Lochmann, T. & Wills, A.J. (2003). Predictive history in an allergy prediction task. In Schmalhofer, F., Young, R. M. & Katz, G. (Eds.) Proceedings of EuroCogSci 03: The European Cognitive Science Congerence. Mahwah, New Jersey: Lawrence Erlbaum Associates.

Predictive History in an Allergy Prediction Task

Timm Lochmann (Timm.Lochmann@urz.uni-heidelberg.de)
Department of Psychology, University of Heidelberg
Hauptstraße 47-51, 69117 Heidelberg. Germany.

A. J. Wills (a.j.wills@ex.ac.uk)
School of Psychology, University of Exeter, Perry Road
Exeter. EX4 4QG. U.K.

Abstract

Two experiments are reported that demonstrate rate of learning in an allergy prediction task can be affected by the predictive history of the cues involved, even if that history relates to outcomes different to those being currently learned about. Predictive history is defined here as a cue's prior status as either a good or a poor predictor of outcomes. Our results are problematic for commonly employed associative theories of human contingency learning but also provide evidence for the sort of associability-change process envisaged by the Mackintosh (1975) theory.

Introduction

In a contingency learning experiment, people are presented with various cues that have some form of predictive relationship to one or more outcomes. A typical example is the allergy prediction task (e.g. Van Hamme & Wasserman, 1994). In this task, the subject is told which foods an imaginary patient has eaten, predicts whether the patient will show an allergic reaction as a result, and receives feedback on that prediction. The nature of what has been learned can then be probed in a number of ways. For example, subjects may be asked for a numerical rating of the strength of the relationship between cue and outcome (e.g. Dickinson & Burke, 1996). Alternatively, one can record the predictions made by subjects about a series of test stimuli (e.g. Shanks, Darby, & Charles, 1998).

In a landmark paper, Dickinson, Shanks and Evenden (1984) suggested that human contingency learning might be explicable in terms of well-established theories of animal learning. They were, however, relatively agnostic about the particular theory to be employed, suggesting the Mackintosh (1975), Pearce and Hall (1980), and Rescorla and Wagner (1972) theories as likely candidates. As the number of experiments on human contingency learning has increased, so has the list of candidate animal learning theories. Whilst the following is by no means exhaustive, the use of Pearce's (1987) configural theory (e.g. by Shanks, Charles, Darby, & Azmi, 1998), a modification of Wagner's (1981) SOP model (Dickinson & Burke, 1996) and modifications of the MKM model (McLaren & Mackintosh, 2000), are three notable examples.

It is possible to identify two main themes in the empirical investigation of associative learning theories of human contingency learning. The first is to provide support for the notion that associative learning theories provide a distinct and superior account of the data to that provided by normative accounts of the type proposed by, for example, Cheng and Novick (1992). Shanks (1995) provides a review of the first decade of research on this issue. This theme is not heavily represented in the current article, primarily because we consider normative and associative accounts to be different levels of explanation rather than directly competing accounts. Nevertheless, normative theorists may wish to consider how the results we report can be accommodated by their theories.

The second main theme (in the empirical investigation of associative learning theories of human contingency learning) is to attempt to reject some of the many competing associative accounts that have been proposed. For example, Rescorla and Wagner (1972), Pearce and Hall (1980), and Pearce (1987) may all be rejected in their original form on the basis that a cue→outcome association can be modulated in the cue's absence by training on a different cue with which it was previously paired. Evidence of such *retrospective revaluation* comes from a number of studies (e.g. Dickinson & Burke, 1996; Le Pelley & McLaren, 2001; Shanks, 1985) but there is no process by which it can happen in the aforementioned theories. Further problems arise from evidence that human contingency learning is more resistant to retroactive interference than most associative accounts would predict (Shanks, Darby et al., 1998).

This second theme - the rejection of specific associative accounts - continues in the current paper. We report a contingency learning experiment whose results appear to create further problems for many of the aforementioned theories. The topic of investigation is how a cue's prior status as either a good or a poor predictor affects the rate at which it forms associations in future.

