Simultaneous Backward Conditioned Inhibition and Mediated Conditioning

Steven Graham[1], Hans Lee Jie[1], Hui-Minn Chan[1] I.P.L. McLaren[2] & Andy J. Wills[2]

[1]Department of Psychology, National University of Singapore, Singapore.

[2]School of Psychology, Exeter University, U.K.

Correspondence should be addressed to SG:

Department of Psychology,

National University of Singapore,

Singapore 117570

Telephone: +65 6516-5000

Facsimile: +65 6773-1843

Email: psysg@nus.edu.sg

Acknowledgments
This research was supported by a British Council PMI2 Connect Research Co-operation Award (RC105) to AJW and SG.

Abstract

Demonstrations of retrospective revaluation suggest that remembered stimuli undergo a reduction in association with the unconditioned stimulus present during learning. Conversely, demonstrations of mediated conditioning in flavor conditioning experiments with rats suggest that remembered stimuli undergo an increase in association with the US present during learning. In a food allergy prediction task with 23 undergraduates we demonstrated simultaneous backward conditioned inhibition and mediated conditioning effects. These results are compatible with the hypothesis that the direction of change (decrease or increase) in associative strength depends on whether the remembered stimulus was of a different category (CS / antecedent) or the same category (US / outcome) as the presented US.

Associative learning theories advance a simple concept – that stimulus representations can become linked (or associated) so that activation of one can affect those associated with it. Accordingly, changes in these associations are the basis of learning in associative theories; with the exact rules governing these associative changes differing from one theory to another. One such influential theory is the 'sometimes opponent processing' (SOP) model (Wagner, 1981). In SOP, stimulus representations are conceived as a set of elements which could be activated either directly (when the stimulus is presented) into a short-term representational store of limited capacity (the 'A1' state) or indirectly retrieved into the 'A2' state via a previously formed association with a currently presented stimulus. In addition, A1 representations would quickly decay into 'A2' (of greater capacity than A1) before finally becoming inactive. Wagner (1981) proposed specific learning rules to govern the formation of associations between activated stimulus representations: two stimuli in A1 form excitatory associations; a stimulus in A1 forms an inhibitory association to a stimulus in A2.

The SOP model predicts many learning phenomena (see e.g. Wagner, 1981); however a series of flavor-conditioning experiments in rats suggested the need for extension of SOP. Holland (1981) showed that when a tone was presented in serial compound with a flavored food and subsequently presented in serial compound with LiCl (a substance which induces nausea) injection, rats acquired an aversion to the food. This result became known as mediated conditioning. If, following tone-LiCl pairings, the tone is presented alone, then the rats' aversion to the food reduces (Holland & Forbes, 1982). To account for these findings, Holland (1983) extended SOP's learning rules such that indirectly activated A2-representations would form

excitatory associations to A1-representations, and that if both stimuli were in A2-states, an inhibitory association would develop between them.

Holland's SOP extensions were, however, incompatible with studies in which human participants re-evaluated causality judgments about an absent stimulus in light of new information about other previously-associated stimuli. For example, Dickinson and Burke (1996) used a scenario in which humans played the role of an allergist asked to judge the likelihood that foods (cues) produced an allergic reaction (outcome). In Phase 1, subjects learned that compound foods AB and CD caused an allergic reaction. In Phase 2, subjects learned that Food A alone resulted in the allergic reaction whereas Food C alone did not (i.e. AB+A+; CD+C-). In the final test phase subjects rated Food B as being less likely to cause the allergic reaction than Food D. Dickinson & Burke (1996) then went on to show that within-compound food associations were responsible for this so-called 'retrospective revaluation' effect and proposed alternative extensions to Wagner's (1981) SOP model: namely, that inhibitory associations would form between an A1-activated representation and an A2-activated representation; whereas excitatory associations would form between simultaneously activated A2 representations.

