicate the likelihood of an allergic reaction, then press OK. Note: the bar must be moved (even if just a bit) before continuing.*" A 15-cm horizontal line was presented approximately 2 cm below the instructions. The line was labeled 0 at its leftmost point and 100 at its rightmost point. Ticks and accompanying numerical labels occurred at equally spaced ten-unit intervals. A small circular button was located at the midpoint of the scale (labeled 50), which participants could slide to the left or right to record their likelihood ratings. Participants were free to move this slider as many times as they liked before confirming their response by clicking a dark grey "*OK*" button located approximately 2 cm below the line and measuring approximately 1.5 cm x 3 cm. Upon pressing the "*OK*" button with the mouse, the screen was replaced with an asterisk at the very center of an otherwise blank grey screen and one asterisk was added per 0.5 seconds of wait time, until 1.5 seconds had elapsed, at which point the next test trial was displayed. This procedure continued until the participant had rated each of the individual foods seen during training. The order of test foods was randomized for each subject.

Finally, demographic data on participant age, gender, and ESL status were collected. A text box appeared with instructions to participants that they should explain any confusion they had experienced during the experiment, using the computer keyboard. The entire experiment took participants approximately 20 minutes to complete.

Experimental design

The top Panel in Table 2 shows the training regimes for each of the three treatment conditions. As shown, participants assigned to the control condition learned *AB+*, *CD+, FG-, IJ-* before going to test. Participants assigned to the *A+* condition learned *AB+*, *CD+, FG-, IJ-* and then *A+, E+, H-, K-* before going to test. Participants in the *A-* condition learned *AB+*, *CD+, FG-, IJ-* and then *A-, E-, H+, K+* before going to test. All participants observed each contingency a minimum of four times and a maximum of six times, or until they made correct predictions for all four meals in one block. The presentation of contingencies was randomized within each block for each participant.

Results

The top Panel of Figure 5 shows the mean ratings of cue validity assigned to foods *A*, *B*, and control foods. The control foods are separated into *Positive Control* foods and *Negative Control* foods. Because participants in all conditions observed foods *C* and *D* as consistently causing an allergic reaction, ratings for foods *C* and *D* were averaged within each participant to produce an overall *Positive Control* rating. Because participants in all conditions observed foods *G*, *H*, *I*, and *J* as consistently not causing an allergic reaction, ratings for foods *G*, *H*, *I*, and *J* were averaged to produce an overall *Negative Control* rating.

As shown, participants revalued food *B* (the first-order associate to food *A*) in a manner only minimally consistent with first-order retrospective revaluation. Whereas participants in the *A+* condition rated food *B* as marginally less valid (M = 66.12, SEM = 4.82) than did those in the control condition (M = 76.94, SEM = 3.58), $F(1, 114) = 3.25$, $p = .07$, participants in the *A-* condition also rated food *B* as less valid (M = 64.76, SEM = 4.88) than did those in the control condition, $F(1, 114) = 4.04$, $p < .05$.

Figure 5. *Behavioural and simulated data from Experiment 2A.*

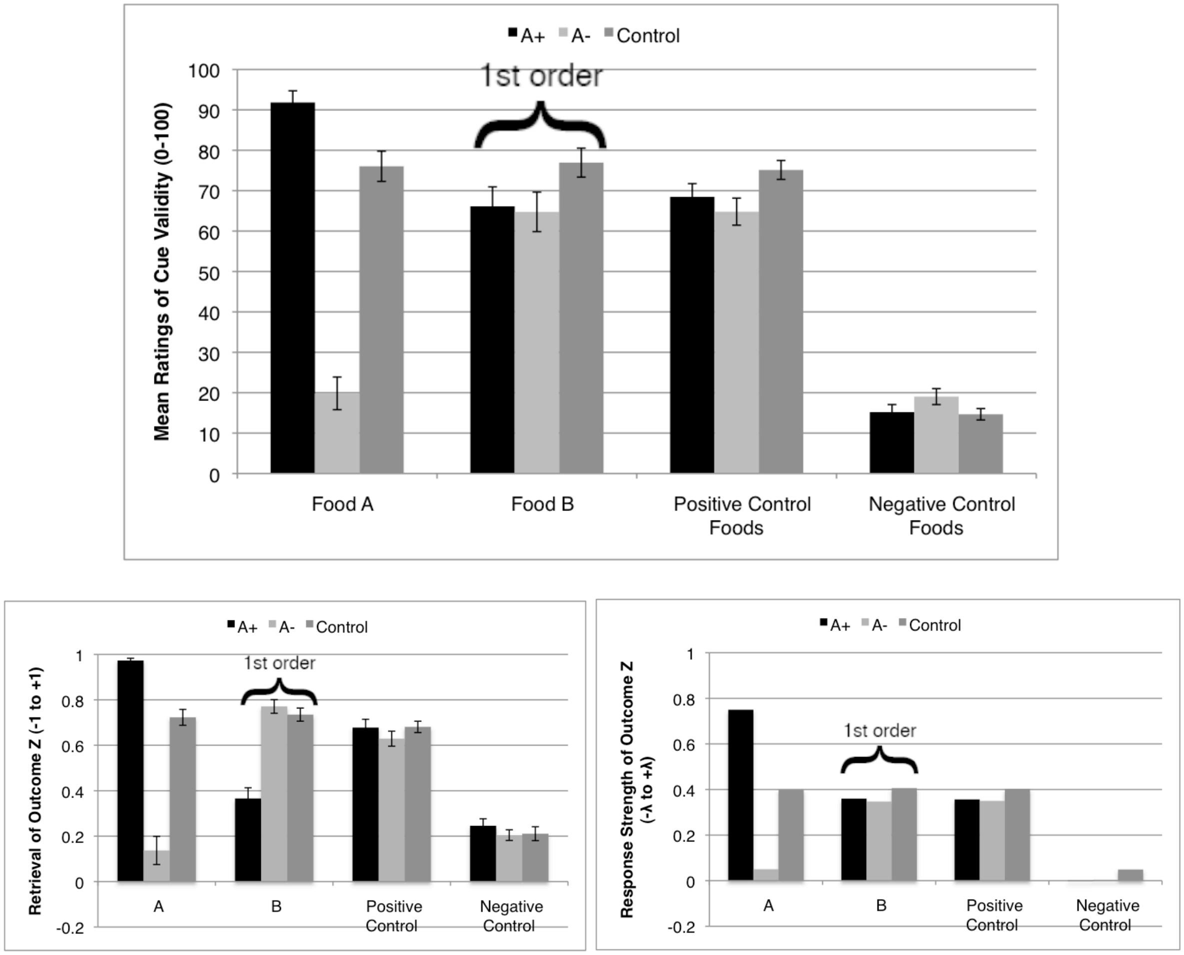

Note. Black bars correspond to the *A+* condition; light grey bars correspond to the *A-* condition; and dark grey bars correspond to the control condition. Error bars indicate 1 standard error. ***Top Panel:*** Mean ratings of cue validity as a function of experimental condition and cue. *Positive Control Foods* represent mean ratings across cues *C* and *D*. *Negative Control Foods* represent mean ratings across cues *F*, *G*, *I*, and *J*. ***Bottom left Panel:*** Mean echo intensity of target outcome, *Z*, as a function of condition and cue used as a probe to memory, using Minerva-AL. ***Bottom right Panel:*** Strength of responding to the outcome, *Z*, as a function of condition and cue in the Ghirlanda model.

According to a liberal interpretation of the results, observed downward revaluation following *A+* training was observed. However, according to a rational interpretation, I observed equal downward revaluation following either *A+* or *A-* training. I conclude a principled failure to observe first-order retrospective revaluation of food *B*.

In addition to the differences in participants' ratings for food *B*, I assessed the pattern of ratings over food *A*, the *Positive Control* foods, and the *Negative Control* foods. In short, participants' ratings of those foods were rational. Participants in the *A+* treatment condition rated food *A* as <u>more</u> valid (M = 91.83, SEM = 2.89) than those in the control condition (M = 76.03, SEM = 3.73), $F(1, 114) = 11.21$, $p < .01$, who in turn rated food *A* as <u>more</u> valid than those in the *A-* condition (M = 19.84, SEM = 4.04), $F(1, 114) = 104.43$, $p < .0001$. Also rational, participants in all three conditions rated *Positive Control* cues as valid and rated *Negative Control* cues as invalid. The failure to obtain clear evidence of first-order revaluation despite a large sample size was therefore not a result of insensitivity to cue contingencies, as participants' ratings for cue *A* and the control cues map clearly onto the frequency structure for those cues in the experimental protocols (see Table 2).

The results do not support the presence of first-order retrospective revaluation in the allergist task, in contrast with previous work that does show evidence of revaluation with the allergist task (e.g., Larkin et al., 1998; Wasserman & Berglan, 1998). Moreover, the results of Experiment 2A contradict the findings of De Houwer and Beckers (2002a) and the results of our own Experiment 1A. That is, despite analogous contingencies, different tasks led to different patterns of results. The inconsistency adds to a growing number of demonstrations of the fragility of retrospective revaluation (e.g., Batson & Batsell, 2000; Mitchell, Killedar & Lovibond, 2005; Vadillo & Matute, 2010; Williams, Sagness, & McPhee).

The failure to produce first-order retrospective revaluation with the allergist task provides little hope for producing second- or third-order revaluation. For the sake of completeness, however, Experiment 2B examines second-order revaluation in the allergist task and Experiment 2C examines third-order revaluation in the allergist task. In short, Experiments 2B and 2C provide no more evidence of retrospective revaluation using the allergist task than Experiment 2A does.

Experiment 2B: Second-Order Allergist Task

Participants were assigned to one of three groups. Participants assigned to a control condition learned *AB+, EF+* then *BC+ FG+* before going to test. Participants assigned to an *A+* condition learned *AB+, EF+,* then *BC+, FG+*, and then *A+, D+* before going to test. Participants assigned to an *A-* condition learned *AB+, EF+,* then *BC+, FG+*, and then *A-, D-* before going to test. All training phases included a number of noncritical foods and food pairs to equate the number of foods that did and did not cause an allergic reaction (see middle Panel of Table 2).

Second-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate food *C* (the second-order associate to food *A*) higher than participants in the control condition and (b) if participants in the *A-* treatment condition rate food *C* lower than participants in the control condition. First-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate food *B* (the first-order associate to food *A*) lower than participants in the control condition and (b) if participants in the *A-* treatment condition rate food *B* higher than participants in the control condition.

Methods

Participants

Fifty-four undergraduate psychology students from the University of Manitoba participated in this experiment in exchange for course credit. Participants were randomly assigned in equal numbers to the control, *A+*, and *A-* treatment conditions.

Apparatus, stimuli, and procedure

The apparatus, stimuli, and procedure were the same as in Experiment 2A.

Experimental design

The middle Panel of Table 2 shows the training regimes for each of the three treatment conditions. As shown, participants assigned to the control condition learned *AB+*, *EF+, IJ-, KL-*, then *BC+, FG+, MN-, QR-* before going to test. Participants assigned to the *A+* condition learned *AB+*, *EF+, IJ-, KL-*, then *BC+, FG+, MN-, QR-*, and then *A+, D+, H-, O-* before going to test. Participants in the *A-* condition learned *AB+*, *EF+, IJ-, KL-*, then *BC+, FG+, MN-, QR-*, and then *A-, D-, H+, O+* before going to test. All participants observed each trial type a minimum of four times and a maximum of six times, or until they made correct predictions for all four meals in one block. The presentation of contingencies was randomized within each block for each participant.

Results

The top Panel of Figure 6 shows the mean ratings of cue validity assigned to foods *A*, *B*, *C,* and control foods, both positive and negative. Because participants in all conditions observed foods *F* and *G* as consistently causing an allergic reaction, ratings for foods *F* and *G* were averaged to produce an overall *Positive Control* rating. Because participants in all conditions observed foods *M*, *N*, *Q*, and *R* as consistently not causing an allergic reaction, ratings for foods *M, N, Q,* and *R* were averaged to produce an overall *Negative Control* rating.

