Transfer of associability and relational structure in human associative learning.

Evan J. Livesey[1], Hilary J. Don[1], Metin Uengoer[2], & Anna Thorwart[2]

[1] The University of Sydney

[2] Philipps-Universität Marburg

Address correspondence to:
Evan Livesey
School of Psychology,
University of Sydney,
Sydney, NSW, 2006,
Australia,
Email: evan.livesey@sydney.edu.au

Abstract

A wealth of recent studies have demonstrated that predictive cues involved in a linearly solvable component discrimination gain associability in subsequent learning relative to non-predictive cues. In contrast, contradictory findings have been reported about the fate of cues involved in learning biconditional discriminations in which the cues are relevant but none are individually predictive of a specific outcome. In three experiments we examined the transfer of learning from component and biconditional discriminations in a within-subjects design. The results show a greater benefit in associability for cues that had previously served as predictive cues in a component discrimination than cues previously used in a biconditional discrimination. Further, new biconditional discriminations were learned faster when they were composed of cues that were previously trained in separate biconditional discriminations. Similarly, new component discriminations were learned faster when they were composed of cues that were previously trained in a separate component discriminations irrespective of whether they were previously predictive or previously non-predictive. These results provide novel evidence that *cue-specific* learning of relational structure affects subsequent learning, suggesting changes in cue processing that go beyond simple changes in cue associability based on learned predictiveness.

Keywords: learned predictiveness, biconditional discrimination, associability, relational learning

Transfer of associability and relational structure in human associative learning.

Many theories of learning make the assumption that acquiring knowledge about the relationship between a cue and its consequences will change the manner in which this cue is processed in the future. For instance, several prominent theories propose that selective attention is directed towards cues that have been useful predictors of meaningful outcomes in the past (Kruschke, 2001; Le Pelley, 2004; Mackintosh, 1975; Pearce & Mackintosh, 2010). This selective processing of previously predictive cues leads to changes in their *associability*, the rate at which those cues enter into new associations, for instance, when the outcomes that they predict change.

Theories of associative learning that appeal to learned changes in selective attention have received considerable support from demonstrations of the *learned predictiveness* (LP) effect. This effect, first reported by Lochman and Wills (2003) and Le Pelley and McLaren (2003), demonstrates that the rate at which people learn about predictive cues depends on whether those cues have been predictive of other task-relevant outcomes in the past. In Le Pelley and McLaren's (2003) procedure (the design of which is shown in Table 1), participants completed a causal learning task in which they played the role of an allergist determining the cause of a patient's allergic reactions. For instance, a participant may have learned that when their patient ate cheese and apple or cheese and bread, they suffered from fever, and when they ate chicken and apple or chicken and bread, they suffered from rash. In this example, two of the cues, cheese and chicken, are perfectly predictive of the outcomes, fever and rash, respectively. In contrast, apple and bread are less predictive (for simplicity, we will say *non-predictive*) as they do not inform the participant whether the outcome will be fever or rash. In a second stage, participants then learn about foods eaten by a different fictitious patient. New cue combinations consisting of one predictive and one non-predictive cue are paired with a new outcome, but unlike the first stage, all food cues are equally relevant and predictive of the symptoms. Despite having objectively equivalent relationships with the outcome, the previously predictive and previously non-predictive cues are not learned about equally in this phase. Participants learn more and faster about the previously predictive cues than the previously non-predictive cues. That is, these cues possess greater associability as a consequence of their predictive status in Stage 1.

Table 1
Stage 1 training, Stage 2 training, and test trials used by Le Pelley & McLaren (2003).

Stage 1	Stage 2	Test	LP effect
AW – O1	AY – O3	AD	Prediction of O3:
AX – O1	BZ – O4	BC	AD > XY
BW – O2	CW – O4	XY	
BX – O2	DX – O3	WZ	Prediction of O4:
			BC > WZ
CY – O1	JK – O3	JM	
CZ – O1	LM – O4	KL	
DY – O2	NO – O3	NO	
DZ – O2	PQ – O4	PQ	

Note: Letters refer to individual cues, A-D: predictive components, W-Z: non-predictive components, J-Q: novel filler cues in Stage 2. O1, O2, O3, and O4 refer to four outcomes.

The LP effect was inspired by predictions of attention-based learning models. Mackintosh (1975) in particular, and has been interpreted as being broadly consistent with several of these models. The central operation of the Mackintosh model is a competitive change in attention based on *relative* predictiveness. During Stage 1, predictive cues gain associability by virtue of being the most predictive stimuli present on each trial and non-predictive cues lose associability by virtue of being poorer predictors.

Considerable research has now been devoted to the LP effect (for a recent review, see Le Pelley et al., 2016). Interestingly, contrary to the predictions of Mackintosh's (1975) theory, competition between potential predictors is not necessary to drive attention change in the manner that the model anticipates. Le Pelley et al. (2010) found that training with individually presented cues during the initial learning phase is sufficient to produce stronger learning for predictive than non-predictive cues when they were subsequently presented in new predictive cue-outcome relationships, also presented individually. In addition, Livesey et al. (2011) found that relative predictiveness, within a trial, produced by presenting a predictive and non-predictive cue simultaneously on each learning trial, was no more or less effective at generating LP effects than absolute predictiveness produced by presenting two equally predictive cues together or two equally non-predictive cues together, a result replicated with a similar design by Kattner (2015). This suggests that the processes responsible for the LP effect are still yet to be fully determined.

In one of the experiments reported by Livesey et al. (2011; Experiment 2), the associability of previously predictive and non-predictive cues was tested against control cues that were used in a biconditional discrimination. In a biconditional discrimination, pairs of cues predict the outcome on each trial but the cues are arranged in such a way that no single cue is correlated with a particular outcome (e.g. AB-O1, BC-O2, CD-O1, DA-O2). This control provides a condition in which the cues are equally relevant to the discrimination and thus must continue to be attended in order to learn the solution but do not individually predict a particular outcome. Livesey et al. found that there was more evidence of learning about predictive cues than biconditional cues when they were paired together in Stage 2. In contrast, non-predictive and biconditional cues appeared to be learned about to the same extent when they were paired together in Stage 2. These results were assessed by recombining cues in a subsequent test, as shown in Table 2, in which participants provided predictive ratings for each of the possible outcomes.