The Holland (1983) and Dickinson and Burke (1996) extensions to SOP make opposing predictions. On the one hand, Holland (1983) proposed the formation of excitatory associations between A1 and A2 stimulus representations, and the formation of inhibitory associations between A2 and A2 stimulus representations. On the other hand, Dickinson and Burke (1996) proposed the formation of inhibitory associations between A1 and A2 stimulus representations, and the formation of excitatory associations between two stimulus representations in the same state (both in A1 or both in A2). The conundrum is that evidence exists to support each of these

proposed modifications. Importantly though, mediated conditioning typically involves associative increases between an evoked US representation and the presented US, whereas decreases are typically observed between an evoked CS representation and the presented US. The difference in results might therefore be attributable to the difference in the category (CS vs. US) or timing (antecedent vs. effect) of the evoked stimulus representation and the presented US (Graham, 1999; Dwyer 2003). The aim of the present experiment was therefore to demonstrate associative decreases (analogous to backward conditioned inhibition, see Table 1; Le Pelley, Cutler & McLaren, 2000) for antecedent CS-type representations and associative increases (akin to mediated conditioning) for US-type effect representations. The aim was also to demonstrate that increases *and* decreases can occur simultaneously, and in a way that could not be explained by appealing to cross-experiment changes in stimulus discriminability (see Balleine *et al.*, 2005; Liljeholm & Balleine 2009).

Methods

Subjects and apparatus. A total of 23 undergraduate psychology students (12 females) between the ages of 21 and 25 participated for course credits. The study was conducted in a quiet testing room using PCs running Presentation® software (Neurobehavioral Systems). The experimental protocol was approved by the local ethics committee and informed consent was obtained prior to experimentation.

Design and Procedure. The participants played the role of a food allergist and were asked to learn which particular allergic reactions a hypothetical person "Mr. X" would develop after eating certain foods. Participants sat about 0.5m away from the

computer screen and were shown white word stimuli (of height 5cm) centered on a black screen. The word stimuli depicting foods (bananas, beef, chicken, coconuts, cheese, lamb, fish, peanuts, shrimp, and oysters) and six allergic reactions (cough, rashes, stomach ache, fever, nausea, and headache) were randomly assigned to A-J and allergic reactions 1-6 in the design (see Table 1). Throughout the experiment, a row of seven white outlined boxes, each containing labels for the six allergic reactions and "No Allergy", were displayed and highlighted whenever their corresponding response key (F1 to F7 on the keyboard) was pressed. Each trial began with the presentation of the food words (displayed one above the other with position counterbalanced across trials) for 2 seconds during which time a response was required. Feedback ("correct", "wrong", or "please respond faster") together with the correct allergic reaction were then presented for 1 second.

Every participant was trained concurrently on five different conditions: Backward Blocking AB+1 A+1, Unovershadowing CD+2 C-, Retrospective Revaluation Control EF+3 E˚, Retrospective Revaluation & Mediated Conditioning GH+4 G+5, and Mediated Conditioning Control IJ˚ I+6 (see Table 1). In Phase 1, participants were shown food pairs (AB, CD, EF and GH) and asked to predict which allergic reaction Mr. X would develop (the correct answers were Allergic Reactions 1, 2, 3 and 4 for AB, CD, EF and GH, respectively). Participants were told they should guess initially, but through trial and error could eventually learn to predict the correct outcomes. An additional food pair, IJ, was not followed by any outcome– the screen simply went blank until the next trial appeared. Before the experiment, participants were told that some trials might terminate prematurely, but the absence of feedback about an allergic reaction did not imply no allergic reaction (rather it could not be determined whether any allergic reaction had developed or not) and judgment about

the effects of these food stimuli should be reserved until a later time. Each of the five compounds was presented once, in a random order, in each block of 5 trials. Phase 1 comprised 250 trials in total; 50 trials for each condition.

At the end of Phase 1, participants received the following instructions, on a sheet of paper:

"Based on what you have learned in this experiment so far, please indicate whether certain foods cause certain allergic reactions in Mr. X. Use a scale from +3 and -3 where:

- positive numbers indicate the food CAUSES the allergy
- leave it blank if you think there is no relationship between the food and allergy
- negative numbers indicate the food PREVENTS the allergy from occurring

Use the scale to reflect the confidence in your answers: more positive or negative numbers (i.e., further away from zero) indicate more certainty of your answer (e.g., +/- 3 indicates that you are very sure of the relationship, +/- 1 indicates you are only somewhat sure). Remember to leave it blank if you think the food and the allergy are not related or if you are unsure about the relationship (blanks will be treated as zeros on the scale). The data analysis for this experiment is based on your answers on this sheet, so PLEASE take your time, and answer carefully and thoughtfully"

On the same sheet of paper, below these instructions, appeared a grid with a different column for each allergy (plus a "No Allergy" column), and a different row for each food. Hence, each participant made up to 70 ratings at the end of Phase 1, self-paced, and in a self-determined order. The sheet was then collected, and the participant began Phase 2 on the computer.