Interest in this question dates back at least as

far as Lawrence's (1949) demonstration that learning of a successive brightness discrimination shows positive transfer to a simultaneous brightness discrimination. Lawrence maintained that this result must be due to increased attention to the stimuli involved, although that interpretation is debatable (see Seigel, 1967). Subsequent investigations tended to concentrate on positive transfer to different stimuli on the same dimension. For example, in the phenomenon of *transfer along a continuum*, training on an easy discrimination (e.g. black vs. white) can facilitate learning of a difficult discrimination (two shades of gray) on the same dimension (e.g. Lawrence, 1952; Mackintosh & Little, 1970; Pavlov, 1927). It is also the case that shifting to a discrimination within the same dimension is easier than shifting to a discrimination on a different dimension (Shepp & Eimas, 1964; Wolff, 1967).

In the current experiments, we introduce a novel design that returns to the issue of the predictive history of specific cues. The basic design is summarized in Table 1. In phase one, all subjects are taught that each of the eight cue pairs is reliably associated with one of two different allergies. The cues shown in bold (A, B, C and D) are always paired with the same outcome and are therefore perfectly predictive of the type of allergic reaction. The cues shown in italics (V, W, X and Y) are paired equally often with each of the two outcomes and are therefore non-predictive of allergy outcome.

Table 1: Experimental design

Phase One	Phase Two	
A*X* → 1	Predictive group	
B*X* → 2	**AB** → 3, **CD** → 4	Familiar
A*Y* → 1		
B*Y* → 2	KL → 3, MN → 4	Novel
	Non-predictive group	
C*V* → 1		
D*V* → 2	*XY* → 3, *VW* → 4	Familiar
C*W* → 1		
D*W* → 2	KL → 3, MN → 4	Novel

In phase two, subjects are taught about two compounds whose components were previously experienced in phase one (*familiar* compounds). They are also taught about two compounds whose components were not previously experienced in phase one (*novel* compounds). There are two between-subject groups in phase two, the *predictive history* group and the *non-predictive history* group. Note that phase two employs different allergy types to phase one, and that all individual cues and their compounds are perfectly predictive of these novel outcomes. The two between-subject groups differ only in the cues' history of predictiveness with previous allergy types.

We predicted that the rate of learning in phase two would be higher for the predictive group than the non-predictive group. We hypothesized that subjects would infer from the cue' s history in phase one that it was either a good or poor predictor. This, in turn, would lead to accelerated learning towards previously good predictors and/or retarded learning towards previously poor predictors in phase two. Whilst this idea seems intuitively plausible, such an effect is not predicted by the associative theories most commonly applied to studies of contingency learning in humans. For example, Rescorla and Wagner (1972) and Pearce (1987), the two most commonly applied associative theories, represent the predictive history of cues solely by the associations formed between cue representations and outcome representations. The fact that different outcomes are used in our two phases would seem to constrain both theories to predict no effect. If one allows for generalization from the outcomes of phase one to the outcomes of phase two, this does not improve matters as each pair of cues in phase two seems likely to evoke representations of outcomes one and two equally. For example, the cue pair AB evokes both outcome 1 because of the AX and AY trials, and outcome 2 because of the BX and BY trials.

Our design also defeats explanation based on the formation of associations between different cues as there are no reliable cue-cue pairs in phase one. This is relevant to, for example, the salience change process proposed by McLaren & Mackintosh (2000) which relies on the formation of cue-cue associations. Salience change processes of this nature therefore do not allow one to predict a difference between our two groups.

Another important aspect of our design is that, in phase one, each trial involves one cue that is perfectly predictive and another that is entirely non-predictive. This causes problems for theories (e.g. Pearce & Hall, 1980) that assume rate of learning is determined by the predictability of the compound rather than of the individual cues. Such theories would predict no difference between our two groups.

In summary, if the predictive group learns the phase-two discrimination in fewer trials than the non-predictive group then this would demonstrate an intuitively plausible predictive history effect under conditions where many common theories of human associative learning seem constrained to predict no effect. Such a demonstration therefore seems of some theoretical importance.

Experiment 1

Method

Subjects and materials The subjects were 30 student volunteers from the University of Exeter. The

experiment was run on a Pentium III PC with 17" monitor. Subjects were tested individually in a quiet cubicle and the monitor was positioned about 80cm in front of them. Responses were collected via a standard two-button mouse. The cue names used were: potatoes, beans, bread, milk, ice cream, oranges, apples, bananas, sprouts, sweetcorn, mushrooms, carrots, pasta, tomatoes, garlic, onions. The allergy words used were: itch, rash, nausea, dizzyness. The allocation of cue words and allergies to the logical design presented in Table 1 was determined randomly for each subject. The large number of cues and outcomes in our design makes a counter-balanced allocation of cues to the logical design impractical.