These results suggest that, despite being relevant to the discrimination, biconditional cues do not retain their associability in the same way as predictive cues that are individually associated with an outcome. In subsequent experiments, Livesey et al. (2011; and Kattner, 2015) also found that associability change is not based on relative predictiveness; cues occurring in compounds where each cue equally predicted the outcome were found to be just as associable as predictive cues trained with non-predictive competitor cues. Therefore the difference in associability between predictive and biconditional cues is unlikely to be a consequence of the two biconditional cues being equally relevant and thus neither winning out over the other in competition for attention. The result suggests that it is the strong association with a specific outcome rather than relevance to solving the task that maintains high associability for predictive cues.

Table 2
Design used by Livesey et al. (2011, Experiment 2)

Stage 1		Stage 2	Test	Effect on
Component	*Biconditional*		*(Summation trials)*	Summation trials
AW – O1	JN – O1	AJ – O3	AD	Prediction of O3:
AX – O1	JO – O2	BK – O4	BC	AD > JM
BW – O2	KN – O2	CL – O4	JM	Prediction of O4:
BX – O2	KO – O1	DM – O3	KL	BC > KL
CY – O1	LP – O1	WN – O3	WZ	Prediction of O3:
CZ – O1	LQ – O2	XO – O4	XY	WZ = NQ
DY – O2	MP – O2	YP – O4	OP	Prediction of O4:
DZ – O2	MQ – O1	ZQ – O3	NQ	XY = OP

Note: Letters refer to individual cues, A-D: predictive components, W-Z: non-predictive components, J-Q: equally relevant stimuli from biconditional discrimination. O1, O2, O3, and O4 refer to four neutral outcomes. The experiment contained three types of test trials (trained, summation and negation) but for simplicity we only show the summation trials here.

A different conclusion can be drawn from a study by Uengoer and Lachnit (2012). Their results indicated that associability is higher for biconditional cues than non-predictive cues, while predictive and biconditional cues did not seem to differ in their associability (see also, Kruschke, 1996). Uengoer and Lachnit trained participants to categorize stimuli that varied on two dimensions (e.g., colour and shape). In Stage 1, participants either received a “component” discrimination for which cues belonging to one dimension were predictive of category membership and cues from the other dimension were non-predictive, or were trained with a biconditional discrimination in which category membership depended on the combination of cues from both dimensions. In Stage 2, all participants were trained with a component discrimination in which the stimuli were characterized by novel cues from one of the previously trained dimensions and cues from a novel dimension (e.g., orientation). In this second discrimination, either the novel or previously trained dimension was predictive, and the other non-predictive, manipulated between experiments. In Experiment 1, the previously trained dimension was predictive in Stage 2 and the novel dimension was non-predictive. Under these transfer conditions, the Stage 2 discrimination was acquired more rapidly when the previously trained dimension had formed part of a biconditional discrimination in Stage 1 rather than the non-predictive dimension in a component discrimination. This result suggests that it was easier to attend to the Stage 1 biconditional dimension than to the Stage 1 non-predictive component dimension. In Experiment 2, the novel dimension was predictive in Stage 2 and the previously trained dimension was non-predictive. In contrast to Experiment 1, now acquisition of the second discrimination proceeded faster when the previously trained dimension constituted the Stage 1 non-predictive component dimension than when it was part of the Stage 1 biconditional discrimination. This result suggests that it was easier to learn to *ignore* the Stage 1 non-predictive component dimension than the Stage 1 biconditional dimension. In both of their experiments, acquisition of the second discrimination was independent of whether one of the dimensions was previously trained as predictive for a component discrimination or as part of a biconditional discrimination (that is, the Stage 1 predictive component dimension and the Stage 1 biconditional dimension appeared to be equally easy to attend and equally easy to ignore). These results suggest that the relevance of a cue to solving the discrimiantion, rather than its correlation with a particular outcome, determines the associability of a cue in new learning (for a similar result in animal learning, see George & Pearce, 1999).

At face value, the results reported by Uengoer and Lachnit (2012; see also, Kruschke, 1996) and Livesey et al. (2011) appear to be in direct conflict. However, there are several important methodological differences between these studies (see Table 3 for a summary). Uengoer and Lachnit used cues that could be grouped according to particular dimensions, and different cues belonging to the same dimension served the same functional role for the discrimination problem (predictive, non-predictive, or biconditional). In contrast, Livesey et al. used discrete cues (line drawings of objects) that were not easily grouped in this way. Uengoer and Lachnit also used a between-subjects design in which participants were only exposed to one type of discrimination during Stage 1, whereas Livesey et al. used a within-subjects design in which participants solved biconditional and component discriminations simultaneously. Importantly, Uengoer and Lachnit used accuracy during a Stage 2 learning task designed to result in positive versus negative transfer to test for associability changes whereas Livesey et al. used a subsequent ratings phase. This is noteworthy because the rating test trials differed to those used during Stage 2 learning and it is possible that generalization to these test trials might be affected by factors other than the strength of the associations of the individual cues.

To provide an example, take two Stage 2 compound trials presented in Livesey et al.'s (2011) Experiment 2, AJ - O3 and DM - O3, as shown in Table 2. On test, there were stronger learning scores for the compound AD than for JM because participants gave higher O3 predictive ratings and lower O4 predictive ratings when presented with AD. From this result we may infer that participants learned more about A and D than about J and M during Stage 2, as is the typical learned predictiveness conclusion. However this relies on an assumption that learning generalises equally to AD and JM. An alternative possibility is that participants recall that the solutions to the biconditional discriminations involving J and M in Stage 1 were not linearly solvable (i.e. the precise configurations of cues mattered for the solution) and thus they display less confidence in their predictions for JM than for AD. Participants may also infer that they should make the opposite prediction (that O4 should occur, following the structural properties of a biconditional discrimination) or that they should not make a differential prediction at all (either O3 or O4 could occur with equal likelihood since the compound has not been seen previously). The learning of the relational structure of the problems in Stage 1 could therefore produce this result without there being a material difference in the attention paid to, or associability of, individual cues during learning.