During Phase 2, participants were shown single foods (A, C, G and I) and asked to predict which allergic reaction Mr. X would develop (the correct answers were Allergic Reaction 1, No Allergic Reaction, Allergic Reaction 5, and Allergic Reaction 6 for A, C, G and I, respectively). For an additional food, E, the screen simply went blank until the next trial appeared. Each of the five foods was presented once, in a random order, in each block of 5 trials. Phase 2, like Phase 1, comprised 250 trials in total; 50 trials for each condition.

At the end of Phase 2, participants received rating instructions identical to those presented at the end of Phase 1. This food-allergy rating sheet was then removed, and replaced by the following instructions:

"Based on what you have learned, please indicate whether allergic reactions below are related to each other. Use a scale from +3 and -3 where:

- positive numbers indicate one allergy CAUSES the other
- leave it blank if you think there is no relationship between the allergies
- negative numbers indicate the allergy PREVENTS the other allergy from occurring.

Use the scale to reflect the confidence in your answers: more positive or negative numbers (i.e., further away from zero) indicate more certainty of your answer (e.g., +/- 3 indicates that you are very sure of the relationship, +/- 1 indicates you are only somewhat sure). Remember to leave it blank if you think the allergies are not related or if you are unsure about the relationship (blanks will be treated as zeros on the scale). The data analysis for this experiment is based on your answers on this sheet, so PLEASE take your time, and answer carefully and thoughtfully".

On the same sheet of paper, below these instructions, appeared a grid with each of the six allergies (plus "No Allergy") as a different column, and each of the six allergies (plus "No Allergy") as a different row. Cells relating an allergy to itself (e.g. nausea-nausea) had "N/A" pre-entered. To the left of the grid was written, "This allergy…", and to the top of this grid was written "…causes this allergy".

Hence, each participant made up to 42 allergy-allergy ratings, in a self-paced, and a self-determined order.

- Table 1 about here –

In this experiment, E° trials were intended as a control, such that reduced ratings for B causing Allergic Reaction 1 compared to ratings for F causing Allergic Reaction 3 would demonstrate backward blocking and increased ratings for D causing Allergic Reaction 2 relative to ratings for F causing Allergic Reaction 3 would demonstrate unovershadowing. Of particular interest however were the trials involving G and Allergic Reaction 5. Food compound GH in Phase 1 was followed by Allergic Reaction 4, whereas G alone in Phase 2 was followed by Allergic Reaction 5. Thus "G+5" trials in Phase 2 should evoke representations of both Food H and Allergic Reaction 4 in the presence of Food G and Allergic Reaction 5. It was predicted that ratings for H causing Allergic Reaction 4 would increase (an effect analogous to unovershadowing) and ratings for H causing Allergic Reaction 5 would decrease (an effect analogous to backward conditioned inhibition) along with increases in the association between Allergic Reaction 4 and Allergic Reaction 5 (an effect analogous to mediated conditioning). That is, we expected simultaneous effects akin to retrospective revaluation and mediated conditioning during G+5 trials.

Results

Terminal accuracy, assessed over the last 10 trials of each condition, was high in both phases (see Table 1). At the end of Phase 1, the mean ratings for foods A-H for their correct outcomes were 2.39, 2.17, 2.26, 2.17, 2.30, 2.22, 2.52, 2.39, respectively (maximum rating was 3). The averaged ratings for the incorrect outcomes for foods A-H were 0.00, 0.04, 0.06, 0.02, -0.01, 0.01, 0.02, -0.02, respectively. I and J had no correct outcome; the average rating across all outcomes were 0.08, and 0.09 respectively. The mean of each of the 70 food-outcome ratings at the end of Phase 1, and the 70 food-outcome ratings and the 42 outcome-outcome ratings at the end of Phase 2, are presented in the Supplementary Materials.