Figure 6. *Behavioural and simulated data from Experiment 2B.*

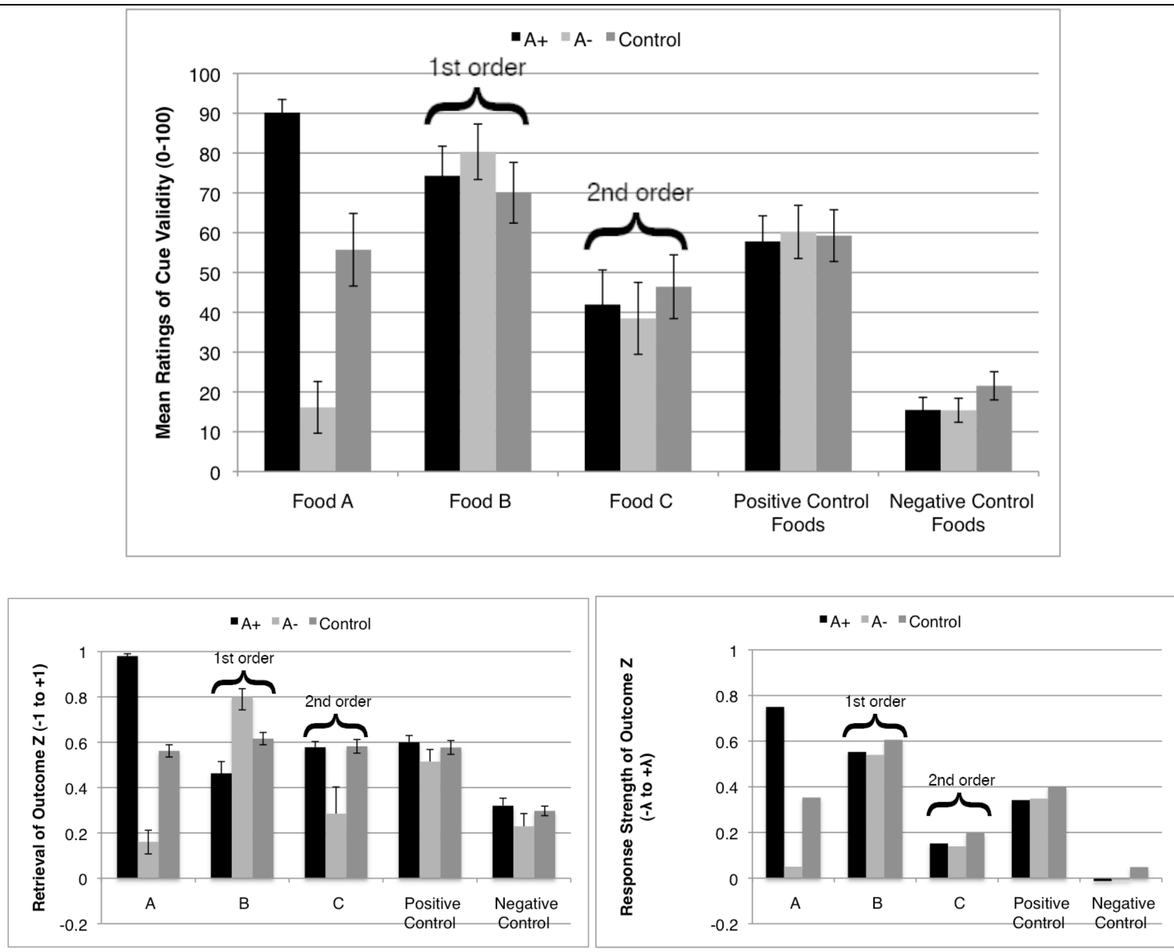

Note. Black bars correspond to the *A+* condition; light grey bars correspond to the *A-* condition; and dark grey bars correspond to the control condition. Error bars indicate 1 standard error. ***Top Panel:*** Mean ratings of cue validity as a function of experimental condition and cue. *Positive Control Foods* represent mean ratings across cues *F* and *G*. *Negative Control Foods* represent mean ratings across cues *M*, *N*, *Q*, and *R*. ***Bottom left Panel:*** Mean echo intensity of target outcome, *Z*, as a function of condition and cue used as a probe to memory, using Minerva-AL. ***Bottom right Panel:*** Strength of responding to the outcome, *Z*, as a function of condition and cue in the Ghirlanda model.

As shown in the top Panel of Figure 6, participants' ratings for food *A*, *Positive Control* foods, and *Negative Control* foods are consistent with the contingency structure of the experimental protocol. In particular, participants in the *A+* treatment condition rated food *A* as more valid (M = 90.15, SEM = 3.31) than those in the control condition (M = 55.71, SEM = 9.11), $F(1, 34) = 12.60$, $p < .01$, who in turn rated food *A* as more valid than did those in the *A-* condition (M = 16.12, SEM = 6.49), $F(1, 34) = 12.51$, $p < .01$. Equally clear, however, is a complete failure of both first- and second-order revaluation.

Participants in both the *A+* (M = 41.95, SEM = 8.67) and *A-* (M = 38.47, SEM = 9.03) conditions rated food *C* (the second-order associate to food *A*) as slightly less valid than those in the control condition (M = 46.42, SEM = 8.01), and not statistically so, both *F*'s < 1.

The same was true of first-order revaluation, except that participants in both the *A+* (M = 74.30, SEM = 7.42) and *A-* (M = 80.33, SEM = 6.97) conditions rated food *B* (the first-order associate of food *A*) as slightly more valid than those in the control condition (M = 70.03, SEM = 7.63), and not statistically so, both *F*'s < 1.

In summary, although participants learned about the presented contingencies, there was no compelling evidence of either first- or second-order revaluation. In Experiment 2C, I extend the second-order analysis of retrospective revaluation to the third order. As in Experiment 2A, which failed to show evidence of first-order revaluation, and Experiment 2B, which failed to show evidence of second-order revaluation, Experiment 2C will fail to show evidence of third-order revaluation.

Experiment 2C: Third-Order Allergist Task

Participants were assigned to one of three groups. Participants assigned to a control condition learned *AB+, EF+*, then *BC+ FG+*, and then *CD+, GH+* before going to test.

Participants assigned to an *A+* condition learned *AB+, EF+,* then *BC+, FG+*, then *CD+, GH+,* and then *A+, I+* before going to test. Participants assigned to an *A-* condition learned *AB+, EF+,* then *BC+, FG+*, then *CD+, GH+,* and then *A-, I-* before going to test. All training phases included a number of noncritical foods and food pairs to equate the number of foods that did and did not cause an allergic reaction (see bottom Panel of Table 2). The reader will note that the design of Experiment 2C corrects for the potential confound in the design of Experiment 1C (i.e., participants were trained on *BC+* and *CD+* in the same phase of learning); third-order revaluation is now separated into four unique phases of trials to facilitate the possibility of observing third-order retrospective revaluation.

Third-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate food *D* (the third-order associate to food *D*) lower than participants in the control condition and (b) if participants in the *A-* condition rate food *D* higher than participants in the control condition. Second-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate food *C* (the second-order associate to food *A*) higher than participants in the control condition and (b) if participants in the *A-* treatment condition rate food *C* lower than participants in the control condition. First-order retrospective revaluation will be observed (a) if participants in the *A+* condition rate food *B* (the first-order associate to food *A*) lower than participants in the control condition and (b) if participants in the *A-* treatment condition rate food *B* higher than participants in the control condition.

Methods

Participants

Fifty-four undergraduate psychology students from the University of Manitoba participated in this experiment in exchange for course credit. Participants were randomly assigned in equal numbers to the control, *A+*, and *A-* treatment conditions.

Apparatus, stimuli, and procedure

The apparatus, stimuli, and procedure were the same as in Experiments 2A and 2B.

Experimental design

The bottom Panel in Table 2 shows the training regimes for each of the three treatment conditions. As shown, participants assigned to the control condition learned *AB+*, *EF+, JK-, LM-*, then *BC+, FG+, JK-, NO-*, then *CD+, GH+, LM-, QR-* before going to test. Participants assigned to the *A+* condition learned *AB+*, *EF+, JK-, LM-*, then *BC+, FG+, JK-, NO-*, then *CD+, GH+, LM-, QR-*, and then *A+, I+, P-, S-* before going to test. Participants in the *A-* condition learned *AB+*, *EF+, JK-, LM-*, then *BC+, FG+, JK-, NO-*, then *CD+, GH+, LM-, QR-*, and then *A-, I-, P+, S+* before going to test. All participants observed each trial type a minimum of four times and a maximum of six times, or until they made correct predictions for all four meals in one block. The presentation of contingencies was randomized within each block for each participant.

Results

The top Panel of Figure 7 shows the mean ratings of cue validity assigned to foods *A*, *B*, *C, D*, and control foods. Because participants in all conditions observed foods *G* and *H* as consistently causing an allergic reaction, ratings for foods *G* and *H* were averaged to produce an overall *Positive Control* rating. Because participants in all conditions observed foods *L*, *M*, *Q*, and *R* as consistently not causing an allergic reaction, ratings for foods *L, M, Q,* and *R* were averaged to produce an overall *Negative Control* rating.

Figure 7. *Behavioural and simulated data from Experiment 2C.*

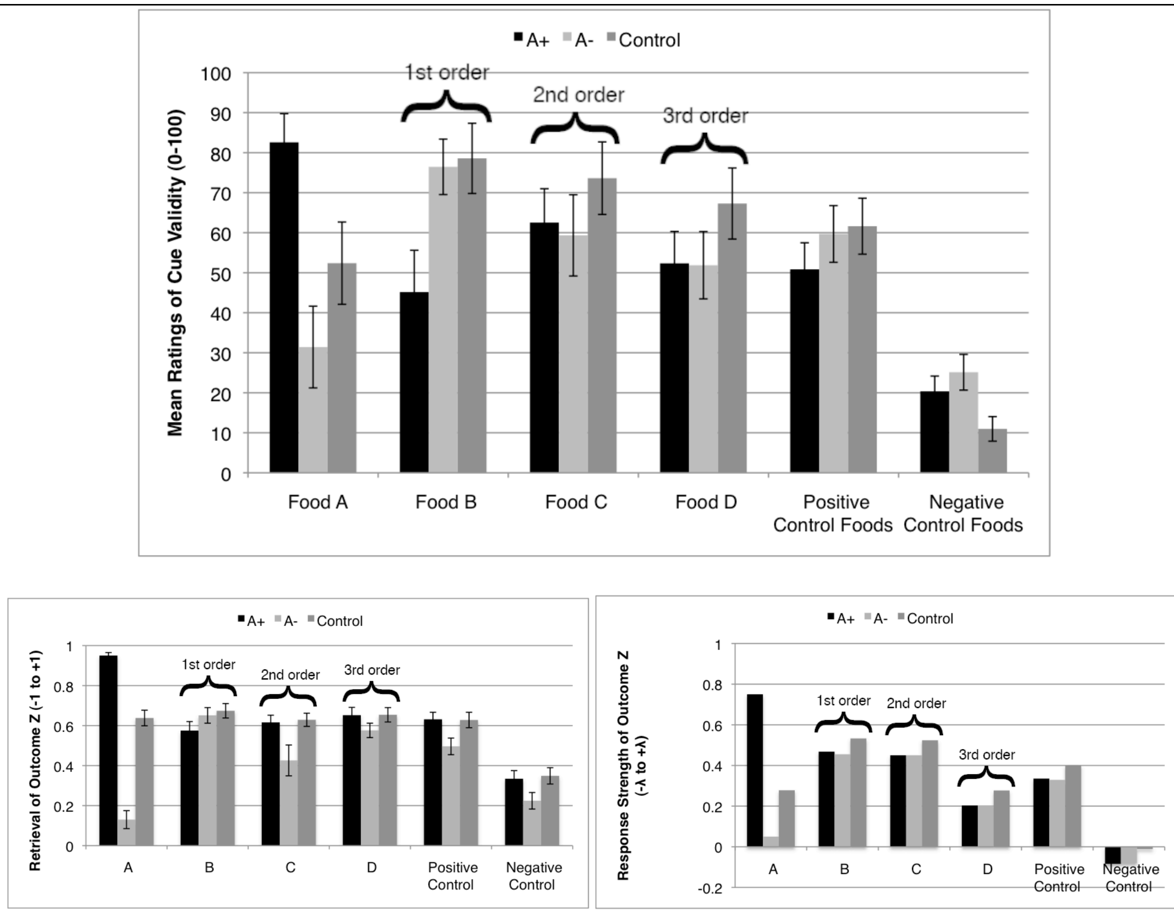

Note. Black bars correspond to the *A+* condition; light grey bars correspond to the *A-* condition; and dark grey bars correspond to the control condition. Error bars indicate 1 standard error. ***Top Panel:*** Mean ratings of cue validity as a function of experimental condition and cue. *Positive Control Foods* represent mean ratings across cues *G* and *H*. *Negative Control Foods* represent mean ratings across cues *L*, *M*, *Q*, and *R*. ***Bottom left Panel:*** Mean echo intensity of target outcome, *Z*, as a function of condition and cue used as a probe to memory, using Minerva-AL. ***Bottom right Panel:*** Strength of responding to the outcome, *Z*, as a function of condition and cue in the Ghirlanda model.