Table 3.
Summary of methodology and results from Livesey et al. (2011) and Uengoer & Lachnit (2012)

Study	**Design**	**Cues**	**Measure**	**Associability results**
Livesey et al. (2011)	Within-subjects	Discrete	Subsequent test phase	Predictive > Biconditional = Non-predictive
Uengoer & Lachnit (2012)	Between-subjects	Dimensions	Subsequent learning phase	Predictive = Biconditional > Non-predictive

In summary then, there are two issues with basing conclusions solely on the test phases used in the Livesey et al. (2011) experiments. The first is that participants may learn some or all of the relational structure of the biconditional discrimination, which may affect the way that the learner generalises new information that is subsequently learned about those cues. Second, there may be a more general reduction in confidence in using new information learned about previous biconditional cues because the biconditional discrimination is more difficult and more complex to solve. The effect reported by Livesey et al. thus requires replication using a test that does not involve generalization of new learning to novel compounds. It also remains to be seen whether participants readily learn the structural properties of the biconditional discrimination and generalize them to new learning.

In three experiments, we investigated transfer of learning from component and biconditional discriminations. Participants first learned several linearly solvable component discriminations, in which one of the cues in each compound is predictive of the correct outcome and the other cue was non-predictive, as well as biconditional discriminations, where only the configuration of both cues is predictive of the correct outcome. In Experiment 1, this Stage 1 training, which is conceptually equivalent to the design used by Livesey et al. (2011), was followed by different Stage 2 tasks designed to test for transfer effects using a similar logic to those employed by Uengoer and Lachnit (2012). In Experiments 2 and 3, we then tested directly for the possibility of changes in generalization and performance based on learning the relational structure of the biconditional and component discriminations. For instance, in Experiment 2, we were interested in whether learning that a cue was involved in a biconditional discrimination in Stage 1 made it relatively easy to learn about that cue in a new biconditional discrimination in Stage 2, compared to a cue that was instead involved in a component discrimination in Stage 1. Experiment 3 then provided a complementary test of transfer to linearly solvable component discriminations in Stage 2. In each experiment, we tested outcome prediction accuracy as well as prediction confidence throughout Stage 2 training. As noted above, differences in relative confidence about certain cues may impact upon how they are learned in the second task. Additionally, some recent studies have focused on the effects of prediction error and uncertainty on the attention paid to predictive cues (e.g. Beesley, Nguyen, Pearson, & Le Pelley, 2015; Griffiths, Johnson & Mitchell, 2011). While all the relationships in this task were objectively deterministic (that is, in all instances, pairs of cues perfectly predicted one outcome), they may not be experienced with equally certainty, especially since we anticipated that biconditional discriminations would be more difficult than component discriminations. We expect transfer effects that are driven by subjective uncertainty to be accompanied by distinct differences in confidence, particularly early in Stage 2 learning. In contrast, transfer effects based on simple attentional changes or relational structure may not have a clear relationship with differences in confidence.

Experiment 1

Experiment 1 re-examined the difference in associability between predictive, non-predictive and biconditional cues reported by Livesey et al. (2011), but this time using performance in Stage 2 learning as the critical test, to be consistent with past studies that have found positive transfer for stimuli relevant to the solution of a biconditional discrimination (Kruschke, 1996; Uengoer & Lachnit, 2012). Participants first completed typical learned predictiveness component discriminations and biconditional discriminations, which were intermixed. In Stage 2, all participants learned four new component discriminations, in which predictive and non-predictive cues from the Stage 1 component discriminations were trained in new compounds with cues used in the Stage 1 biconditional discriminations. In two of these discriminations, the Stage 1 biconditional cues were now the predictive cues to which participants must attend to learn the new outcome associations, and the cues from the Stage 1 component discriminations (either predictive or non-predictive) were uninformative. In the other two discriminations, the cues used in the Stage 1 component discrimination (either predictive or non-predictive) were now the relevant cues and the biconditional cues were uninformative.

Table 4.
Stage 1 and Stage 2 training contingencies used in Experiment 1.

Stage 1		Stage 2		
Component	Biconditional		Stage 1 component is now predictive	Stage 1 biconditional is now predictive
AW – O1	JN – O1	*Stage 1 predictive*	AJ – O3	BK – O3
AX – O1	JO – O2	*components and*	AL – O3	BM – O4
BW – O2	KN – O2	*Stage 1 biconditional*	CJ – O4	DK – O3
BX – O2	KO – O1		CL – O4	DM – O4
CY – O1	LP – O1	*Stage 1 non-predictive*	WN – O4	XO – O4
CZ – O1	LQ – O2	*components and*	WP – O4	XQ – O3
DY – O2	MP – O2	*Stage 1 biconditional*	YN – O3	ZO – O4
DZ – O2	MQ – O1		YP – O3	ZQ – O3

Note: Letters refer to individual cues, A-D: predictive components, W-Z: non-predictive components, J-Q: equally relevant stimuli from biconditional discrimination. O1, O2, O3, and O4 refer to four neutral outcomes.

Method

Participants. A total of 61 undergraduate students from the University of Sydney (n=33) and the University of Marburg (n=28) participated in the experiment in return for course credit or payment. In this and the subsequent experiments, we applied a Stage 1 performance criterion to ensure that only participants who learned about both the component and biconditional discriminations were used in the final analysis (we would not expect to see transfer effects in Stage 2 from participants who did not learn these discriminations in Stage 1). If a participant did not reach at least 60% mean accuracy on both the component and biconditional discriminations in the final quarter of Stage 1 then they were excluded from the analysis. Eleven participants failed to reach one or both of these 60% cut-offs. Data collection continued until there were 50 participants who passed both of these criteria (41 female, mean age = 22.4 years). Ethical approval for the methods used all experiments was obtained from the University of Sydney Human Research Ethics Committee.