- Figure 1 about here -

Figure 1 summarizes the participants' mean ratings of interest at the end of Phase 2. We found significant backward blocking - ratings for B causing Allergic Reaction 1 (B1) were significantly lower than for F causing Allergic Reaction 3 (F3), $F(1,22) = 8.72$, $p<0.01$. We also found significant unovershadowing (D2 > F3), $F(1,22) = 10.14$, $p<0.01$. Combined these ratings indicated significant retrospective revaluation (B1 < D2), $F(1,22) = 24.14$, $p<0.001$. In addition, we found backward conditioned inhibition (H5 < H6), $F(1,22) = 5.29$, $p<0.05$, and an extinction-like effect (G4 < E3), $F(1,22) = 5.15$, $p<0.05$. Another possible test of backward condition inhibition, J5 vs. J6, was not significant, $F(1,22) < 1$. Le Pelley & McLaren (2001) also failed to demonstrate backward conditioned inhibition with a design of this type (i.e., one where Phase 1 involved non-reinforced pre-exposure of the compound). One

possible explanation of this absence of a difference is that the salience of J reduces during Phase 1, as a result of its non-reinforced pre-exposure, though Le Pelley and Mclaren favour the idea that only an existing inhibitory association could be revalued to produce backward conditioned inhibition, and pre-exposure (IJ followed by a blank screen) prevents the compound becoming slightly inhibitory as would be the case if the compound was simply presented with no outcome (IJ-).

Ratings for H causing Allergic Reaction 4 were numerically higher than F causing Allergic Reaction 3 (i.e., an effect similar to unovershadowing) but this difference failed to reach significance ($F(1,22) = 2.68$, $p=0.12$). We note that the ratings for H4 and D2 are numerically similar, (and do not differ significantly, $F(1,22) < 1$), indicating that the numerical size of the unovershadowing effect in the Unovershadowing and RR & MC conditions are similar. We also note that the rating for F3 is numerically higher than the rating for E3, suggesting that F may have undergone some overshadowing with respect to Allergic Reaction 3; making overshadowing of H with respect to Allergic Reaction 4 harder to detect against an F3 control. However, the F3-E3 difference is not statistically significant, $F(1,22) < 1$. An alternative test of unovershadowing in the Retrospective Revaluation and Mediated Conditioning condition is H4 vs. E3, which is statistically significant, $F(1,22) = 4.56$, $p < 0.05$. However, this comparison is likely to over-estimate the size of any unovershadowing effect of H with respect to Allergic Reaction 4, as E3 may have undergone some extinction as a result of E˚ trials.

Ratings for allergy-allergy associations at the end of Phase 2 were also examined. Ratings for the likelihood that Allergic Reaction 4 predicted Allergic Reaction 5 were significantly higher than those for Allergic Reaction 4 predicting an unrelated outcome of similar novelty and treatment during Phase 2 (i.e., Allergic

Reaction 6: $F(1,22) = 4.43$, $p<0.05$). We shall return to this later, but this effect seems analogous to mediated conditioning. Ratings for Allergic Reaction 5 predicting Allergic Reaction 4 were also significantly higher than for Allergic Reaction 6 predicting Allergic Reaction 4 ($F(1,22) = 7.91$, $p<0.01$). The direction of the association between Allergies 4 and 5 appeared to be stronger for Allergic Reaction 5 predicting Allergic Reaction 4, but failed to reach significance ($F(1,22) = 1.13$).

Discussion

The present study demonstrated simultaneous excitatory (mediated conditioning) and inhibitory (backward conditioned inhibition) associative changes between presented and evoked stimuli in the same experiment, in fact resulting from the same trial. It clearly makes the point that both types of effect are possible, and may enable us to shed some light on how and under what circumstances they may occur. Although retrospective revaluation type effects were typically found in humans and elusive in infra-humans (see Balleine *et al.*, 2005) whereas mediated conditioning was typically demonstrated in rodents (Holland, 1981, 1983), this study adds further support (see also Balleine et al., 2005; Liljeholm & Balleine 2009) to the idea that explanations in terms of species differences are unlikely to reconcile the opposing SOP modifications proposed by Holland (1983) and Dickinson & Burke (1996).