As shown in the top Panel of Figure 7, participants' ratings for food *A*, *Positive Control* foods, and *Negative Control* foods were consistent with the contingency structure of the experimental protocol. In particular, participants in the *A+* treatment condition rated food *A* as <u>more</u> valid (M = 82.59, SEM = 7.16) than those in the control condition (M = 52.39, SEM = 10.28), $F(1, 34) = 16.82$, $p < .001$, who in turn rated food *A* as <u>more</u> valid than did those in the *A-* condition (M = 31.43, SEM = 10.21), $F(1, 34) = 2.09$, $p = .16$. Equally clear, however, is a complete failure of first- and second-, and third-order revaluation.

As shown, participants in the control condition rated foods *D* (the third-order associate to food *A*), *C* (the second-order associate to food *A*), and *B* (the first-order associate to food *A*) as <u>more</u> valid than participants in either the *A+* or *A-* condition. There was only one statistically reliable difference: participants in the control condition rated food *B* as <u>more</u> valid (M = 78.58, SEM = 8.77) than those in the *A+* condition (M = 45.17, SEM = 10.44), $F(1, 34) = 6.24$, $p < .05$. No other differences were reliable, all p's > .10.

In summary, although participants learned about the presented contingencies, there was no compelling evidence of first-, second-, or third-order revaluation. Namely, although there was evidence for downward revaluation, the effect was uniform across all conditions and orders of association and was inconsistent with the seesaw pattern definitive of higher-order retrospective revaluation (see Experiment 1C).

Discussion of Experiments 2A, 2B, and 2C

Experiments 2A, 2B, and 2C were designed as analogues to Experiments 1A, 1B, and 1C, to examine first-, second-, and third-order retrospective revaluation using the allergist task. In addition to solving the major confound I saw in traditional explorations of revaluation using the allergist task, namely the lack of independent control groups, these last three experiments sought to assess whether the results of De Houwer and Beckers (2002a) could generalize to a different procedure. Overwhelmingly, the results show that they cannot.

Considered in isolation, the results of Experiments 2A, 2B, and 2C are a collection of unreliable results and large standard errors that pose challenges to interpretation. There is, however, a striking consistency in the results over the three experiments. Namely, with one exception, participants in the control condition always rated a food that was a first-, second-, or third-order associate of food *A* as more valid than participants in either the *A+* or *A+* treatment condition. The only exception is in the results of Experiment 2B, where participants in both the *A+* and *A-* conditions rated food *B* as more rather than less valid than those in the control condition.

The results of Experiments 1A, 1B, and 1C differ from those of Experiments 2A, 2B, and 2C. But why? First, the learning demands are substantially different between the two tasks. Participants in Experiment 1A learned about 4 weapons and participants in Experiment 2A learned about 11 foods. Participants in Experiments 1B learned about 5 weapons and participants in Experiment 2B learned about 17 foods. Participants in Experiments 1C learned about 6 weapons and participants in Experiment 2C learned about 19 foods, respectively. It is possible that the learning demands in the allergist task were too high to support revaluative learning. However, others have observed retrospective revaluation with even stronger learning demands (e.g., Larkin et al., 1998; Mutter et al., 2012).

Second, the nature of training differed between the tank task and the allergist task. In the tank task, participants observed the contingencies. In the allergist task, participants predicted the contingencies. However, in an additional experiment conducted but not reported here, participants observed rather than predicted the contingencies in the allergist task. The results of that experiment were consistent with the results of Experiment 2A (see Figure 5).

Third, the tasks also differed in terms of cover story and materials. It is possible that participants imported pre-existing knowledge or beliefs about foods into their test ratings in the allergist task, resulting in a contamination of the results (see Jamieson & Mewhort, 2010).

A fourth difference, and I speculate the one most likely responsible for the discrepancy of results, was a difference in task instructions. In the instructions for the tank task, participants were told that when two weapons fired simultaneously and the tank exploded, only one weapon was responsible. This was reinforced during the training phase by the presentation of a “IMPACT 10/20” feedback message. In the allergist task, there was no such instruction. The instructions in the tank task encouraged an *elemental* assessment of the weapons’ efficacy. The instructions in the allergist task, by omission, allowed for a *configural* assessment of the foods’ predictive validity. For example, when apples and bananas caused an allergic reaction, participants were free to draw any number of conclusions. Perhaps apples caused an allergic reaction independent of bananas; perhaps bananas caused an allergic reaction independent of apples; or perhaps apples and bananas together interacted to cause an allergic reaction. Because participants’ interpretation was not under experimental control, it is not entirely surprising that their behaviour was not under experimental control. In fact, this insight has been noted previously by Williams et al. (1994) and by De Houwer, Beckers, and Glautier (2002).

In retrospect, it is not entirely surprising that the results of Experiments 1A, 1B, and 1C differ from the results of Experiments 2A, 2B, and 2C. The two tasks differ in so many ways that retrospective revaluation would need to be an extremely robust phenomenon to generalize across the two, whereas it is known to be a fragile phenomenon.

In spite of their structural similarities, the two sets of results stand in stark opposition to one another. Whereas Experiments 1A, 1B, and 1C showed consistent first-order revaluation and weakening of revaluation at the higher orders, the allergist task failed to show reliable revaluation at even the first order of learning. The one consistency between the two sets of experiments was a general trend for increasing standard error with increasing order of learning, indicating that participants agreed less with one another as the learning became more complex. This is hardly a surprising result, but it is one that De Houwer and Beckers (2002a) fail to acknowledge in their own work—nor do they need to, because their more liberal comparisons between ratings in the *A+* and *A-* conditions allow them to overlook the dilution of retrospective revaluation with increasing order of association.

The full set empirical results do not serve to clarify the empirical reliability of retrospective revaluation: depending on the procedure used, the results do or do not show compelling evidence for revaluation. However, the findings provide constraints for computational models seeking to explain learning. Namely, the full set of results suggests a need for a model that either succeeds or fails to accommodate first-, second-, and third-order retrospective revaluation, depending on the task used to invoke learning. For obvious reasons, a model that explains contradictory results on the basis of changes in task detail is neither feasible nor particularly desirable; at the very least, any such model would not be addressing the structure of basic mechanisms of human learning, but rather task-specific peculiarities.

In order to understand the mechanisms that may be necessary to make sense of revaluation, I present two formal models of learning and memory and explore their respective abilities to account for the two sets of results. I contrast a post-neoclassical theory of learning based on the Rescorla-Wagner (1982) model of learning against an unorthodox model of learning based on the Minerva 2 model of human episodic memory (Hintzman, 1986). The computational findings bear important ramifications for theoretical considerations of retrospective revaluation learning in humans.

Part 2: Computational Modelling

Associative learning has been described in various terms. The classical theories model learning as the formation and updating of associative bonds between stimulus elements—the scheme well described by the Rescorla-Wagner (1972) model and those that follow from it (e.g., Ghirlanda, 2005; Van Hamme & Wasserman, 1994; Witnauer & Miller, 2011). An alternative approach supposes that learning is a corollary of memory. For example, Jamieson et al. (2010, 2012) argued that associative learning is a process of cued recall, in which learning about a cue-outcome relationship is measured by the cue's ability to retrieve the target outcome from memory. Still another approach involves the formation and use of symbolic propositional structures (De Houwer, 2009). Despite differences amongst the various models, they all consider associative learning a primitive of cognition in both human and nonhuman animals.

In the section that follows, I will test two models of associative learning against the experimental data collected in Part 1. The first is a classical learning approach recast as a single layer neural network, an elaborated and modern extension of the classical Rescorla-Wagner (1972) account (Ghirlanda, 2005; Ghirlanda & Enquist, 1999, 2007). The other is Jamieson et al.'s (2010, 2012) memory-based theory of associative learning, an elaborated and modern extension of Hintzman's (1986) Minerva 2 model. The contrast is particularly meaningful in light of the empirical problem I am trying to solve. In short, do we need a theory of memory to account for retrospective revaluation? Or is a post-neoclassical theory of learning sufficient?

I simulate all six of the experiments reported in Part 1 with both models. I begin by describing the models mechanistically and computationally, and then report computer simulations with the two models. As will be illustrated, Minerva-AL fits the patterns of data observed in Experiments 1A, 1B, and 1C, whereas the Ghirlanda model fits the patterns of data observed in Experiments 2A, 2B, and 2C. I will explain the differences between the models and their patterns of behaviour after the simulation work is reported.

Computational Models

Minerva-AL

Minerva-AL (Associative Learning; Jamieson, Crump, & Hannah, 2012) is an instance-based model of associative learning built on the theoretical framework provided by the Minerva 2 (Hintzman, 1984) model of human episodic memory.

According to Minerva-AL, each individual experience is represented in memory by a unique trace. All retrieval from memory is cue-driven and parallel. When a cue (i.e., a memory probe) is presented, it activates all traces in memory. Each trace's activation is in proportion to its similarity to the probe. The information retrieved from memory is the sum of the activated traces, called the echo. Because the probe retrieves traces similar to it, a probe will retrieve a representation of itself from memory. This is how the model accomplishes recognition. Because it retrieves whole traces, a probe also retrieves events it has co-occurred with in the past. This is how the model accomplishes cued recall, the process it uses to model associative learning.

In the Minerva-AL model, a stimulus or event is represented by a vector of *n* elements. Each feature takes one of two discrete values: +1 or 0. A value of +1 indicates the feature is relevant to the stimulus description; a value of 0 indicates the feature is either indeterminate or irrelevant to the stimulus description. Co-occurrence of events is represented by summing event representations to form a single vector. For example, if two events A = [0,0,0,0,0,1,1,0,0,0,0,0] and B = [0,0,0,0,0,0,0,0,0,1,1,0] co-occur, their co-occurrence is represented as $AB = A + B$ = [0,0,0,0,0,1,1,0,0,1,1,0].

Memory, M, is a two dimensional matrix. Each row in the matrix stores an instance. Each column corresponds to a feature. Memory for a trial is determined by taking the difference between the event vector and the echo retrieved. By encoding differences between the event vector and the echo, memory of preceding trials (i.e., represented in the echo) has influence on what is learned on the current trial (i.e., represented in the event vector). Unexpected information (i.e., information in the event vector that is not retrieved in the echo) is encoded more strongly than expected information (i.e., information in the event vector that is retrieved in the echo).

This operation, known as discrepancy-encoding, is implemented using subtraction,

$$M_{ij} = E_j - C'_j,$$

where i indexes the row in memory, j indexes the features of the vector representations, M is the memory matrix, E is the event vector, and C' is the echo, normalized to the range {-1, +1}. Variation in encoding quality is controlled by a parameter, L, whereby $M_{ij} = E_j - C_j'$ with probability L and $M_{ij} = 0$ with probability $1 - L$.

In the model, all retrieval is cued. When a cue is presented, it activates each memory trace in proportion to its similarity to the cue. The similarity of the probe, P, to trace i in memory, M_i, is computed as,

$$S_i = \frac{\sum_{j=i}^{n} P_j \times M_{ij}}{\sqrt{\sum_{j=1}^{n} P_j^2} \times \sqrt{\sum_{j=1}^{n} M_{ij}^2}},$$

where P_j is the value of the j^{th} feature in the probe, M_{ij} is the value of j^{th} feature of the i^{th} row in memory, and n is the number of features in the vectors under comparison.

Trace i's activation, A_i, is a nonlinear function of its similarity to the probe,

$$A_i = S_i^3 .$$

In principle, the probe activates all traces in memory. However, the nonlinear activation function ensures that traces very similar to the probe are activated much more strongly than traces that are moderately similar or that are dissimilar to the probe.