Procedure Subjects were presented with on-screen instructions that asked them to take the role of an allergist treating an allergic patient. They were told that their task was to predict which allergy the patient would develop if they were exposed to the cues shown.

On each trial, the two cue words were presented towards the top of the screen and were horizontally aligned. The two allergies appropriate for the phase (i.e. allergies 1 and 2 for phase one, allergies 3 and 4 for phase two) were presented towards the bottom center of the screen and were vertically aligned. In the center of the screen, a large rectangle contained the phrase "Please state your diagnosis".

Subjects indicated their diagnosis by clicking on one of the two allergy words and then clicking an "OK" button at the bottom right of the screen. Feedback was provided by the "state your diagnosis" rectangle turning red and displaying the word "false" or turning green and displaying the word "correct". A small blue arrow indicated the correct prediction by pointing to the appropriate allergy. The subject moved on to the next trial by clicking on the feedback rectangle.

The eight trial types for phase one (see Table 1) were presented sequentially and in a random order. Trial order randomisation was via eight-item blocks, each block containing exactly one instance of each trial type. Ends of blocks were not signalled to subjects.

Training was to criterion - once the subject had reached a criterion of two consecutive errorless blocks, they moved on to phase two. Phase one was also terminated if a subject completed 240 trials. At the beginning of phase two, subjects were told they were moving on to a second patient whose allergies were different. The four trial types appropriate to the subject's group (see Table 1) were then presented sequentially and in a random order. Trial order randomisation was via eight-item blocks, each block containing exactly two instances of each trial type. The phase two procedure was otherwise identical to phase one, with training to the same criterion and termination of the experiment after the same number of trials.

The left-right position of cue words was randomly determined for each trial and subject. The position of allergy words was randomly determined for each phase and subject.

Results and discussion

Six subjects failed to reach criterion in phase one and hence were excluded from all analysis. A Type 1 error rate of 5% ($\alpha = .05$) was used for all statistical tests.

Subjects in the predictive condition took a mean of 10.6 blocks to reach criterion in phase 2, whilst subjects in the non-predictive condition took a mean of 6.5 blocks. This effect, whose trend is opposite to our predictions, did not approach significance, $t(22) = 0.98$

Faced with this absence of information, we derived the following post-hoc hypothesis: Given that all cues are fully diagnostic in phase 2, it would seem likely that if there are any effects of predictive history they would be easiest to detect early in phase 2. We defined “early” as the first block of phase 2.

Figure 1 shows, for each condition, the number of familiar and novel stimuli responded to correctly in the first block of phase two. Analysis of variance with one between-subjects variable (predictive history) and one within-subjects variable (novelty) confirms that the observed interaction is reliable, $F(1,22) = 8.56$.

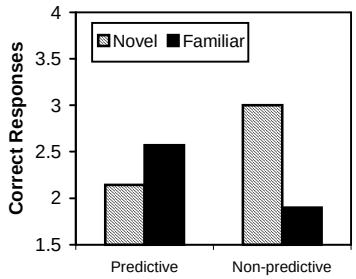

Figure 1: Mean number of correct responses in the first block of phase two of Experiment 1.

Predictive history does not approach significance as a main effect, $F(1,22) = 0.33$, and the same is true for novelty, $F(1, 22) = 1.20$. Performance on novel compounds is significantly greater in the non-predictive condition than the predictive condition, $t(22) = 2.14$. The trend to greater performance on familiar compounds in the predictive condition is not significant, $t(22) = 1.44$.

These data provide some indication that the rate of learning in a human contingency learning task can be affected by the predictive history of the cues involved, even if that history relates to outcomes different to those currently being learned about. However, a critic might justifiably point to at least three shortcomings. First, there is no reliable effect of predictive history on blocks-to-criterion. Second the analysis of block 1 does not reveal any significant effect

for familiar compounds, only for novel compounds. Third, the choice of block 1 as the unit of analysis, rather than a smaller or larger sub-section of phase 2, is post-hoc and seems a little arbitrary. Further work was therefore required to establish the validity of the hypothesized predictive history effect.

Table 2. Design of Experiment 2.