Balleine et al. (2005) demonstrated that retrospective revaluation (i.e., training of AB+ and CD+ followed by A+ and C-, leading to B < D at test) depends upon the cues of the compounds being relatively discriminable (e.g., A being discriminable from B). Where the cues of the compounds are highly confusable, the opposite result is found (i.e., B > D). Although this anti-retrospective revaluation result is consistent

with Holland's modification to SOP, an alternative explanation, which Balleine et al. (2005) offered, is that the conditioning to the single cue in Phase 2 (e.g. A) generalizes to the other cue in the Phase 1 compound via common representational elements at test.

A generalization-based account does not explain the mediated conditioning result reported in the current paper. Specifically, whilst one could posit that Allergic Reaction 5, presented in Phase 2, shares some common elements with Allergic Reaction 4, there is no reason to suppose this is any different for Allergic Reaction 6 (allocation of allergies to outcomes was randomized in the current study). Therefore a generalization-based account doesn't explain why the 4-5 rating is higher than the 4-6 rating in the current experiment.

In the present study, simultaneous backward conditioned inhibition and mediated conditioning arising from the same trial (G+5) require a different explanation. Specifically, GH and Allergic Reaction 4 (GH+4) pairings followed by G and Allergic Reaction 5 (G+5) pairings decreased the ratings for H causing Allergic Reaction 5 (backward conditioned inhibition) but also increased the ratings that Allergic Reaction 4 caused Allergic Reaction 5 (mediated conditioning). Neither Dickinson and Burke's, nor Holland's, modifications to SOP can, without further modification, explain both the backward conditioned inhibition and mediated conditioning results found in the current study. Within an SOP framework, the Dickinson and Burke modification is required to explain backward conditioned inhibition, whilst the Holland modification is required to explain the mediated conditioning effects observed here.

One way forward within an associative framework might be to speculate that the difference between the nature of the associatively evoked stimulus representation (CS

or US) or its timing (antecedent cue vs. subsequent outcome) might determine whether it decreases or increases its association with the presented outcome (see also Graham 1999; Dwyer 2003). Notably, mediated conditioning in our study involves the evoked representation of a US (US in A2 state) whereas backward conditioned inhibition involves the evoked representation of a CS (CS in A2 state). An account of the results presented in this paper might therefore be sought in terms of whether the associatively remembered (but absent) stimulus could be considered as a member of the same category be that physical, temporal or logical (e.g., food vs. allergy, cue vs. outcome, CS vs. US) as the presented stimulus with which it was being re-valued. In the present study, G+5 trials would have allowed Food H (a 'CS' in A2) to be associatively activated in the presence of Allergic Reaction 5 (a 'US' in A1), resulting in a decrease in their association (the direction of associative change analogous to backward conditioned inhibition). At the same time, G would have associatively activated Allergic Reaction 4 (a 'US' in A2) in the presence of Allergic Reaction 5 (a 'US' in A1), resulting in an increase in their association (an associative change analogous to mediated conditioning). Thus, decreases in association were observed when a CS was associatively activated in the presence of a US (different category). Conversely, increases in association were observed when a US was associatively activated in the presence of another US (same category).

An alternative possibility is that our results are better explained outside an associative framework (and, indeed, some other results in human contingency learning seem resistant to any simple associative explanation, e.g. De Houwer, 2002; Mitchell, Livesey & Lovibond, 2007). One way of distinguishing associative and inferential accounts of the current phenomena might be to manipulate the causal scenario employed. It seems reasonable to suggest that, in the allergy prediction task

we employed, people assume that foods cause allergies but that allergies do not cause other allergies. It has previously been demonstrated that the causal scenario employed can affect the results found in cue competition experiments (e.g. Waldmann & Holyoak, 1992), although some phenomena seem relatively insensitive to the choice of causal scenario (Arcediano, Matute, Escobar & Miller, 2005). It is an open question whether the current results would generalize to a procedure employing abstract stimuli in a causally neutral scenario. The modified-SOP account offered in the current paper predicts no qualitative change in the results as a consequence of the change of causal scenario.