Stimuli, Apparatus and Design. The learning task was presented in the context of a food allergist scenario. Each food cue was presented as a 300 x 300 pixel colour image, with the name of the food printed underneath. The foods used were Apple, Avocado, Banana, Bread, Cheese, Chicken, Coffee, Corn, Eggs, Fish, Garlic, Lemon, Milk, Mushrooms, Pasta, and Peanuts, randomly allocated to serve as cues A-Z for each participant. The four allergic reaction outcomes were Fever, Headache, Nausea, and Rash randomly allocated to serve as O1-O4 for each participant. The food cue and allergy outcome names, as well as all other text and instructions, were presented in English for the University of Sydney students and in German for the University of Marburg students.

The design of the experiment, which is shown in Table 4, was completely within-subjects. In Stage 1, there were four discriminations (two component, two biconditional), and each discrimination involved four compounds with overlapping sets of cues such that each cue appeared in two different compounds. Thus the only difference between the component discriminations and the biconditional is that the component discriminations are linearly solvable (e.g. A predicts O1, B predicts O2) whereas the biconditional discriminations are only solvable on the basis of the compounds (e.g. JN predicts O1, JO predicts O2, KN predicts O2, KO predicts O1). In Stage 2, a further four discriminations were presented, each comprising a recombined set of the cues used in Stage 1. This time, all four were linearly solvable component discriminations with two predictive and two non-predictive cues. The four discriminations differed in terms of the roles that the predictive and non-predictive cues served in Stage 1.

Procedure. Participants were first given typical "allergist task" instructions asking them to play the role of a doctor trying to predict a patient's allergic reactions based on the foods that they have eaten. On each trial of Stage 1, they observed two foods eaten by Patient X and were asked to predict which allergic reaction the patient would suffer as a consequence. Each trial began with a blank intertrial interval of 0.5 seconds, followed by the presentation of the food cues, then followed 0.5 seconds later by the presentation of the two outcome choices (O1 and O2). Participants used the mouse to click on their preferred option. The outcome options then disappeared and the allergic reaction suffered by Patient X was shown along with corrective feedback ("correct" / "incorrect"), which was presented for 2 seconds. Participants completed 16 blocks of trials in Stage 1, with each of the 16 Stage 1 trial types shown in Table 4 presented once per block in a randomised order. Participants were thus learning two component discriminations and two biconditional discriminations concurrently, with 64 trials presented for each discrimination.

At the beginning of Stage 2, participants were told that they now have a new Patient Y who is suffering from different allergic reactions. They were instructed that they would again be trying to predict which allergic reactions Patient Y would suffer after eating foods but that the reactions would be different, and that the way in which they made outcome predictions would also be different. In Stage 2, instead of choosing between two vertically positioned outcomes, participants indicated both their prediction and confidence using a split ratings scale. The scale was displayed horizontally, with confident choices in O3 and O4 at the left and right extremes, and unconfident guess responses in the middle of the scale. However, the scale was split in half in the middle to make it clear that participants were still predicting O3 or O4 even if they were just guessing (see Figure 1). Instructions and visual examples showing how to make ratings were shown at the beginning of Stage 2. Participants then completed 12 blocks of the Stage 2 trials, with each block containing one presentation of each of the Stage 2 trial types shown in Table 4 in a randomised order.

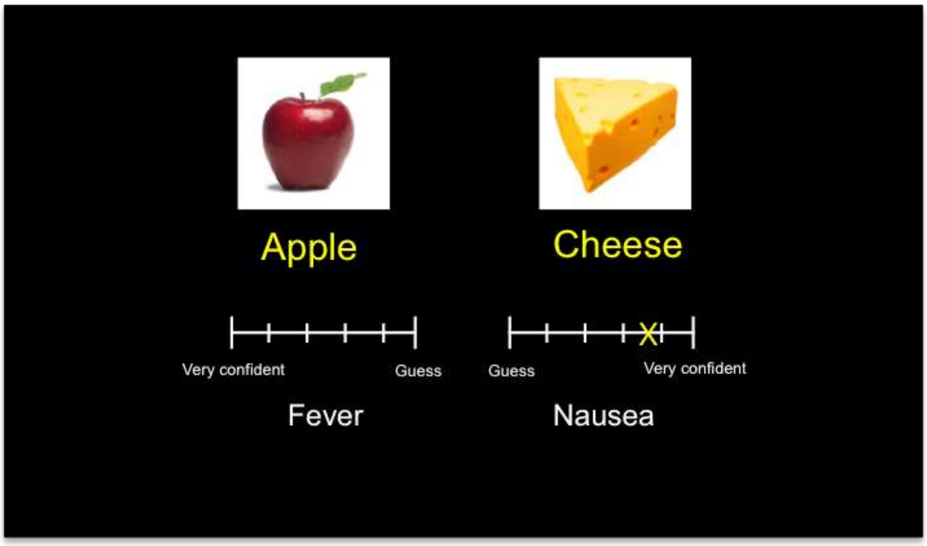

Figure 1. Example of stimulus presentation and response scales used in Stage 2 of each experiment. Participants used the scale to indicate both their outcome prediction and their confidence in that prediction.