The information that a probe retrieves from memory is a vector, C, called the echo. Element j in the echo is equal to the sum of the corresponding weighted elements in the $i = 1...m$ traces in memory,

$$C_j = \sum_{i=1}^{m} A_i \times M_{ij} .$$

A randomly sampled value from the interval [-0.001, +0.001] is added to each element in the echo to introduce noise and slow the learning. This better reflects the multi-trial nature of associative learning than the single-trial learning for which the original Minerva 2 model was designed. Without the addition of noise, the learning is too fast.

Strength of responding is indexed in two steps. Values in the echo are normalized,

$$C'_j = \frac{C_j}{\max\left|C_{1,n}\right|},$$

then similarity is computed between the normalized echo and the relevant target associate,

$$Z|P = \frac{\sum_{j=1}^{n} Z_j \times C'_j}{n_R},$$

where P is the probe, Z is the target associate, and n_R is the number of nonzero features in Z

and *C'*. The value *Z|P* is read "retrieval of *Z* given *P*." The larger the absolute value of *Z|P,* the better that *Z* is retrieved by *P*. If the probe retrieves *Z* perfectly, then *Z|P* = 1. If the probe retrieves an inverse (i.e., opposing) representation of *Z*, then *Z|P* = -1. If the probe does not retrieve *Z*, then *Z|P* = 0.

I will adopt a single representation scheme over all the simulations that I report with Minerva-AL. In all cases, events of a trial are coded in an event vector composed of twenty-six successive 20-element subfields (i.e., an event vector has 520 elements in total). Each subfield corresponds to a cue, the representation of a context, or an outcome. For example, the first, second, third, and fourth subfields correspond to cues *A*, *B*, *C*, and *D* respectively, and so on to the twenty-fourth subfield. The twenty-fifth subfield, elements 481 through 500, corresponds to a context cue, *Y*, which is present on all trials. The twenty-sixth subfield, elements 501 through 520, corresponds to the target outcome *Z*. The representation of a cue, context, or outcome is accomplished by assigning the value +1 to the relevant elements and 0 to all others. Thus, cue *A* is represented by a vector of 520 elements, where elements 1 through 20 take a value of +1 and all other elements take a value of 0. The context is represented by a vector of 520 elements, where elements 481 to 500 take a value of +1 and all other elements take a value of 0. The outcome, *Z*, is represented by a vector of 520 elements, where elements 501 to 520 take a value of +1 and all other elements take a value of 0. Whereas the presence cues and outcomes will vary from trial to trial in a procedure, the context is constant. For all simulations, the parameter *L* was set to 1.0, indicating no loss of information.

Ghirlanda model

The Ghirlanda model (2005, but see also Ghirlanda & Enquist, 1999, 2007) is a recent adaptation of the original Rescorla-Wagner (RW; Rescorla & Wagner, 1972) model and fits firmly within the post-neoclassical approach to associative learning.

However, whereas the original RW model posited a single node to represent each stimulus, Ghirlanda's revision adopts a vector-based representation scheme. The change in representation brings the model more in line with models of memory, for which vector representation is the norm.

According to the model, each individual cue is a vector of 460 elements. Element *i* in a stimulus representation, *S*, is defined by the equation,

$$S_i(\beta,\gamma,\mu,\sigma) = \beta + \gamma e^{-(\mu-i)^2/2\sigma^2},$$

where β is baseline activity in the network, γ is the maximum intensity of the stimulus element, μ is the peak element of the distribution, σ is the standard deviation of the distribution, and *e* is Euler's constant (in the present simulations, the constant is precise to ten decimal places). For all stimuli, I hold $\beta = .10$, $\gamma = .35$, and $\sigma = 5.0$ constant.

For present purposes, I specify representations for cues *A* through *V*, with an additional cue *Y* to represent the background context in which each trial occurs. Each cue is assigned to a number: cue *A* = 1, cue *B* = 2, … cue *V* = 22. The value of μ for individual cues is determined as $\mu = 20S$, where *S* indicates the number assigned to the cue. Cue *Y* is assigned as the 23rd cue, and $\mu_Y = \beta$. A compound stimulus, for example *AB*, is represented by forming a new vector, *S*, where element *i* in *S* is defined as the maximum value of *S* for its constituent elements, $S_i = Max(A_i, B_i)$. In this way, learning about a compound stimulus impacts the representations of all individual elements in the compound. Figure 8 presents concrete illustrations of stimulus representations in the Ghirlanda model: the left Panel shows the representation for stimulus *A*, the middle Panel shows the representation for stimulus *B*, and the right Panel shows the representation for the compound *AB,* determined by the equation $S_i = Max(A_i, B_i)$.

Figure 8. *Stimulus representations in the Ghirlanda (2005) model.*

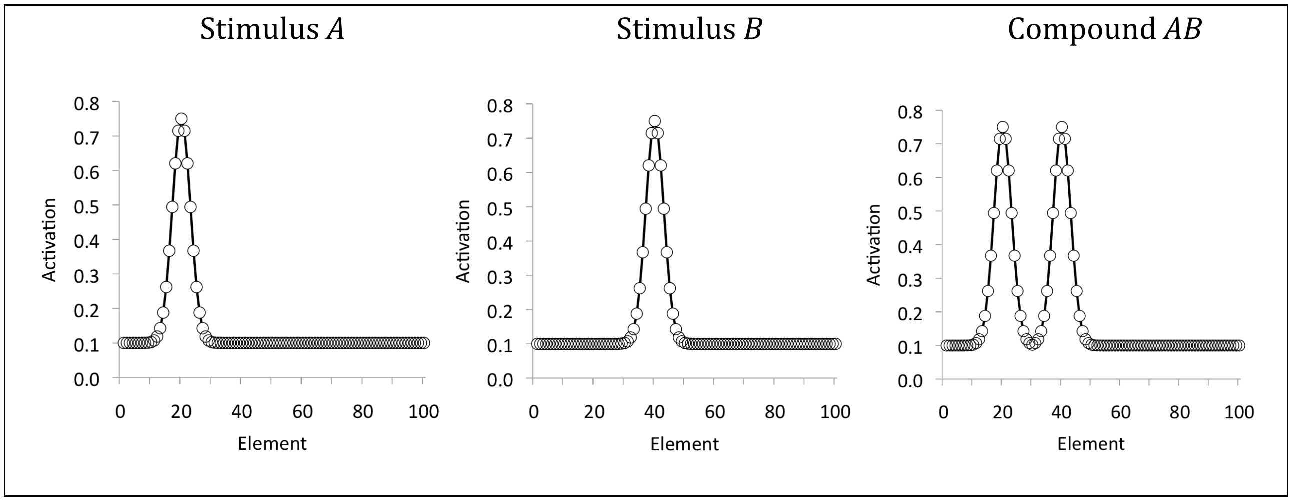

Ghirlanda's representation scheme makes four assumptions. First, the activity of each element must be a positive number. Second, a stimulus elicits varying degrees of activation from different elements. Third, stimuli that are physically more similar will activate more common elements, while stimuli that are physically more distinct will activate fewer common elements. Finally, the strength of element activity is a direct function of the intensity of the stimulus.

The model is a fully-connected single-layer feed-forward artificial neural network with 520 input units and 1 output unit. The strength of responding to a given stimulus is measured by a dot-product,

$$r_s = \sum W_i S_i,$$

where S_i is the activity induced in element i of the input layer by stimulus S; W_i is the connection weight from input unit i of the input layer to the output unit; and r_S is the sum of activation at the output unit caused by applying stimulus S to the input layer. In the simulations that I report, r_S is the measure of response strength.

Learning in the model uses Rescorla-Wagner learning rule. Expressed in the terms of the network, learning operates by,

$$W_i^{n+1} = W_i^n + \alpha S_i(\lambda - r_s),$$

where W_i^{n+1} is the connection weight from input unit i of the input layer to the output unit on the current trial; W_i^n is the connection weight from input unit i of the input layer to the output unit on the previous trial; α is a constant that specifies the speed of learning; S_i is the value of unit i in the input layer; λ is a constant that specifies the maximum possible value for r_S; and r_S is the strength of response on a given trial. Learning is a process of updating the weights connecting the input units to the output unit, where the update to weight W_i is equal to the product of the learning rate, the value of input unit i, and the discrepancy between λ and r_S. Therefore, when λ is equal to r_S, no learning takes place.

Before continuing, the model is constrained by a basic computational issue. To simulate even the simplest learning the model must be presented with not only the learning contingencies, but also the context Y in the absence of other stimulation (i.e., Y-). As described above, $\mu_Y = \beta$.

Ghirlanda (2005) has used the model to simulate a number of classical learning phenomena that have posed challenges to the RW model. Germane to the present discussion, he used the model to simulate examples of first-order retrospective revaluation. For the model to fit a result, there are two constraints that need to be concurrently satisfied. To illustrate, consider the *A+* condition from first-order retrospective revaluation (see Experiments 1A and 2A). Under this learning condition *AB+, Y-* learning is followed by *A+, Y-* learning. The compound *AB* provokes responding from the participant, such that the weights must satisfy the equation

$$\sum_i W_i AB_i = r_{AB} >> 0,$$

where ">>" means *substantially greater than*. Likewise, the participant will not respond when neither *A* nor *B* is present (i.e., when only the background context is present), such that the weights must satisfy the equation

$$\sum_i W_i Y_i = r_Y \approx 0,$$

where "$\approx$" means *close to*. Since one assumption stated that that the activity of each element must be positive, and since it is necessarily the case that $Y_i = 0$, some weights must be positive and others negative to satisfy the equation for r_y above. However, elements that carry negative weights cannot be the same elements most excited by the stimuli *A* and *B*, because r_{AB} must be substantially greater than zero. Therefore, in order to satisfy the equations for r_y and r_{AB} simultaneously, those elements most excited by *A* and *B* must carry positive weights. A distribution of stimulus representations at this point shows equal levels of positive activation for cues *A* and *B* (see the left Panel of Figure 9).

Figure 9. *Changes in stimulus representations in the Ghirlanda (2005) model over the course of retrospective revaluation learning.*

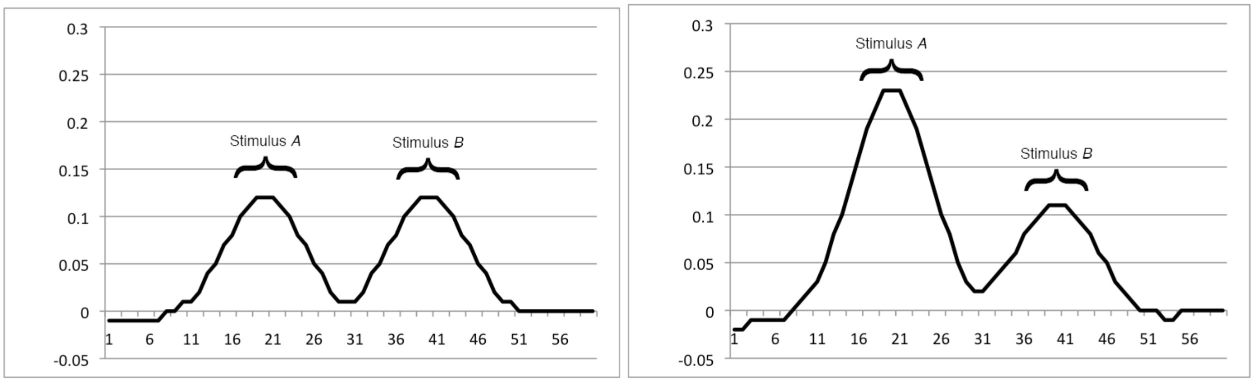

Left panel: Representation for cues *A* and *B* following *AB+, Y-* training. ***Right panel:*** Representation for cues *A* and *B* following *AB+, Y-* and subsequent *A+, Y-* training.

During the second phase of learning (i.e., *A+, Y-*) responding to cue *A* will increase. To continue to satisfy the equation for r_y, however, other weights must decrease as weights for *A* increase. Thus, some of the weights for elements related to cue *B* will decrease in activation, leading to a downward revaluation of cue *B* (see the right Panel of Figure 9).

In the simulations that follow, the maximal value of responding to *S* is set to $\lambda = 0.75$ for trials on which the outcome occurs (i.e., "+" trials) and $\lambda = 0.05$ for trials on which the outcome does not occur (i.e., "-" trials). The model undergoes 10,000 blocks of learning, each of which involves one instance of each trial. Importantly, each block involves a trial of learning about the context cue *Y*, with $\lambda = 0.05$. Associative learning is measured by the strength of responding to the outcome *Z*, given a stimulus.