A related point is that the current experiment placed participants under a moderate degree of time pressure (2 seconds per training trial; the ratings were self-paced). Previous work in categorization and contingency learning indicates that time pressure can sometimes qualitatively change peoples' performance compared to a self-paced condition (e.g. Karazinov & Boakes, 2007; Milton, Longmore & Wills, 2008). However, other phenomena, for example forward cue competition, have been demonstrated under both time pressured (Wills, Lavric, Croft & Hodgson, 2007) and self-paced (e.g. Dickinson & Burke, 1996) conditions. It nevertheless remains an interesting question whether the results reported in the current paper would also be found under self-paced conditions.

In summary, the results of the present study are consistent with the idea that, when a US-like representation is associatively activated, increases in association can be expected with any US-like stimulus present on that trial. Conversely, when a CS-like representation is associatively activated, decreases in association can be expected with any US-like stimulus present on that trial. What constitutes a US or CS should depend on the motivational significance of that stimulus, but perhaps in the case of

humans it might be interpreted more loosely as whether a stimulus is perceived to be a cause or an effect, a food or an allergic reaction, or even a member of the same or a different category.

References

Arcediano, F., Matute, H., Escobar, M., & Miller, R. R. (2005). Competition between antecedent and between subsequent stimuli in causal judgments. *Journal of Experimental Psychology-Learning Memory and Cognition, 31*(2), 228-237.

Balleine, B. W., Espinet, A., & Gonzalez, F. (2005). Perceptual learning enhances retrospective revaluation of conditioned flavor preferences in rats. *Journal of Experimental Psychology: Animal Behavior Processes, 31*(3), 341-350.

De Houwer, J., & Beckers, T. (2002). Higher-order retrospective revaluation in human causal learning. *Quarterly Journal of Experimental Psychology, 55B*(2), 137-151.

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

Dwyer, D. M. (2003). Learning about cues in their absence: Evidence from flavour preferences and aversions. *Quarterly Journal of Experimental Psychology B, 56*, 56–67.

Graham, S. (1999). Retrospective revaluation and inhibitory associations: does perceptual learning modulate our perception of the contingencies between events? *Quarterly Journal of Experimental Psychology B, 52*(2), 159-185.

Holland, P.C. (1981). Acquisition of representation-mediated conditioned food aversions. *Learning and Motivation*, 12, 1-18.

Holland, P.C. (1983). Representation-mediated overshadowing and potentiation of conditioned aversions. *Journal of Experimental Psychology: Animal Behavior Processes*, 9, 1-13.

Holland, P.C., & Forbes, D.T. (1982). Representation-mediated extinction of conditioned flavor aversions. *Learning and Motivation*, **13**, 454-471.

Karazinov, D. M., & Boakes, R. A. (2007). Second-order conditioning in human predictive judgements when there is little time to think. *Quarterly Journal of Experimental Psychology, 60*(3), 448-460.

Le Pelley, M.E., Cutler, D.L. & McLaren, I.P.L. (2000). Retrospective effects in human causality judgment. In L.R. Gleitman & A.K. Joshi (Eds.). *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society* (pp.782-787). Hillsdale, NJ: Lawrence Erlbaum Associates.

Le Pelley, M. E., & McLaren, I. P. L. (2001). Retrospective revaluation in humans: Learning or memory? *Quarterly Journal of Experimental Psychology, 54B*(4), 311-352.

Liljeholm, M., & Balleine, B. W. (2009). Mediated conditioning versus retrospective revaluation in humans: the influence of physical and functional similarity of cues. *Quarterly Journal of Experimental Psychology, 62*(3), 470-482.

Milton, F., Longmore, C. A., & Wills, A. J. (2008). Processes of overall similarity sorting in free classification. *Journal of Experimental Psychology: Human Perception and Performance, 34*(3), 676-692.

Mitchell, C. J., LIvesey, E., & Lovibond, P. F. (2007). A dissociation between causal judgement and the ease with which a cause is categorized with its effect. *Quarterly Journal of Experimental Psychology, 60*(3), 400-417.