Results

Figure 2 shows the prediction accuracy for each type of discrimination in Stage 1 averaged into four equal blocks of training. Each of these blocks thus contains four repeats of each trial type, which amounts to 32 trials when collapsed into Component versus Biconditional in Stage 1. Figure 3 displays prediction accuracy and confidence in Stage 2 averaged into four equal blocks for four within-subjects conditions. As noted above, in addition to outcome predictions, in Stage 2 we also measured confidence using continuous ratings, where ratings in the centre of the scale indicated lowest confidence and ratings at the left or right extremes indicated maximum confidence. We converted these confidence ratings, irrespective of accuracy, to a simple from 0-100. Accordingly, mean confidence ratings across Stage 2 are also shown in Figure 3.[1] Discriminations involving Stage 1 Predictive Component and Biconditional cues, where the Component is predictive in Stage 2 and discriminations involving Stage 1 Predictive Component and Biconditional cues, where the Biconditional is predictive in Stage 2, are shown in panels A (prediction accuracy) and C (confidence ratings). Discriminations involving Stage 1 Non-Predictive Component and Biconditional cues, where the Component is predictive in Stage 2, and discriminations involving Stage 1 Non-predictive Component and Biconditional cues, where the Biconditional is predictive in Stage 2 are shown in panels B (prediction accuracy) and D (confidence ratings). Each of these blocks contains three repeats of each trial type, which amounts to 12 trials for each of the four conditions.

Stage 1 prediction accuracy. Looking at Figure 2, and consistent with past results (e.g. Livesey et al., 2011; Saavedra, 1975), prediction accuracy in Stage 1 was consistently higher for the Component discriminations than for the Biconditional discriminations. A 2 x 4 repeated measures ANOVA with discrimination (Component vs Biconditional) and Block (1-4) as factors confirmed that this difference between the Stage 1 discriminations was significant, $F(1,49) = 28.34$, $p < .001$, $\eta_p^2 = .366$. In this and all subsequent ANOVA, we report the linear and quadratic trends across Block rather than the main effect, as these trends are more readily interpretable as evidence of improvement across trials. The linear and quadratic trends across Block were both highly significant, $F(1,49) = 432.37$, $p < .001$, $\eta_p^2 = .898$, $F(1,49) = 10.31$, $p = .002$, $\eta_p^2 = .174$, respectively. The quadratic trend in block also interacted with Discrimination, $F(1,49) = 7.67$, $p = .008$, $\eta_p^2 = .135$.

[1] We also used the combination of prediction accuracy and confidence to create a continuous learning score, as has been used in some previous studies. However, since the results from this learning score yielded essentially the same pattern as prediction accuracy alone, we have not reported them here.

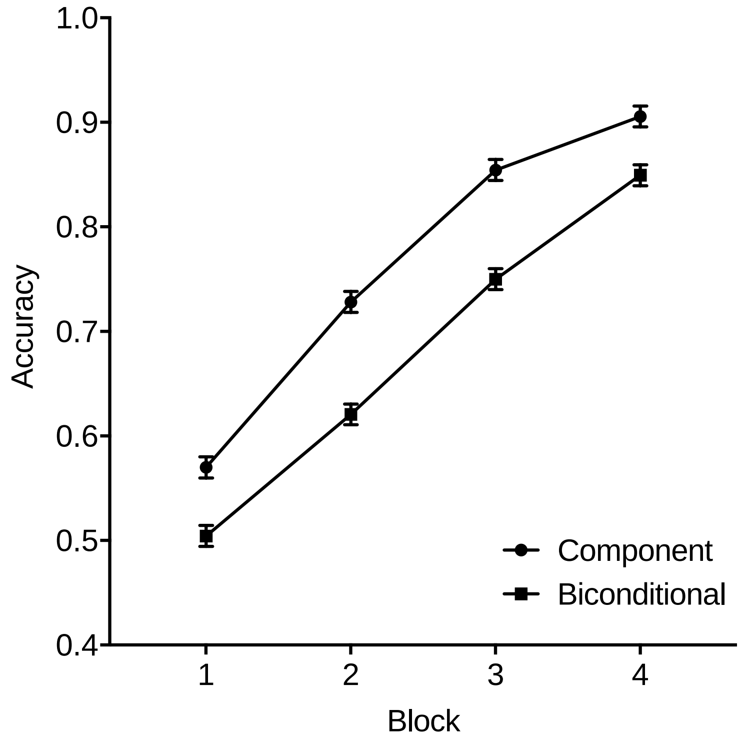

Figure 2. Prediction accuracy for component and biconditional discriminations across Stage 1 of Experiment 1, divided into four equal blocks. Within-subjects error bars were calculated here and in all subsequent figures following Cousineau (2005).

Stage 2 analysis. In Stage 2, a comparison between the discriminations in which Stage 1 biconditional cues are now predictive and the discriminations in which the Stage 1 component cues are now predictive provides the critical indication of transfer effects caused by the learning history of the cues in each type of discrimination. We will refer to this as an effect of Stage 1 Discrimination. A learned predictiveness effect manifests as positive transfer for previously predictive cues and/or negative transfer for previously non-predictive cues (e.g. Livesey & McLaren, 2007; Lochman & Wills, 2003). We thus have reason to expect an interaction between Stage 1 Discrimination and whether the component cue included in each Stage 2 compound was previously predictive or previously non-predictive. We will refer to this second factor as Stage 1 Component Predictiveness. Based on Livesey et al.'s (2011) results, we would expect to see an interaction between these two factors, and in particular, higher accuracy for the Stage 2 discrimination in which the Component cue was predictive in Stage 1 and predictive again in Stage

2. For both prediction accuracy and confidence, we tested this interaction using 2 x 2 x 4 repeated measures ANOVAs with Stage 1 Discrimination (Component vs Biconditional), Stage 1 Component Predictiveness (Predictive vs Non-predictive) and block (1-4) as factors.