Experimental simulations

Experiment 1A

Experiment 1A tested first-order retrospective revaluation in the tank task. The experiment included three experimental conditions. In a control condition, participants learned *AB+*, *C-*, *D+* before going to test. In an *A+* condition participants learned *AB+*, *C-*, *D+* and then *A+*, *C-*, *D+* before going to test. In an *A-* condition, participants learned *AB+*, *C-*, *D+* and then *A-*, *C-*, *D+* before going to test (see top panel in Table 1). Participants' behaviour showed evidence of first-order retrospective revaluation. Specifically, participants in the *A+* treatment condition revalued cue *B* (the first-order associate to cue *A*) downward following *A+* training and participants in the *A-* condition revalued cue *B* upward following *A-* training.

I applied the Minerva-AL and Ghirlanda models to the three treatment protocols from Experiment 1A. Simulation with the Minerva-AL model involved 25 independent simulations of each training protocol (L = 1.0), with model performance reported as an average over the 25 simulations. Simulation with the Ghirlanda model involved just one simulation of each protocol (*n blocks* = 10,000), because the model always arrives at the same outcome.[2]

Results of simulations with both models are shown in Figure 2. Empirical data are shown in the top Panel. Data from the simulations of Minerva-AL are shown in the bottom

---

[2] The difference in numbers of simulations reflects a difference in the models. The Minerva-AL model involves a stochastic component so that it is necessary to run multiple simulations and consider the performance averaged over those simulations. The Ghirlanda model, by contrast, is run to asymptote so that the result of one simulation is identical that of the next.

left Panel. Data from the simulations of the Ghirlanda model are showing in the bottom right Panel.

Minerva-AL fit the pattern of empirical data better than the Ghirlanda model (R^2 = .94 versus R^2 = .92, respectively). However, only the Minerva-AL model accommodates first-order revaluation. According to the Minerva-AL model, *A+* training forces a downward revaluation of cue *B* (the first-order associate to cue *A*) and *A-* training forces an upward revaluation of cue *B*. According to the Ghirlanda model, training with both *A+* and *A-* does not force any compelling revaluation of cue *B*. Whereas training with *A+* forces a downward revaluation of cue *B* relative to control, and training with *A-* forces an upward revaluation of cue *B* relative to control, the magnitude of that revaluation is in no way similar to the magnitude of revaluation seen in the empirical data.

On the topic of magnitude, Minerva-AL anticipates stronger downward revaluation of *B* following *A+* training than upward revaluation of *B* following *A-* training. However, the empirical showed equal upward and downward revaluation of *B* following *A+* and *A-* training. The discrepancy shows that whereas Minerva-AL predicts both upward and downward first-order revaluation, it does not do so perfectly.

Whereas the two models did disagree on the revaluation of cue *B*, they were in strong agreement on cues *A*, *C*, and *D*. Both models rated cue *A* as stronger following *A+* training and weaker following *A-* training, both rated cue *B* as strongest following *A-* training, and both rated cue *C* as poorly predictive of and cue *D* as strongly predictive of the outcome.

In summary, the Minerva-AL and Ghirlanda models offer the same predictions except for the retrospective revaluation of cue *B*. Whereas the Minerva-AL model predicts revaluation of *B*, the Ghirlanda model does not. Although I will return to a discussion of how the two models differ later, the difference in the two models' predictions is a reflection of how the two theories conceive of learning. The Minerva-AL model treats learning as a corollary of memory. The Ghirlanda model—like the model of Rescorla and Wagner (1972)—does not (see Miller, 2006; Miller, Barnet, & Grahame, 1995). The results suggest that an explanation of retrospective revaluation requires a consideration of memory.

Experiment 1B

Experiment 1B tested second-order retrospective revaluation in the tank task. The experiment included three experimental conditions. In a control condition, participants learned *AB+*, *D-*, *E+* and then *BC+, D-, E+* before going to test. In an *A+* condition participants learned *AB+*, *D-*, *E+,* then *BC+, D-, E+,* and then *A+*, *D-*, *E+* before going to test. In an *A-* condition, participants learned *AB+*, *D-*, *E+,* then *BC+, D-, E+,* and then *A-*, *D-*, *E+* before going to test (see middle Panel in Table 1). Participants' behaviour showed evidence of both first- and second-order retrospective revaluation. Specifically, participants in the *A+* treatment condition revalued cue *C* (the second-order associate to cue *A*) upward following *A+* training and participants in the *A-* condition revalued cue *C* downward following *A-* training. Likewise, participants in the *A+* treatment condition revalued cue *B* (the first-order associate to cue *A*) downward following *A+* training and participants in the *A-* condition revalued cue *B* upward following *A-* training.

I applied the Minerva-AL and Ghirlanda models to the three treatment protocols from Experiment 1B. Simulation with the Minerva-AL model involved 25 independent simulations of each training protocol (L = 1.0). Simulation with the Ghirlanda model involved just one simulation of each protocol (*n blocks* = 10,000).

Results of simulations with both models are shown in Figure 3. Empirical data are shown in the top Panel. Data from the simulations of Minerva-AL are shown in the bottom left Panel. Data from the simulations of the Ghirlanda model are showing in the bottom right Panel.

Whereas both models fit the overall empirical data fairly well (R^2 = .97 and R^2 = .85, for Minerva-AL and the Ghirlanda model, respectively), only Minerva-AL accommodated first- and second-order revaluation.

As shown in the bottom left panel of Figure 3, the Minerva-AL model anticipates stronger downward revaluation of *C* (the second-order associate to cue *A*) following *A*- training than upward revaluation of *C* following *A+* training. However, the match to empirical data is not perfect as participant behaviour showed the opposite trend: participants showed stronger upward revaluation of *C* following *A+* training than downward revaluation of *C* following *A*- training. The Minerva-AL model also anticipates stronger downward revaluation of *B* (the first-order associate to cue *A*) following *A+* training than upward revaluation of *B* following *A*- training. However, participants showed equally strong upward and downward revaluation of *B* following *A+* and *A*- training. In conclusion, whereas Minerva-AL predicts retrospective revaluation at the first- and second orders, it fails to match the empirical data on differences in magnitudes of upward and downward revaluation.

The Ghirlanda model, on the other hand, failed to capture the empirical data at all. First, there was no sign of first- or second-order revaluation following *A+* training. Second, the model's performance following *A*- training showed a confusing pattern of revaluation.

Namely, the model revalued both cues *B* and *C* downward following *A*- training, at odds with the seesaw pattern that is the signature of second-order revaluation.

Whereas the two models did disagree on the revaluation of cues *B* and *C*, they were in strong agreement on cues *A*, *D,* and *E*. Both models rated cue *A* as stronger following *A+* training and weaker following *A*- training, and both rated cue *D* as poorly predictive of and cue *E* as strongly predictive of the outcome.

In summary, the Minerva-AL and Ghirlanda models offer the same predictions except for the retrospective revaluation of cues *B* and *C*. Whereas the Minerva-AL model predicts first- and second-order revaluation, the Ghirlanda model does not.

Experiment 1C

Experiment 1C tested third-order retrospective revaluation in the tank task. The experiment included three experimental conditions. In a control condition, participants learned *AB+*, *E-*, *F+,* and then *BC+, CD+, E-, F+* before going to test. In an *A+* condition participants learned *AB+*, *E-*, *F+,* then *BC+,CD+, E-, F+,* and then *A+*, *E-*, *F+* before going to test. In an *A*- condition, participants learned *AB+*, *E-*, *F+,* then *BC+, CD+, E-, F+,* and then *A-*, *E-*, *F+* before going to test (see bottom Panel in Table 1).

Participants' behaviour showed partial support for third-order retrospective revaluation. Specifically, participants in the *A+* treatment condition revalued cue *D* (the third-order associate to cue *A*) downward following *A+* training and participants in the *A*- condition revalued cue *D* upward following *A*- training. The data moreover showed partial evidence of second-order retrospective revaluation. Participants in the *A+* treatment condition revalued cue *C* (the second-order associate to cue *A*) upward following *A+* training and participants in the *A*- condition revalued cue *C* downward following *A*- training.

Likewise, the data showed partial evidence of first-order retrospective revaluation. Participants in the *A+* treatment condition revalued cue *B* (the first-order associate to cue *A*) downward following *A+* training and participants in the *A-* condition revalued cue *B* upward following *A-* training.

I applied the Minerva-AL and Ghirlanda models to the three treatment protocols from Experiment 1C. Simulation with the Minerva-AL model involved 25 independent simulations of each training protocol (L = 1.0). Simulation with the Ghirlanda model involved just one simulation of each protocol (*n blocks* = 10,000).

Results of simulations with both models are shown in Figure 3. Empirical data are shown in the top Panel. Data from the simulations of Minerva-AL are shown in the bottom left Panel. Data from the simulations of the Ghirlanda model are showing in the bottom right Panel.

Whereas both models fit the overall empirical data fairly well (R^2 = .99 and R^2 = .96 for Minerva-AL and the Ghirlanda model, respectively), neither did an excellent job of accommodating third-order revaluation. Targeted inspection of the model predictions shows that whereas the Minerva-AL model anticipates a reversal of third-order revaluation, the Ghirlanda model fails to predict any revaluation.

As shown in the bottom left Panel of Figure 4, the Minerva-AL model incorrectly predicts upward revaluation of cue *D* (the third-order associate to cue *A*) following *A+* training and downward revaluation of cue *D* following *A-* training; the empirical data showed the opposite. The Minerva-AL model correctly anticipates revaluation of cues *B* (the first-order associate to cue *A*) and *C* (the second-order associate to cue *A*) following *A-*

and *A+* training. However, the magnitude of revaluation is in all cases weaker than the revaluation seen in the empirical data.

In contrast to the Minerva-AL model, the Ghirlanda model failed to show any degree of revaluation of cues *D, C,* or *B* following *A+* or *A-* training. In fact, additional training with *A+* or *A-* had only a modest uniform influence on ratings of the first-, second-, and third-order associates of cue *A*.

Whereas the two models did disagree on the revaluation of cues *B*, *C*, and *D*, they were in strong agreement on cues *A*, *E,* and *F*. Both models rated cue *A* as stronger following *A+* training and weaker following *A-* training, and both rated cue *E* as poorly predictive of and cue *F* as strongly predictive of the outcome.

In summary, the Minerva-AL and Ghirlanda models offer the same predictions except for the retrospective revaluation of cues *B*, *C,* and *D*. Whereas the Minerva-AL model predicts first- and second-order revaluation, the Ghirlanda model does not. Both models fail to capture third-order revaluation.

Experiment 2A

Experiment 2A tested first-order retrospective revaluation in the allergist task. The experiment included three experimental conditions. In a control condition, participants learned *AB+*, *CD+* before going to test. In an *A+* condition participants learned *AB+*, *CD+* and then *A+*, *E+* before going to test. In an *A-* condition, participants learned *AB+*, *CD+* and then *A-*, *E+* before going to test (see top Panel in Table 2). Participants' behaviour showed some evidence of first-order retrospective revaluation. Specifically, participants in the *A+* treatment condition revalued cue *B* (the first-order associate to cue *A*) downward following

A+ training, but contrary to predictions of revaluation participants in the *A-* condition also revalued cue *B* downward following *A-* training.

I applied the Minerva-AL and Ghirlanda models to the three treatment protocols from Experiment 2A. Simulation with the Minerva-AL model involved 25 independent simulations of each training protocol (L = 1.0). Simulation with the Ghirlanda model involved just one simulation of each protocol (*n blocks* = 10,000).

Results are shown in Figure 5: empirical data in the top Panel, simulated data from the Minerva-AL model in the bottom left Panel, and simulated data from the Ghirlanda model in the bottom right Panel. The Minerva-AL model fits the pattern of empirical data better than the Ghirlanda model (R^2 = .99 versus R^2 = .96, respectively). However, targeted inspection of the model predictions shows that whereas the Minerva-AL predicts first-order revaluation, the Ghirlanda model better captures the participants' failure to upwardly revaluate cue *B* after *A-* training.

The Minerva-AL model anticipates stronger downward revaluation of *B* following *A+* training than upward revaluation of *B* following *A-* training. However, participants downwardly revalued cue *B* following *A-* training. In contrast to the Minerva-AL model, the Ghirlanda model predicted (weak) downward revaluation of cue *B* of equal strength following *A+* and *A-* training, consistent with the empirical data.