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

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and Diagnostic Learning within Causal-Models - Asymmetries in Cue Competition. *Journal of Experimental Psychology-General, 121*(2), 222-236.

Wills, A. J., Lavric, A., Croft, G. S., & Hodgson, T. L. (2007). Predictive learning, prediction errors, and attention: Evidence from event-related potentials and eye tracking. *Journal of Cognitive Neuroscience, 19*(5), 843-854.

Table 1. Outline of *Experimental Design, and Mean Accuracy*

Condition	*Phase 1*		*Phase 2*	
Backward Blocking	AB+1	(95%)	A+1	(95%)
Unovershadowing	CD+2	(96%)	C-	(94%)
RR Control	EF+3	(96%)	E°	(N/A)
RR & MC	GH+4	(92%)	G+5	(95%)
MC Control	IJ°	(N/A)	I+6	(96%)

Note. Letters denote different foods; "+1" to "+6" denote different allergic reactions; "-" denotes "No Allergic Reaction"; and "°" denotes that the screen simply went blank screen after presentation of the food(s), RR – retrospective revaluation, MC – mediated conditioning . Numbers in brackets indicate percent correct, by trial type, across the last 10 trials of each condition in each training phase.

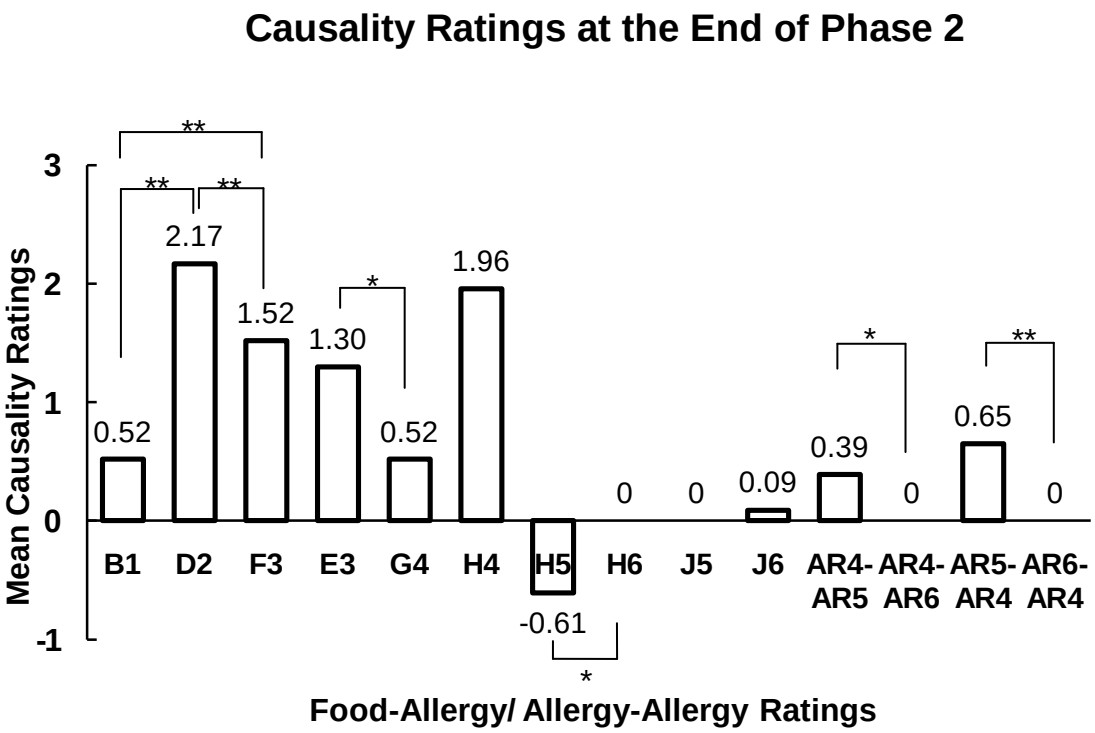

Figure 1. Mean causality ratings across all subjects for different food-allergy (denoted by letter-digit labels) and allergy-allergy pairs. In letter-digit labels, letters denotes the type of food, AR denotes allergic reaction, and the number that follows indicates the allergic reaction being rated. The causality ratings ranged from +3 (the food causes the allergic reaction) through zero (the food and allergic reaction are unrelated) to -3 (the food prevents the allergic reaction). ** indicates $p < .01$ and * indicates $p < .05$.