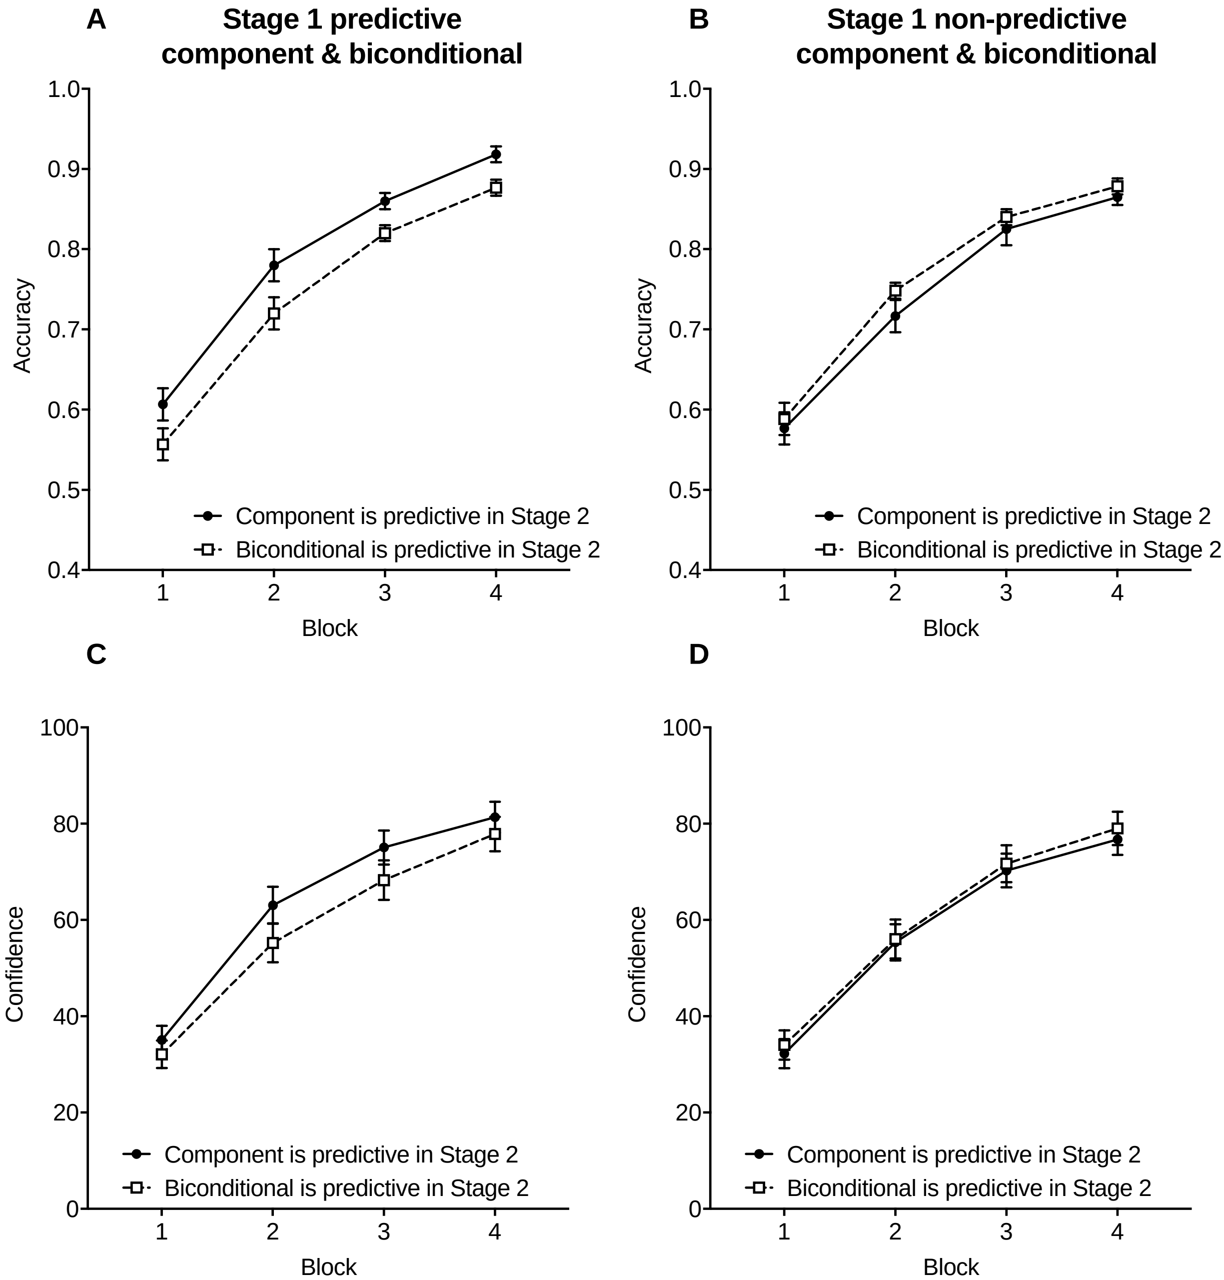

Figure 3. Prediction accuracy (panels A & B) and confidence (panels C & D) in Stage 2 of Experiment 1. Each panel shows one condition in which Stage 1 component cues were predicitve and one condition in which Stage 1 biconditional cues were predictive. Panels A & C show conditions containing a Stage 1 predictive component, panels B & D show conditions containing a Stage 1 non-predictive component. Error bars indicate within-subjects error (Cousineau, 2005).

Stage 2 prediction accuracy. Examining Figure 3A, accuracy was consistently higher for the discrimination in which the Component cue was predictive in both Stage 1 and Stage 2. There were significant linear and quadratic trends in block, $F(1,49) = 229.09$, $p < .001$, $\eta_p^2 = .824$, $F(1,49) = 32.21$, $p < .001$, $\eta_p^2 = .397$, respectively. These trends did not interact with the other factors, Fs < 1. Importantly, while neither the main effect of Stage 1 Discrimination nor Stage 1 Component Predictiveness were significant on their own, larger $F(1,49) = 1.34$, $p = .252$, there was an interaction between these factors, $F(1,49) = 5.86$, $p = .019$, $\eta_p^2 = .107$. This indicates that the relative difficulty of a discrimination in which the Stage 1 Component cue was predictive versus a discrimination in which the Stage 1 Biconditional cue was predictive depended on whether the Stage 1 Component cue had been predictive or non-predictive. Analysis of simple effects showed an effect of Stage 1 Discrimination when the Component cue was predictive, $F(1,49) = 5.89$, $p = .019$, $\eta_p^2 = .107$, indicating the Stage 2 discrimination was easier when the Stage 1 Component cue was predictive and the Stage 1 Biconditional cue was non-predictive. In contrast, when the Stage 1 component cue had been non-predictive there was no effect of Stage 1 Discrimination, $F(1,49) = 1.02$, $p = .318$, $\eta_p^2 = .02$. We ran a follow up Bayes Factor (BF) t-test using Rouder et al.'s (2009) suggested JZS prior with scaling factor r = .707 (all BFs reported below use these same assumptions), to compare mean accuracy for the two discriminations that included Stage 1 non-predictive component cues. This yielded a BF of 4.09 in favor of the null hypothesis, that the two discriminations were solved at the same rate.