Phase One	Phase Two	
	Predictive group	
	AB → 3, **CD** → 4	Familiar
A*X* → 1		
B*X* → 2	GH → 3, IJ → 4	Novel
A*Y* → 1	KL → 3, MN → 4	
B*Y* → 2	OP → 3, QR → 4	
	Non-predictive group	
C*V* → 1	*XY* → 3, *VW* → 4	Familiar
D*V* → 2		
C*W* → 1	GH → 3, IJ → 4	Novel
D*W* → 2	KL → 3, MN → 4	
	OP → 3, QR → 4	

Experiment 2

In our second experiment, we attempted to make phase two more difficult by increasing the number of novel compounds seen by the subject to six. Table 2 shows the modified design. We hypothesized that this design would increase the likelihood of finding a predictive history effect as indexed by blocks-to-criterion as it would slow down learning in phase two. One likely side-effect of such a change is that any effects that may occur early on in phase 2 would become harder to detect as subjects would be nearer chance (due to the increased difficulty).

Method

Subjects and materials The subjects were 51 student volunteers from the University of Exeter. Subjects were tested in groups of up to 16 using a suite of identical Pentium 4 PCs with 17" monitors. Each computer and subject was positioned in a different, semi-enclosed cubicle in a manner such that the subjects could not observe each other. The cue names used were: plum, dust, pear, apple, grapefruit, lemon, lime, perfume, tangerine, grape, avocado, peach, pollen, melon, orange, turnip, parsnip, beetroot, carrot, paint. The allergy words used were: itch, rash, nausea, dizziness. The allocation of cue words and allergies to the logical design presented in Table 2 was determined randomly for each subject, within the constraints that a) all cues in phase 1 were fruit, b) each allergy in phase 2 was associated with one novel compound of two fruits, one novel compound of two vegetables, and one novel compound of two non-foods. These constraints mean that subjects should be roughly comparable in terms of how similar novel cues are to cues seen in phase one.

Procedure Subjects were familiarized with each of the 20 cue words by presenting each of the 190 possible pairs of cues in turn and requesting a similarity rating for each pair. The rating requested ranged from 1 ("not similar") to 9 ("very similar"). The task was self-paced and order of presentation was randomized for each subject. The allergy prediction task immediately followed. The procedure was identical to Experiment 1 apart from the greater number of novel cues in phase 2 (see table 2), and that phase 2 was terminated after a maximum of 160 trials (rather than the 240 in Experiment 1).

Results and discussion

Twenty subjects failed to reach criterion in phase one and hence were excluded from all analysis. A Type 1 error rate of 5% ($\alpha = .05$) was used for all statistical tests.

In accordance with our hypothesis, subjects in the predictive condition reached criterion in phase two significantly more quickly, $t(29) = 2.78$, predictive condition mean = 13.35 blocks, non-predictive condition mean = 18.21 blocks. The effect remained significant when assessed non-parametrically, Mann-Whitney $U_{14,17} = 60$. These data provide further support for the idea that rate of learning can be affected by the predictive history of the cues, even if that history relates to outcomes different to those currently being learned about.

We also analyzed proportion of correct responses over the first sixteen trials of phase 2. This corresponds to the sub-section of the data analyzed for Experiment 1 in the sense that it includes the first two occurrences of each trial type in phase 2. It is also the largest data set over which proportion correct can be calculated in an unbiased way due to our criterion-determined termination of phase 2 (from trial 17 onwards, some subjects may have completed the experiment). However, analysis of variance failed to reveal any significant main effects of predictive history , $F(1,29) = 0.12$, or stimulus type, $F(1, 29) = 1.12$, in this data set. The interaction term did not approach significance, $F(1,29) = 0.16$.

One reason for the absence of any significant effects early in training may be that the increased difficulty of this variant of the task means that very little is learned in the first two presentations of each trial type.

General discussion

The present findings indicate that the rate of learning in a contingency learning task can be affected by the predictive history of the cues involved, even if that history relates to outcomes different to those currently

being learned about. This effect was demonstrated under conditions where a number of commonly employed theories of associative learning seem constrained to predict the absence of an effect (e.g. McLaren & Mackintosh, 2000; Rescorla & Wagner, 1972; Pearce, 1987; Pearce & Hall, 1980).