Supplementary Materials

Simultaneous Backward Conditioned Inhibition and Mediated Conditioning

Steven Graham[1], Hans Lee Jie[1], Chan Hui Minn[1] I.P.L. McLaren[2] & Andy J. Wills[2]

[1]Department of Psychology, National University of Singapore, Singapore.

[2]School of Psychology, Exeter University, U.K.

In these Supplementary Materials, we present the means of all ratings collected at the end of phase 1, and at the end of phase 2.

Mean Food-Allergy Causality Ratings at the end of Phase 1

	AR 1	AR2	AR3	AR4	AR5	AR 6	No AR
A	2.39	0.00	0.00	0.00	0.00	0.00	0.00
B	2.17	0.13	0.00	0.09	0.00	0.00	0.04
C	0.30	2.26	0.04	0.00	0.00	0.00	0.00
D	0.13	2.17	-0.04	0.00	0.00	0.00	0.00
E	0.00	-0.04	2.30	0.00	0.00	0.00	0.00
F	0.00	0.04	2.22	0.00	0.00	0.00	0.04
G	0.09	0.00	0.00	2.52	0.00	0.00	0.00
H	-0.09	0.00	0.00	2.39	0.00	0.00	0.00
I	-0.04	-0.04	0.30	0.00	0.09	0.09	0.17
J	-0.04	-0.04	0.30	-0.04	0.09	0.09	0.30

Note. Letters denote foods, AR denotes Allergic Reaction. Mean causality ratings ranging from -3 (the food prevents the allergic reaction) through zero (the food and allergic reaction are not related) to +3 (the food causes the allergic reaction).

Mean Food-Allergy Causality Ratings at the end of Phase 2

	AR 1	AR2	AR3	AR4	AR5	AR 6	No AR
A	2.83	0.00	0.00	0.00	0.00	0.00	0.00
B	0.52	0.00	0.00	0.13	0.00	0.00	0.87
C	0.00	0.13	0.00	0.00	0.00	0.00	2.96
D	0.00	2.17	0.00	0.13	0.00	0.00	0.00
E	0.00	0.04	1.30	0.00	0.00	0.04	0.04
F	0.00	0.00	1.52	0.00	0.00	0.00	0.00
G	0.00	0.00	0.00	0.52	2.87	0.00	0.00
H	0.00	0.00	0.00	1.96	-0.61	0.00	0.04
I	0.00	0.00	0.13	0.00	0.00	2.70	0.00
J	0.00	0.00	0.13	-0.04	0.00	0.09	0.35

Note. Letters denote foods, AR denotes Allergic Reaction. Mean causality ratings ranging from -3 (the food prevents the allergic reaction) through zero (the food and allergic reaction are not related) to +3 (the food causes the allergic reaction).

Mean Ratings for Allergy-Allergy Associations at the end of Phase 2

		…causes this allergy →						
		AR 1	**AR2**	**AR3**	**AR4**	**AR5**	**AR 6**	**No AR**
This allergy… ↓	**AR 1**	N/A	0.13	0.04	0.04	0.00	0.09	0.13
	AR 2	0.17	N/A	0.00	0.09	0.00	0.13	0.26
	AR 3	0.00	0.00	N/A	0.13	0.00	0.04	0.00
	AR 4	0.04	0.00	0.00	N/A	0.39	0.00	0.00
	AR 5	0.00	0.17	0.00	0.65	N/A	0.13	0.13
	AR 6	0.13	0.13	0.13	0.00	0.13	N/A	0.04
	No AR	0.04	0.13	0.00	0.00	0.13	0.04	N/A

Note. The mean ratings for one allergy (row) causing another allergy (column). Mean causality ratings ranging from -3 (one allergic reaction prevents the other allergic reaction) through zero (the two allergic reactions are not related) to +3 (one allergic reaction causes the other allergic reaction). AR denotes Allergic Reaction.