Stage 2 confidence. Confidence followed the same pattern as accuracy. There were significant linear and quadratic trends in block, $F(1,49) = 176.77$, $p < .001$, $\eta_p^2 = .783$, $F(1,49) = 35.73$, $p < .001$, $\eta_p^2 = .422$, respectively, and interaction between Stage 1 Discrimination and Stage 1 Component Predictiveness, $F(1,49) = 5.36$, $p = .025$, $\eta_p^2 = .025$. There was an effect of Stage 1 Discrimination when Component was previously predictive, $F(1,49) = 5.93$, $p = .019$, $\eta_p^2 = .108$, but no difference when component was previously non-predictive, $F(1,49) = .83$, $p = .366$, BF = 4.39 in favor of the null.

Relationship with Stage 1 performance. As expected, there was a notable difference in Stage 1 accuracy for learning the component and biconditional discriminations. This raises the possibility that the advantage observed in Stage 2, when a previously predictive component cue again served as the predictive cue, is a direct consequence of this difference in mastering the Stage 1 task. For reasons that will become clearer in the context of Experiment 2, we were also interested in how the Stage 2 transfer effects related to *overall* performance in Stage 1. To examine these relationships, we computed a difference score that reflects the key finding from Stage 2; for each participant, we took the mean difference, collapsing across blocks, between the two lines illustrated in Figure 3A. That is, we took the mean accuracy for the Stage 2 discrimination in which the predictive cues had also been predictive in Stage 1, and subtracted mean accuracy for the Stage 2 discrimination in which the predictive cues had been biconditional cues in Stage 1 and the non-predictive cues had been predictive in Stage 1. We ran two simple liner regressions with this Stage 2 difference score as the dependent variable regressed against 1) the *difference* in mean accuracy for Stage 1 component and Stage 1 biconditional discriminations and 2) overall mean accuracy for all Stage 1 discriminations. Difference in accuracy for component and biconditional discriminations in Stage 1 did not predict the magnitude of the critical transfer effect, $F(1,48) < 0.1$, $R^2 < .001$, nor did overall mean accuracy, $F(1,48) = 0.948$, $p = .335$, $R^2 = .019$. It thus appears that the associability advantage observed here for previously predictive cues over biconditional cues is not closely related to overall Stage 1 performance.

Discussion

The result of main interest in Experiment 1 is the relative difficulties of the Stage 2 discriminations, especially those evident in Figure 3A. Accuracy was relatively high for the Stage 2 discrimination in which the predictive components had also served as predictive components in Stage 1, and the non-predictive components had served as biconditional cues in Stage 1. This discrimination generated significantly higher performance than one in which the functional role of these cues in Stage 2 was swapped (i.e. the predictive components in Stage 2 had been biconditional cues in Stage 1 and the non-predictive components had been predictive components in Stage 1). This result, which suggests higher associability for predictive components than for biconditional cues, was not apparent in the discriminations containing other Stage 1 biconditional cues paired with Stage 1 non-predictive components. For these discriminations, we found equivalent performance on the discrimination in which the Biconditional cues now served as predictive components relative to the discrimination in which the Biconditional cues now served as non-predictive components.

These results are consistent with those of Livesey et al. (2011, Experiment 2), which found a similar advantage in Stage 2 learning for Stage 1 predictive components trained with Stage 1 biconditional cues, but equivalent learning for Stage 1 non-predictive components trained with Stage 1 biconditional cues. However, whereas their result relied upon predictive ratings made at test in response to new combinations of cues, here we have demonstrated the effect in performance during learning. In doing so, we can rule out an explanation in terms of the participant generalising differently to new compounds at test depending on the functional roles of the cues in Stage 1. As noted in the introduction, learning the relational structure of the component and biconditional discriminations may provide an impetus to be more cautious or even to make completely different inferences about the recombination of biconditional cues compared to component cues, even though the Stage 2 discriminations are linearly solvable and can be learned via a simple summation of the associations of the cues with the outcomes.

The differences between discriminations in Stage 2 reveal transfer effects caused by the functional role that cues played in Stage 1. Specifically, what differed between the predictive and non-predictive cues was whether they had been used in a component or biconditional discrimination in Stage 1 and, if a component discrimination, whether they were predictive or non-predictive in

Stage 1. But in Stage 2, all four discriminations were linearly solvable component discriminations with two predictive cues and two non-predictive cues. This is important to note because it is still possible that learning the structure of the biconditional discrimination at a more relational level may hinder later learning about the cues on which that structure is based. For instance, learning that cue J is involved in a biconditional discrimination in Stage 1 and that its predictive properties are complex and conditional may hinder the individual's ability to use that cue as a predictive component in Stage 2. For this reason, it is important to examine whether relational learning is widespread in this task and whether we can find direct evidence that it affects subsequent learning about specific cues.

Thus while the main result from Experiment 1 strongly suggests differences in the rate of learning about previously predictive component cues compared to biconditional cues, it does not rule out the possibility that participants learn something about the relational structure of these discriminations, and that this affects the way they learn and make predictions during Stage 2 training. This possibility was investigated in the following two experiments.