Le Pelley & McLaren (2003), using a design similar to and directly inspired by our procedure, have subsequently demonstrated that a similar effect can also be found if one employs rating scales rather than trials-to-criterion as the dependent measure. Taken together, their study and ours provide strong empirical support for the reality of a predictive history effect in human contingency judgments.

Le Pelley & McLaren (2003) also suggest a mechanism for the production of such an effect based on the Mackintosh (1975) theory of associative learning. In this theory, associations are assumed to form between each cue and outcome. Associative strength changes according to the rule

$$\Delta V_A = \alpha_A \theta (\lambda - V_A) \qquad (1)$$

where V_A is the associative strength between cue A and an outcome, α_A is the associability of cue A, θ is a fixed learning rate parameter and λ is the limit of associative strength. The cue-specific associability parameter α is not fixed; its value varies on each trial according to the cue's predictiveness. Specifically, $\Delta\alpha_A$ is positive if

$$| \lambda - V_A | < | \lambda - V_X | \qquad (2)$$

and negative otherwise, where V_X is the sum of the associative strengths of all cues other than A present on the trial.

Applying the Mackintosh (1975) theory to our experiment, phase one training should result in the associability of cues A to D being higher than the associability of cues V to Y because cues A to D perfectly predict their outcome and so will acquire greater associative strengths. This, in turn, will lead to V_X being smaller than V_A for predictive compounds and equal to V_A for non-predictive compounds, resulting in associability changes via Equation 2. In phase two, this will result in the faster learning to familiar compounds in the predictive group than in the non-predictive group. Such a trend is observed, albeit non-significantly, in the first block of phase two of our Experiment 1 and could plausibly underlie the effect, in our Experiment 2, that the predictive group reaches criterion faster in phase two than the non-predictive group. Kruschke (2001) has recently proposed a model which, as he notes, is under certain conditions strikingly similar to the Mackintosh (1975) theory.

Whilst Le Pelley & McLaren's result, our trend with familiar compounds in Experiment 1 and our trials-to-criterion difference in Experiment 2 can all be potentially explained by the Mackintosh (1975) theory, an alternative explanation in terms of proactive interference (Underwood, 1957) is also possible.

In phase one, cues are associated with outcome 1, outcome 2 or both. In phase two, these pre-existing associations might proactively interfere with the formation of associations in phase two, hence retarding the rate of learning. The non-predictive group is retarded more than the predictive group because, in the former, each cue has two potential sources of interference (e.g. X→3 receives interference from X→1 and X→2) whilst, in the latter, each cue has just one source of interference (e.g. A→3 just receives interference from A→1). The validity of this argument rests on the assumption that it is the number of sources of interference, rather than the number of times each source occurs, which is the dominant effect. If one assumes that strength of a memory is an increasing, negatively accelerated function of number of presentations then number of sources will dominate number of presentations. An increasing, negatively accelerated function is consonant with the assumptions of most associative learning theories.

This proactive interference explanation does not appear to account for the effect we observed for novel compounds. In contrast, it seems possible that the Mackintosh (1975) theory can account for this effect if one employs a slightly more complex version of the theory also discussed in the Mackintosh (1975) paper.

Mackintosh suggests that the change in associative strength between a presented cue and an outcome generalizes *to* other cues to the extent that they are similar to the target cue. Specifically, in the special case of an experiment with just two cues,

$$\Delta V_B = S_{A,B} \alpha_A \theta (\lambda - V_A) \qquad (3)$$

where V_B is the associative strength of the cue changing through generalization (cue B), and $S_{A,B}$ represents the similarity between cue A and cue B. With regard to our Experiment 1, two important aspects of our design need to be underlined. First, the outcomes in phase two are different to those in phase one. We assume this means that V_A is zero for all cues at the start of phase two. Second, all cues in phase two are perfectly predictive of their outcome so cues in phase two do not differ in the $\lambda - V_A$ term of Equation 3. Finally, all cues are assumed to be, on average, of equal similarity to each other due to the randomized allocation of cue names to our design. It is therefore only the α_A term of Equation 3 that determines the extent to which the change in associative strength between a presented cue and an outcome generalizes *to* other cues. The α_A terms for cues are determined by their prior predictive history in phase one.