Outside of differences in how they did (or did not) revalue cue *B*, the two models are in strong agreement. Both models rated cue *A* as stronger following *A+* training and weaker following *A-* training, and both rated *Negative Control* cues as poorly predictive of and *Positive Control* cues as strongly predictive of the outcome.

In summary, the Minerva-AL and Ghirlanda models offer the same predictions except for the first-order revaluation cue *B*. Whereas the Minerva-AL model predicts revaluation of *B* following both *A+* and *A-* training, the Ghirlanda model only does so following *A+* training. The failure to upwardly revaluate cue *B* following *A-* training was a critical feature of the participant data, and only the Ghirlanda model predicts it.

Experiment 2B

Experiment 2B tested second-order retrospective revaluation in the allergist task. The experiment included three experimental conditions. In a control condition, participants learned *AB+*, *EF+* and then *BC+, FG+* before going to test. In an *A+* condition participants learned *AB+*, *EF+,* then *BC+, FG+,* and then *A+*, *D+* before going to test. In an *A-* condition, participants learned *AB+*, *EF+,* then *BC+, FG+,* and then *A-*, *D-* before going to test (see middle Panel in Table 2). Participants' behaviour showed insufficient evidence to support second-order retrospective revaluation. Specifically, participants in both the *A-* and *A+* conditions revalued cue *C* (the second-order associate to cue *A*) downward following the final phase of training, but in neither case was the difference significant. The data moreover showed insufficient evidence to support first-order retrospective revaluation. Again, participants in both the *A-* and *A+* conditions revalued cue *B* (the first-order associate to cue *A*) upward following the final phase of training, but in neither case was the difference significant.

I applied the Minerva-AL and Ghirlanda models to the three treatment protocols from Experiment 2B. Simulation with the Minerva-AL model involved 25 independent simulations of each training protocol (L = 1.0). Simulation with the Ghirlanda model involved just one simulation of each protocol (*n blocks* = 10,000).

Results are shown in Figure 6: empirical data in the top Panel, simulated data from the Minerva-AL model in the bottom left Panel, and simulated data from the Ghirlanda model in the bottom right Panel. In short, the Ghirlanda model fits the pattern of empirical data better than the Minerva-AL model (R^2 = .96 versus R^2 = .95, respectively). Targeted inspection of the model predictions shows that whereas the Minerva-AL predicts partial second-order revaluation and full first-order revaluation, the Ghirlanda model better captures the participants' failure to reliably revaluate cues *C* and *B* in either group.

The Minerva-AL model predicted a downward revaluation of *C* (the second-order association to cue A) following *A-* training but no upward revaluation of *C* following *A+* training. However, participants downwardly revalued cue *C* in both conditions, though the differences were not statistically reliable. In contrast to the Minerva-AL model, the Ghirlanda model predicted (weak) downward revaluation of cue *C* of equal strength following both *A+* and *A-* training, consistent with the empirical data (see bottom right Panel of Figure 6). The Ghirlanda model moreover captures the depressed level of responding to cue *C* compared to other cues, as observed in the participant data. The Minerva-AL model fails to capture this result.

The Minerva-AL model also anticipates downward revaluation of *B* (the first-order associate to cue *A*) following *A+* training and upward revaluation of *B* following *A-* training to be equally reliable. However, participants upwardly revalued cue *B* in both conditions, though the differences were not statistically reliable. In contrast to the Minerva-AL model, the Ghirlanda model predicted (weak) downward revaluation of cue *B* of equal strength following both *A+* and *A-* training. Although this is not consistent with participant data, the pattern of observable but nominal differences between the three conditions is more consistent than the reliable revaluation predicted by the Minerva-AL model.

Outside of differences in how they did (or did not) revalue cues *C* and *B*, the two models are in strong agreement. Both models rated cue *A* as stronger following *A+* training and weaker following *A-* training, and both rated Negative Control cues as poorly predictive of and Positive Control cues as strongly predictive of the outcome.

In summary, the Minerva-AL and Ghirlanda models offer the same predictions except for first- and second-order retrospective revaluation. Whereas the Minerva-AL model predicts revaluation of *C* following *A-* training and revaluation of *B* following both *A+* and *A-* training, the Ghirlanda model predicts weak downward revaluation of *C* and *B* following both *A+* and *A-* training. Participants' failure to revaluate at either the first or second order was evident, and only the Ghirlanda model predicts the failures.

Experiment 2C

Experiment 2C tested third-order retrospective revaluation in the allergist task. The experiment included three experimental conditions. In a control condition, participants learned *AB+*, *EF+,* then *BC+, FG+,* and then *CD+, GH+* before going to test. In an *A+* condition participants learned *AB+*, *EF+,* then *BC+, FG+,* then *CD+, GH+,* and then *A+*, *I+* before going to test. In an *A-* condition, participants learned *AB+*, *EF+,* then *BC+, FG+,* then *CD+, GH+,* and then *A-*, *I-* before going to test (see bottom Panel in Table 2).

Participants' behaviour showed insufficient evidence to support third-order retrospective revaluation. Specifically, participants in both the *A-* and *A+* conditions revalued cue *D* (the third-order associate to cue *A*) downward following the final phase of training, but in neither case was the difference significant. The data likewise showed insufficient evidence to support second-order retrospective revaluation. Again, participants in both the *A*- and *A+* conditions revalued cue *C* (the second-order associate to cue *A*) downward following the final phase of training, but in neither case was the difference significant. The data was showed evidence in partial support of first-order retrospective revaluation. In particular, participants in the *A+* condition revalued cue *B* (the first-order associate to cue *A*) downward following *A+* training, but participants in the *A*-condition revalued cue *B* downward following *A*- training, although this difference was not statistically reliable.

I applied the Minerva-AL and Ghirlanda models to the three treatment protocols from Experiment 2C. Simulation with the Minerva-AL model involved 25 independent simulations of each training protocol (L = 1.0). Simulation with the Ghirlanda model involved just one simulation of each protocol (*n blocks* = 10,000).

Results are shown in Figure 7: empirical data in the top Panel, simulated data from the Minerva-AL model in the bottom left Panel, and simulated data from the Ghirlanda model in the bottom right Panel. Whereas the Ghirlanda and Minerva-AL models both fit the overall pattern of results well (R^2 = .91 versus R^2 = .89, respectively), targeted inspection shows that the Ghirlanda model offers a slightly more accurate fit.

The Minerva-AL model anticipates downward revaluation of cue *D* (the third-order associate to cue *A*) following *A*- training, but no revaluation of cue *D* following *A+* training. However, participant behaviour revealed (weak) downward revaluation of *D* following both *A+* and *A*- training, although neither difference was statistically reliable. In contrast to the Minerva-AL model, the Ghirlanda model predicted (weak) downward revaluation of cue *D* of equal strength following both *A+* and *A-* training, consistent with the empirical data (see bottom right Panel of Figure 6).

The Minerva-AL model likewise anticipates downward revaluation of cue *C* (the second-order associate to cue *A*) following *A-* training, but no revaluation of cue *C* following *A+* training. However, participant behaviour revealed (weak) downward revaluation of *C* following both *A+* and *A-* training, although neither difference was statistically reliable. In contrast to the Minerva-AL model, the Ghirlanda model predicted (weak) downward revaluation of cue *C* of equal strength following both *A+* and *A-* training, consistent with the empirical data.

The Minerva-AL model predicted (weak) downward revaluation of cue B (the first-order associate to cue A) following *A+* training, but no revaluation of cue *B* following *A-* training. This prediction was consistent with participant behaviour, which showed reliable downward revaluation of *B* following *A+* training, but no revaluation of cue *B* following *A-* training. In contrast to the Minerva-AL model, the Ghirlanda model predicted (weak) downward revaluation of cue *B* of equal strength following both *A+* and *A-* training, a finding that is not entirely consistent with the empirical data.

Outside of differences in how they did (or did not) revalue cues *D, C* and *B*, the two models are in strong agreement. Both models rated cue *A* as stronger following *A+* training and weaker following *A-* training, and both rated *Negative Control* cues as poorly predictive of and *Positive Control* cues as strongly predictive of the outcome.

In summary, the Minerva-AL and Ghirlanda models offer the same predictions except for first-, second-, and third-order retrospective revaluation. Whereas Minerva-AL predicts revaluation of cues *D* and *C* following *A-* training and revaluation of *B* following *A+*

training, Ghirlanda predicts equally weak downward revaluation of cues *D, C* and *B* following both *A+* and *A-* training. Participants' failure to revaluate at either the third or second order was evident, and only the Ghirlanda model predicts the failures. It is notable, however, that the Minerva-AL model better predicts participants' behaviour at the first order of learning.

Discussion of Computational Modelling

I have applied a post-neoclassical learning model based on the Rescorla-Wagner (1972) model and an unorthodox model of associative learning based on a model of episodic memory to the empirical data. The results from the Minerva-AL model (Jamieson, Crump, & Hannah, 2012) map on to the pattern of results from first- (Experiment 1A) and second-order revaluation (Experiment 1B), but fail to clearly capture third-order revaluation (Experiment 1C). On the other hand, the Ghirlanda model (Ghirlanda, 2005; Ghirlanda & Enquist, 1999, 2007) captures the particular patterns of failure in first- (Experiment 2A), second- (Experiment 2B), and third-order (Experiment 2C) revaluation in the allergist task.

I conclude that in order to accommodate retrospective revaluation, one must consider the role of memory. The conclusion is consistent with De Houwer's (2009) theoretical account of revaluation, but has a distinctly different flavor. In De Houwer's theory, retrospective revaluation is the outcome of explicit propositional reasoning. In Minerva-AL, that reasoning manages to emerge from the mechanics of the storage and retrieval of instances. In essence, Minerva-AL manages to behave *as if it were* conducting propositional reasoning, when in fact it is not. But how does behaviour that is consistent with an account based on propositional reasoning emerge from the storage and deployment of instances?

To explain, consider how Minerva-AL explains the devaluation of cue *B* following *AB+, A+* training in first-order revaluation. In phase 1 of training, cue *A* is established as a retrieval cue for *B* (i.e., a within-compound association). Consequently, in phase 2, cue *A* retrieves *B* into the echo. Because cue *B* was retrieved into the echo but was not presented, however, the trace stored to memory records the absence of *B* (i.e., its inverse, -*B*) paired with the presence of *Z* (i.e., +*Z*). This information establishes the conditions for producing retrospective revaluation at test.

When cue *B* is presented at test, the traces that contain its inverse (i.e., the traces that code -*B*) are activated. But, because the similarity of *B* to -*B* is a negative value, and during retrieval a trace is multiplied by its activation, traces with a negative representation of cue *B* and a positive representation of *Z* are inverted (i.e., a trace that joins -*B* and +*Z* in memory is activated as +*B* and -*Z*). The sum of the inverted traces collected into the echo yields a negative representation of *Z* in the echo, thereby producing the result.

Critically, the model's explanation for this revaluation effect is that the retrieval process collects inverse representations of *Z* into the echo, even though the memory matrix contains no inverse representations of *Z* itself. Thus, at a deeper level, the example underscores an important insight from the Minerva-AL model. Learning is not an encoding effect or a retrieval effect alone. Rather, learning emerges from an interaction of encoding and retrieval.

At an even deeper level, this provides a description of how Minerva-AL "reasons" outside of symbolic structure. The formation of expectations not only allows the model to anticipate events that are presented, but also to recognize that events ought to have been presented on the basis of previous experience. By recognizing violations of its expectancies, the model can update its understanding of those events that were expected but unpresented. Critically, these behaviours of the model do not require a homunculus or a phenomenological description. Rather, they are a corollary of the mathematics describing the storage and retrieval from memory.

The Ghirlanda model, on the other hand, is unable to accommodate retrospective revaluation. By my analysis, it fails because it has no memory for specific trials. The argument has been made previously by Miller (Miller, 2006; Miller, Barnet, & Grahame, 1995) under the *assumption of path independence*. Namely, the RW model and the neoclassical and post-neoclassical models derived from it deny that a learner remembers the events of separate learning trials. The assumption is important because it divorces the behaviour of learning from memory of what has been learned. For example, after learning that cue *A* predicts the outcome *Z*, then learning that cue *A* does not predict the outcome *Z*, the RW model and its successors fail to capitalize on previous learning when given the task of re-learning that cue *A* predicts the outcome *Z*. That is, a failure to remember previous learning. This assumption has obvious implications for the models' ability to accomplish retrospective revaluation. If the model cannot remember the cues that are not being presented, how can it revalue its knowledge of those cues?