In summary, the critical prediction of Equation 3 is that changes in associative strength in phase two generalize more effectively from cues with high

associability than from cues with low associability. Therefore, KL and MN receive more generalized change in associative strength from AB and CD in group predictive than they do from XY and VW in group non-predictive. Receiving more generalized associative strength retards learning in this instance because generalization will occur equally to cues with the same outcome as the target cue and to cues with a different outcome than the target cue. Under a variety of performance rules, including the ratio rule, difference rule and a winner-take-all system, adding an fixed amount to all associative strengths reduces response accuracy (see e.g. Wills, Reimers, Stewart, Suret, & McLaren, 2000). One may therefore predict that the novel compounds in our experiment will be responded to less accurately in the predictive group than in the non-predictive group which is, of course, what is observed. Whilst generalized changes in associative strength from KL and MN will also retard learning of the familiar compounds AB, CD, XY and VW, the amount of retardation is predicted to be equivalent in both conditions. In summary, then, our effect with novel cues, although based on a post-hoc analysis of the data, nevertheless provides some evidence that may favor the Mackintosh (1975) account over an account based on proactive interference.

Thus far, the reader might be left with the impression that, from the plethora of associative accounts available, the Mackintosh (1975) theory is the only adequate associative model of contingency learning in humans. Such an impression would be inaccurate due to the model's inability to account for retrospective revaluation effects (e.g. Dickinson & Burke, 1996) and for results that imply a role for configural processing (e,g, Shanks, Charles et al., 1998). Nevertheless, from the alternatives considered, the Mackintosh (1975) theory seems to provide the most adequate account of the sort of predictive history effect demonstrated in the current experiments.

References

Cheng, P.W., & Novick, L.R. (1992). Covariation in natural causal induction. *Psych. Review, 99*, 365-382.

Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of casuality judgements. *Quart. J. Exp. Psy, 49B*, 60-80.

Dickinson, A., Shanks, D., & Evenden, J. (1984). Judgement of act-outcome contingency: The role of selective attribution. *QJEP, 36A*, 29-50.

Kruschke, J.K. (2001). Toward a unified model of attention in associative learning. *J. Math. Psychol., 45*, 812-863.

Lawrence, D.H. (1949). Acquired distinctiveness of cues. *J. Exp. Psychol., 39*, 770-784.

Lawrence, D.H. (1952). Transfer of a discrimination along a continuum. *J. Comp. Phys. Psychol., 45*, 511-516.

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

Le Pelley, M., & McLaren, I.P.L. (2003). Learned associability and associative change in human causal learning. *QJEP, 56B*,1, 68-79.

Mackintosh, N.J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psych. Review, 82*, 276-298.

Mackintosh, N. J. & Little, L. (1970). An analysis of transfer along a continuum. *Canadian Journal of Psychology, 24*, 362-369.

McLaren, I. P. L., & Mackintosh, N. J. (2000). An elemental model of associative learning. *Animal Learning & Behavior, 28*(3), 211-246.

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

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

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning. *Psychological Review, 87*, 532-552.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II.* New York: Appleton-Century-Crofts.

Shepp, B.E. & Eimas, P.D. (1964). Intradimensional and extradimensional shifts in the rat. *J. Comp. Phy. Psychol., 57*, 357-361.

Siegel, S. (1967). Overtraining and transfer processes. *J. Comp. Phy. Psychol., 64*, 471-477.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *QJEP, 37B*, 1- 21.

Shanks, D. R. (1995). *The psychology of associative learning*. Cambridge: Cambridge University Press.

Shanks, D. R., Charles, D., Darby, R. J., & Azmi, A. (1998). Configural processes in human associative learning. *J. Exp. Psychol: Learn., Mem. & Cog., 24*(6), 1353-1378.

Shanks, D. R., Darby, R. J., & Charles, D. (1998). Resistance to interference in human associative learning. *J. Exp. Psych.: An. Behav. Proc., 24*(2), 136-150.

Underwood, B. J. (1957). Interference and forgetting. *Psychological Review, 64*(1), 49-60.

Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgements. *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* Hillsdale, NJ: L.E.A.

Wills, A. J., Reimers, S., Stewart, N., Suret, M., & McLaren, I. P. L. (2000). Tests of the ratio rule in categorization. *QJEP, 53A*(4), 983-1011.

Wolff, J.L. (1967). Concept-shift and discrimination-reversal learning in humans. *Psychol. Bull.,68,* 369-408.