Nevertheless, when retrospective revaluation is not observed in the empirical data (i.e., Experiments 2A, 2B, and 2C), the RW model, the Ghirlanda model, and others do an excellent job of tracking behaviour. The distinction sheds light on the very nature of retrospective revaluation as a learning phenomenon that is mediated by memory or, if De

Houwer (2009) is correct, symbolic reasoning.

Recent interest in the neurological plausibility of computational models demands a critical consideration of the Ghirlanda (2005) and Minerva-AL (Jamieson et al., 2012) models. Briefly, the evidence for both models is mixed. Path independence and catastrophic forgetting in single-layer artificial neural networks like the RW and Ghirlanda models pose obvious hurdles to neurological compatibility (French, 1999). Nevertheless, accumulating evidence that the nucleus accumbens is sensitive to prediction error suggests a possible biological mechanism for the error correction algorithm underlying classical associative learning models (e.g., Rodriguez, Aron, & Poldrack, 2006). Likewise, whereas the content-addressable nature of memory in the Minerva family of models casts doubt on their neurological plausibility, Hintzman (1990) showed that Minerva 2 could be recast as a multilayer artificial neural network with an associative memory framework. Given its foundations in Minerva 2, it is not unexpected that the Minerva-AL could be similarly recast. Moreover, the expectancy-encoding function of Minerva-AL allows it to capitalize on the aforementioned biological mechanism of prediction error in the nucleus accumbens.

General Discussion

Retrospective revaluation is the greatest challenge to classical theories of learning (e.g., Rescorla & Wagner, 1972). It challenges those theories because they suppose an organism only learns about information presented directly to it. Whereas theorists updated the classical models to account for learning about unpresented but implied stimuli (e.g., Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994), demonstrations of higher-order retrospective revaluation reissued the old challenge (De Houwer & Beckers, 2002a, b). However, the empirical reliability of revaluation learning remains contentious even at the first order of association (e.g., Batson & Batsell, 2000; Mitchell, Killedar, & Lovibond, 2005; Vadillo & Matute, 2010; Williams et al., 1994). The possibility of a file drawer effect suggests a level of discretion in our willingness to develop alternative theoretical accounts.

In Part 1, I examined retrospective revaluation at the first, second, and third orders of association. In Experiments 1A, 1B, and 1C, I found evidence for revaluation at all three orders in the tank task. These findings reinforced the work of De Houwer and Beckers (2002a, b) and uncovered a horizon on revaluation at higher orders, defined by an overall dilution of learning. In Experiments 2A, 2B, and 2C, analogous work with the allergist task failed to show compelling evidence for retrospective revaluation at any order of association and direct comparisons to successful revaluation in the tank task yielded a number of methodological considerations that may bear on the reliability of the phenomenon. Overall, the experimental work provides mixed empirical evidence for retrospective revaluation that is difficult to interpret in isolation.

In Part 2, I applied two computational models to the problem. Results from simulations with Minerva-AL, a modification of a classic model of human episodic memory (Hintzman, 1986), mapped on to participant behaviour in the tank task by successfully predicting first- and second-order revaluation (results from third-order simulations were less clear, although see below). Results from simulations with the Ghirlanda (2005) model, a post-neoclassical update to the Rescorla-Wagner (1972) model, mapped on to the particular patterns of failed revaluation behaviour observed in the allergist task. The historical origins and theoretical mechanisms of the two models shed light on the distinction. Namely, whereas a post-neoclassical model of learning does not accommodate retrospective revaluation, a model that derives learning as a corollary of memory does. I conclude along with others that retrospective revaluation forces a consideration of the role that memory plays in learning (Melchers, Lachnit, & Shanks, 2004).

Beyond the relative fit of the two models to the two sets of empirical data, the fact remains that neither model can explain the full array of results uncovered in this thesis. To my view, this state of affairs presents three possibilities. In the first possibility, the two tasks are both appropriate measures of retrospective revaluation, and a model that can explain both sets of results would be necessary to truly capture the phenomenon. However, such a model would require at least one (but likely more than one) additional free parameter to accomplish the divergent predictions, and would as a result be open to criticisms of over-fitting and unfalsibiablity (Sternberg & McClelland, 2012). In the second possibility, only one of the two tasks is an actual index of retrospective revaluation. On the basis of the behavioural data, it is arguable that the tank task is the more likely candidate. This possibility, however, bears the question: if the tank task is an index of retrospective revaluation, what is the allergist task measuring? There have been numerous successful demonstrations of retrospective revaluation in the allergist task, so a wholesale discarding of the task seems premature (e.g., Dickinson & Burke, 1996; Larkin, Aitkin, & Dickinson, 1998; Mutter, Atchley, & Plumlee, 2012; Wasserman & Berglan, 1998). In the third possibility, retrospective revaluation is not a window into the fundamental mechanisms of human learning but rather an anomalous blip in an otherwise sensible field, at best unpredictable and at worst illusory. The prevalence of complications in obtaining revaluation certainly hint at a file drawer effect, a possibility that is corroborated by personal communications with researchers in the field and by my own laboured attempts to bring the phenomenon under experimental control (e.g., Batson & Batsell, 2000; Mitchell, Killedar, & Lovibond, 2005; Vadillo & Matute, 2010; Williams, Sagness, & McPhee, 1994).

Of course none of these possibilities is particularly satisfying, as each leaves the question of retrospective revaluation in a state of uncertainty. At this point, researchers must begin to focus on the nuances and finer points of the phenomenon to develop an informed stance on the three possibilities. Until revaluation is brought under firm experimental control, no progress can be made on this front. Thankfully, computational modelling can inform and guide empirical work by providing candidate explanations for existing behaviours and predictions for further research.

Additional simulations with Minerva-AL, for example, have revealed insightful predictions regarding third-order revaluation. The protocol of Experiment 1C was informed by De Houwer and Beckers' (2002a) third-order design, in which participants learned *AB+* in a first phase, both *BC+* and *CD+* in a second phase, and either *A+* or *A-* in a third phase. Because *BC+* and *CD+* trials both occurred in the second phase of learning and in random order, the relationships between cue *A* and cues *C* and *D* were not under experimental control and were unclear with regard to revaluation. Whereas Minerva-AL did not accommodate revaluation under these contingencies, the model does predict third-order revaluation once *BC+* and *CD+* trials are separated into distinct learning phases, *BC+* and then *CD+*. It is not controversial to suggest that participants trained under such contingencies would show clearer patterns of higher-order revaluation than were observed in Experiment 1C of the present thesis. This testable prediction would moreover speak to the crucial role of memory in revaluative learning. Namely, in order learn retrospectively about unpresented but implied cues, those cues must first be stored in memory.

In fact the prediction draws a distinction between memory-based and reasoning-based explanations of the phenomenon. There is no reason, for example, that De Houwer's (2009) propositional model would anticipate different results from the two experimental protocols, a fact that others have recently appreciated as well (Sternberg & McClelland, 2012). Even more damaging, if participants are applying logic to symbolic representations of relationships between stimuli, then the order of *BC+* and *CD+* trials should have no impact on revaluation. Reasoning in a third-order revaluation task would proceed as follows: "if weapons *A* and *B* together, weapons *B* and *C* together, weapons *C* and *D* together, and weapon *A* alone all cause the tank to explode, *B* is unlikely to be an effective weapon. If *B* is not effective, then *C* must be. If *C* is effective, then it is unlikely that *D* is as well," and the behavioural outcome is identical whether *BC+* learning precedes *CD+* learning, or vice versa. By contrast, it has long been known that trial sequencing does affect human contingency learning (e.g., Collins & Shanks, 2002; Lopez et al., 1998). People do not implicitly apply pure Aristotelian logic to problems (Henle, 1962) and are notoriously susceptible to cognitive biases, framing effects, and failures of rationality (e.g., Kahneman & Tversky, 1973; Oberauer, 2006).

It is difficult to accept symbolic reasoning as a viable explanation for my participants' behaviours, not least of all because it implies a memory store from which to draw inferences without postulating a memory mechanism. Minerva-AL, on the other hand, describes both the behaviours and the mechanisms in a working process model. Whereas reasoning is an explicit and deliberate behaviour that by necessity assumes some foundation of memory, memory itself operates all the time and does not necessitate any additional processes to account for retrospective revaluation. Indeed, I argue that the appearance of symbolic reasoning in revaluation emerges from memory processes. The preference for basic mechanistic accounts over more complex ones is twofold. First, laws of parsimony suggest that explanations with fewer assumptions should be exhausted before more sophisticated explanations are considered. Second, retrospective revaluation is a phenomenon not limited to humans, but occurs even among species unlikely to possess Aristotelian reasoning capacities, such as honey bees (Giurfa & Bernard, 2006) and rats (i.e., Baker & Mercier, 1989; Miller & Matute, 1996; although see Beckers, Miller, De Houwer, & Urushihara, 2006). Whereas syllogistic reasoning might be a controversial idea in comparative cognition, the idea that animals both human and otherwise possess memory is not.

Of course, the evidence provided in this thesis does not force a single-system view. It is possible that there are two approaches to completing retrospective revaluation tasks: one that recruits memorial processes common across species, and one that recruits the sorts of syllogistic reasoning processes proposed by De Houwer (2009). In this analysis, nonhuman animals may rely on the former strategy, whereas human participants may be free to make flexible use of either. Although it is impossible to prove the nonexistence of a second system, a dilemma previously addressed in other areas of cognitive science (see for example Berry et al., 2010), the results of the current work are consistent with the predictions of a single-system model of memory.

The notion that complex cognitions like symbolic reasoning might be a corollary of simpler processes operating in concert is an old idea. Simon (1969) argued that sophisticated behaviour could emerge from unsophisticated systems when those systems operate within a complex environment. Accordingly, he warned against empiricists attributing apparently complex behaviours to complex mechanisms rather than to complex environments. In a compelling demonstration of the principle, Tero et al. (2010) recently showed that a slime mold (*Physarum polycephalum*) placed in a structured environment could develop networks that rival human engineered transport systems in measures of efficiency, fault resistance, and cost. Although it is tempting to attribute the behaviour to sophisticated mechanisms, a closer analysis reveals that the behaviour follows instead from the application of a simple foraging algorithm to a complex environment. Likewise, retrospective revaluation in the present thesis can be considered in light of emergent processes. When a system—whether a human participant or a computational model—is placed in a structured environment (see Tables 1 and 2), what appears at first glance to be complex behaviour may not in fact reflect complex processes. Indeed, insight into the mechanisms of Minerva-AL speaks to this fact: the model is decidedly simple, yet it accomplishes behaviours that researchers have taken as evidence of symbolic propositional reasoning abilities. I argue that retrospective revaluation is an illustration of *elegant* rather than *sophisticated* cognition.

The value of computational theory becomes clear in a consideration of emergent behaviours. Take for example Minerva-AL's account of higher-order retrospective revaluation. Whereas it is easy to track the model's behaviour at the first order (see Discussion of Computational Modelling), the exercise becomes more complicated with increasing orders of association. Part of the problem is that the interaction between encoding and retrieval in the model is a stochastic process, rendering a clean and stable description laborious after only three iterations (i.e., three phases of learning, as in second-order revaluation). Yet revaluation behaviour reliably emerges from that interaction at the level of a single simulated subject. Computational theory provides invaluable means for testing otherwise intractable theoretical mechanisms with the press of a button. Moreover, computational theory imparts an avenue for objective comparison between theoretical stances. Whereas verbal theorists must argue the semantics of associative learning mechanisms, comparisons are finessed and resolved cleanly by mathematics. In this thesis I compared predictions from two contrasting theoretical accounts against behavioural data—a precision of analysis that would not be possible independent of mathematical models.

There are nevertheless limitations to my analysis that must be acknowledged. Namely, the two models selected for comparison were a purposeful and explicit decision to contrast a theory of learning (Ghirlanda, 2005; Ghirlanda & Enquist, 1999, 2007) against a theory of memory (Jamieson et al., 2010, 2012). The purpose of course was to evaluate whether memory is a necessary construct to explain retrospective revaluation. There are, however, a number of other models that have not been considered in this thesis but that may have served the same purpose. Miller's Sometimes-Competing Retrieval model (SOCR; Stout & Miller, 2007; see also Miller & Matzel, 1988) is an alternative memory-based model that is known to account for retrospective revaluation at both the first order and higher orders (see Discussion of Experiments 1A, 1B, and 1C). But the use of Minerva-AL reflects ongoing theoretical work in my laboratory, and future work comparing the two models would be beneficial. In addition to the Ghirlanda (2005) model described in this thesis, Ghirlanda and Enquist (1998) have used a multilayer artificial neural network with backpropagation algorithms to simulate a variety of associative learning effects. Likewise, Delamater (2012) recently proposed a multilayer connectionist network model and illustrated its proficiency with a number of classical Pavlovian learning phenomena. It is unclear whether either of these multilayer network models could account for revaluation, but the question should certainly be explored in light of the failures of the post-neoclassical account assessed in this thesis.

The behavioural and computational analyses of this thesis provide a number of new avenues for additional empirical work. The first avenue involves a variation on the design of the third-order tank task that separates learning into four phases (i.e., *AB+, BC+, CD+, A+ or A-*) rather than three (*i.e., AB+, BC+/CD+, A+ or A-*). The predictions of Minerva-AL on this variation present a unique opportunity for comparison between memory- and reasoning-based theories of revaluation. When the design of Experiment 1C and De Houwer and Beckers' (2002a) analogous experiment is considered critically, it is a wonder that some participants in both experiments showed evidence of third-order revaluation. Nevertheless, evidence for third-order revaluation in my Experiment 1C was incomplete. This allows for the possibility that separating the third-order task into four phases of learning will facilitate both retrospective revaluation and a cleaner fit between behavioural data and model predictions. The second avenue explores one of the major differences between the tank task and allergist task that may have contributed to the divergent sets of results (see Discussion of Experiments 2A, 2B, and 2C). Namely, the instructions of the tank task and the subsequent presentation of "IMPACT 10/20" feedback messages during training encouraged an elemental consideration of compound weapons. By contrast, the instructions in the allergist task, by omission, allowed for a configural interpretation of compound foods. The difference between elemental and configural processing in both human (e.g., De Houwer, Beckers, & Glautier, 2002; Williams et al., 1994) and nonhuman learning (e.g., Pearce, 1987; Rescorla, Grau, & Durlach, 1985) has long been appreciated. The possibility that elemental interpretations of compound cues are necessary for revaluation learning presents an interesting empirical question. Namely, if "IMPACT 10/20" feedback and accompanying instructions are removed from the tank task, will participants cease to revaluate cue *B*? Analogously, if "REACTION STRENGTH 10/20" feedback and accompanying instructions are added to the allergist task, will participants begin to revaluate cue *B*? This "Mad Hatter" hypothesis proposes that the results of Experiments 1A and 2A may be interchangeable around an axis of configural versus elemental processing. If so, it will provide compelling insight into the mechanisms that support retrospective revaluation. In this style, even the failures of revaluation in the allergist task provide meaningful opportunities for a deeper understanding of revaluative learning in general, and the interactions of human learning and memory more broadly.

Beyond its value to theoretical and computational considerations of retrospective revaluation, the current work speaks more generally to philosophical concerns of human mind and cognition. Indeed, the emergence of seemingly complex behaviours from simple memory mechanisms need not be limited to revaluative learning behaviours. I propose that the encoding and retrieval of memories constitutes the basic building blocks of cognition and that behaviours as apparently complex as consciousness are emergent properties of these mechanisms operating in parallel. Though not explicitly informed by the evidence provided in this thesis, these possibilities are suggested by the ideas contained within and lay the groundwork for a broader theory of human cognition.

References

Baker, A. G. & Mercier, P. (1989). Attention, retrospective processing and cognitive representations. In S. B. Klein & R. R. Mowrer (Eds.). *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 85-101). Hillsdale, NJ: Lawrence Erlbaum.

Batson, J. D. & Batsell., W. R. Jr. (2000). Augmentation, not blocking, in a A+/AX+ flavor-conditioning procedure. *Psychonomic Bulletin & Review, 7*, 466-471.

Beckers, T., Miller, R. R., De Houwer, J., & Urushihara, K. (2006). Reasoning rats: Forward blocking in Pavlovian animal conditioning is sensitive to constraints of causal inference. *Journal of Experimental Psychology: General, 131*, 92-102.

Berry, C. J., Shanks, D. R., Li, S., Rains, L. S., & Henson, R. N. A. (2010). Can "pure" implicit memory be isolated? A test of a single-system model of recognition and repetition priming. *Canadian Journal of Experimental Psychology, 64*, 241-255.

Collins, D. J. & Shanks, D. R. (2002). Momentary and integrative response strategies in causal judgment. *Memory & Cognition, 30*, 1138-1147.

Daw, N. & Courville, A. (2008). The Pigeon as Particle Filter. *Advances in Neural Information Processing Systems, 20*, 1-8.

Dawson, M. R. W. (2008). Connectionism and classical conditioning. *Comparative Cognition and Behavior Reviews, 3* (Monograph), 1–115.

Delamater, A. R. (2012). On the nature of CS and US representations in Pavlovian learning. *Learning & Behavior, 40*, 1-23.

De Houwer, J. (2009). The propositional approach to associative learning as an alternative for association formation models. *Learning & Behavior, 37*, 1-20.

De Houwer, J., & Beckers, T. (2002a). Higher-order retrospective revaluation in human causal learning. *Quarterly Journal of Experimental Psychology, 55B,* 137-151.

De Houwer, J., & Beckers, T. (2002b). Second-order backward blocking and unovershadowing in human causal learning. *Journal of Experimental Psychology, 49,* 27-33.

De Houwer, J., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology, 55A,* 965-985.

Dickinson, A. & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgements. *Quarterly Journal of Experimental Psychology, 49B*, 60-80.

Dickinson, A., Shanks, D. R., & Evenden, J. L. (1984). Judgement of act-outcome contingency: the role of selective attribution. *Quarterly Journal of Experimental Psychology, 36A*, 29-50.

French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Science, 3*, 128-135.

Ghirlanda, S. (2005). Retrospective revaluation as simple associative learning. *Journal of Experimental Psychology: Animal Behavior Processes, 31*, 107-111.

Ghirlanda, S. & Enquist, M. (1998). Artificial networks as models of stimulus control. *Animal Behaviour, 56*, 1383-1389.

Ghirlanda, S. & Enquist, M. (1999). The geometry of stimulus control. *Animal Behaviour, 58*, 695-706.

Ghirlanda, S. & Enquist, M. (2007). Electronic supplementary material. How training and testing histories affect generalization: A test of simple neural networks. *Philosophical Transactions of the Royal Society B362*, 449–454.

Giurfa, M. & Bernard, J. (2006). Prospective and retrospective learning in honeybees. *International Journal of Comparative Psychology, 19*, 358-367.

Haselgrove, M. (2010). Reasoning rats or associative animals? A common-element analysis of the effects of additive and subadditive pretraining on blocking. *Journal of Experimental Psychology: Animal Behavior Processes, 36*, 296-306.

Henle, M. (1962). On the relation between logic and thinking. *Psychological Review, 69*, 366-378.

Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behaviour Research Methods, Instruments, & Computers*, *16*, 96-101.

Hintzman, D. L. (1990). Human learning and memory: Connections and dissociations. *Annual Review of Psychology, 41*, 109-139.

Jamieson, R. K., Hannah, S. D., & Crump, M. J. C. (2010). A memory based account of retrospective revaluation. *Canadian Journal of Experimental Psychology, 64*, 153-164.

Jamieson, R. K., Crump, M. J. C., & Hannah, S. D. (2012). An instance theory of associative learning. *Learning & Behavior, 40*, 61-82.

Jamieson, R. K., & Mewhort, D. J. K. (2010). Applying an exemplar model to the artificial-grammar task: String-completion and performance on individual items. *Quarterly Journal of Experimental Psychology, 63*, 1014-1039.

Kahneman, D. & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*, 237-251.

Kamin, L. J. (1969). Predictability, surprise, attention and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279-296). New York: Appleton-Century-Crofts.

Kruschke, J. K. (2006). Locally Bayesian learning with applications to retrospective revaluation and highlighting. *Psychological Review, 113*, 677-699.

Larkin, M. J. W., Aitken, M. R. F., & Dickinson, A. (1998). Retrospective revaluation of causal judgements under positive and negative contingencies. *Journal of Experimental Psychology: Learning and Memory, 26*, 1331-1352.

Le Pelley, M. E., Cutler, D. L., & McLaren, I. P. L. (2000). *Retrospective effects in human causality judgment*. Poster presented at the 22nd annual meeting of the Cognitive Science Society.

Le Pelley, M. E., Oakeshott, S. M. & McLaren, I. P. L. (2005). Blocking and unblocking in human causal learning. *Journal of Experimental Psychology: Animal Behavior Processes, 31*, 56-70.

Lopez, F. J., Shanks, D. R., Almaraz, J, & Fernandez, P. (1998). Effects of trial order on contingency judgments: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 24*, 672-694.

Melchers, K.G., Lachnit, H., & Shanks, D. R. (2004). Within-compound associations in retrospective revaluation and in direct learning: A challenge for comparator theory. *Quarterly Journal of Experimental Psychology, 56B*, 25-53.

Miller, R. R. (2006). Challenges facing contemporary associative approaches to acquired behavior. *Comparative Cognition & Behavior Reviews, 1*, 77-93.

Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the Rescorla-Wagner model. *Psychological Bulletin, 117*, 363-386.

Miller, R. R. & Matute, H. (1996). Biological significance in forward and backward blocking: Resolution of a discrepancy between animal conditioning and human causal judgement. *Journal of Experimental Psychology: General, 125*, 370-386.

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 51-92). San Diego, CA: Academic Press.

Miller, R. R., & Schachtman, T. R. (1985). Conditioning context as an associative baseline: Implications for response generation and the nature of conditioned inhibition. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 51 - 88). Hillsdale, NJ: Erlbaum.

Mitchell, C. J., Killedar, A., & Lovibond, P. F. (2005). Inference-based retrospective revaluation in human causal judgments requires knowledge of within-compound associations. *Journal of Experimental Psychology: Animal Behaviour Processes, 31*, 418-424.

Mutter, S. A., Atchley, A. R., & Plumlee, L. M. (2012). Aging and retrospective revaluation of causal learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *38,* 102-117.

Oberauer, K. (2006). Reasoning with conditionals: A test of formal models of four theories. *Cognitive Psychology, 53*, 238-283.

Pavlov, I. P. (1927). *Conditioned Reflexes*. London: Oxford University Press.

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

Rescorla, R. A., Grau, J. W., & Durlach, P. J. (1985). Analysis of the unique cue in configural discriminations. *Journal of Experimental Psychology: Animal Behavior Processes, 11*, 356-366.

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

Rodriguez, P. F., Aron, A. R., & Poldrack, R. A. (2006). Ventral-striatal/nucleus-accumbens sensitivity to prediction errors during classification learning. *Human Brain Mapping, 27*, 306-313.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *Quarterly Journal of Experimental Psychology, 37B*, 1-21.

Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.

Sternberg, D. A. & McClelland, J. L. (2012). Two mechanisms of human contingency learning. *Psychological Science, 23*, 59-68.

Stout, S. C. & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review, 114*, 759-783.

Tangen, J. M & Allan, L. G. (2004). Cue interaction and judgments of causality: contributions of causal and associative processes. *Memory & Cognition, 32*, 107-124.

Tero, A., Takagi, S., Saigusa, T., Ito, K., Bebber, D. P., Fricker, M. D., et al. (2010). *Science, 327*, 439-442.

Vadillo, M. A. & Matute, H. (2010). Augmentation in contingency learning under time pressure. *British Journal of Psychology, 101*, 579-589.

Van Hamme, L. J. & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and Motivation, 25*, 127-151.

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

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

Wasserman, E. A. & Berglan, L. R. (1998). Backward blocking and recovery from overshadowing in human causal judgement: The role of within-compound associations. *Quarterly Journal of Experimental Psychology, 51B*, 121-138.

Williams, D. A., Sagness, K. E., & McPhee, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 20*, 694-709.

Witnauer, J. E. & Miller, R. R. (2011). The role of within-compound associations in learning about absent cues. *Learning & Behavior, 39*, 146-162.