Experiments 2 and 3 were designed to test for positive transfer based on the relational structure of the discriminations in Stage 1 rather than the associability of the individual cues. The two experiments tested complementary questions. First, are Stage 1 biconditional cues easier to learn about than Stage 1 component cues when used in separate biconditional discriminations in Stage 2? Second, are Stage 1 component cues (irrespective of their functional role as predictive or non-predictive) easier to learn about than Stage 1 biconditional cues when used in separate component discriminations in Stage 2? These questions are addressed by Experiments 2 and 3, respectively. In both experiments, Stage 1 cues were recombined into novel compounds in Stage 2. Two discriminations in Stage 2 were composed solely of Stage 1 biconditional cues, and two discriminations were composed solely of Stage 1 component cues, one of which contained predictive component cues, and the other non-predictive component cues.

Experiment 2

In Experiment 2, Stage 2 comprised four biconditional discriminations, as shown in Table 5. If participants have learned about the structural qualities of the solution that is necessary for the biconditional discrimination then we might observe positive transfer for the discriminations that contain Stage 1 biconditional cues. This learning may be complex and relational in a true sense, for instance an understanding that the cues are organised in groups of four and that the correct outcome is determined by the individual combinations within those groups, or it could be based on a more simple heuristic, for instance, "I can't assume that cheese eaten in one food combination will cause the same allergic reaction as cheese eaten in a different combination, I need to learn about the combinations".

Importantly, in this study, evidence of positive transfer in Stage 2 necessitates that structural learning is cue-specific to some degree. Other studies have examined the learning and use of structural knowledge of this form in patterning and biconditional discriminations (e.g. Cobos et al., 2017; Don et al., 2015, 2016; Harris & Livesey, 2008; Shanks & Darby, 1998; Wills et al., 2011a, 2011b), but usually in a context in which *all* cues in the initial stage of learning abide by (or at least are potentially consistent with) the underlying relational rules of the complex discrimination. Understanding of the structural rules is assessed by examining how well participants transfer these rules to new cues that may or may not relate to one another in the same way. Thus there is an assumption that such learning is abstracted and does not rely on the specific cues in question. In these experiments, we are testing for cue-specific transfer, which may differentially affect how Stage 1 biconditional and Stage 1 component cues are learned about. While it is clear from previous studies that healthy adults are capable of learning transferrable relational knowledge about the structure of complex discriminations, in this case any relational knowledge that is completely abstract would be equally applicable to all cues in Stage 2 and thus transfer effects would not be observed because each of the three classes of biconditional discrimination would benefit equally from that knowledge.

Table 5

Stage 1 and Stage 2 training contingencies used in Experiment 2.

Stage 1			Stage 2	
Component	Biconditional		Component to biconditional	Biconditional to biconditional
AW – O1	JN – O1	*(predictive*	AB – O3	JK – O3
AX – O1	JO – O2	*components)*	AD – O4	JM – O4
BW – O2	KN – O2		CB – O4	LK – O4
BX – O2	KO – O1		CD – O3	LM – O3
CY – O1	LP – O1	*(non-predictive*	WX – O4	NO – O4
CZ – O1	LQ – O2	*components)*	WZ – O3	NQ – O3
DY – O2	MP – O2		YX – O3	PO – O3
DZ – O2	MQ – O1		YZ – O4	PQ – O4

Note: Letters refer to individual cues, A-D: predictive components, W-Z: non-predictive components, J-Q: equally relevant stimuli from biconditional discrimination. O1, O2, O3, and O4 refer to four neutral outcome

Regardless of the potential for this type of learning and transfer, we can assume that individual cue associability will be different for predictive, non-predictive and biconditional cues. However, because the cues are recombined in independent discriminations, these differences do not predict transfer effects in the same way as in Experiment 1. It is possible that cues of higher associability (e.g. Stage 1 predictive component) may be faster to learn about than cues of lower associability (e.g. Stage 1 non-predictive component cues) but since the individual cues do not predict a specific outcome in Stage 2, this will not necessarily benefit acquisition or performance and, indeed, could have the opposite effect if attention to individual cues interferes with learning about cue compounds. Importantly, any direct influence of individual cue associability on the learning of the biconditional discrimination should be evident in a difference in performance between the Stage 2 discriminations using Stage 1 predictive versus Stage 1 non-predictive cues. If these two discriminations yield equal performance but learning is faster for the discrimination using Stage 1 biconditional cues then it suggests positive transfer based on the learning of something other than simple cue predictiveness.

Work on relational rule discovery in other learning paradigms has clearly demonstrated that not all individuals identify abstract relations or transfer them to new situations to the same degree (e.g. Goldwater et al., 2018; McDaniel et al., 2014). Indeed, Shanks & Darby (1998) found that efficiency in learning patterning discriminations during initial training predicted the transfer of the structural patterning rule to new cues. Efficient learners, who performed well during training showed much stronger evidence of rule transfer. Other studies using the Shanks-Darby patterning task have found that the learning of relational patterning rules is associated with higher working memory capacity, stronger cognitive reflection, and greater model-based choice on instrumental decision tasks (Don et al., 2015, 2016; Wills et al., 2011a, 2011b). The structural transfer effect that we tested for in this experiment may have different properties since it is necessarily cue-specific rather than completely abstract. Nevertheless an effect may be confined to a subset of individuals who engage in the task differently (or alternatively, simply learn faster) than others. To investigate this, we ran a simple regression testing for the relationship between Stage 2 transfer and Stage 1 accuracy, similar to that performed in Experiment 1, but for illustrative purposes, we also examined groups divided by a median split based on performance in Stage 1.

Method

Participants. For each of Experiments 2 and 3, we continued testing until we had data from 52 usable participants. In Experiment 2, a total of 62 undergraduate students (31 from the University of Marburg and 31 from the University of Sydney) participated in the experiment. Ten participants failed to satisfy the Stage 1 learning criteria used in the previous experiments (60%

accuracy in the final quarter for both component and biconditional discriminations). All analyses were run on the resulting sample of 52 participants (34 female, mean age = 20.9 years).

Stimuli, Apparatus, Design and Procedure. Experiment 2 used the same stimuli and apparatus as Experiment 1. The design was the same with the exception that cues were grouped differently in Stage 2. Following Table 5, four biconditional discriminations were constructed by recombining either Stage 1 biconditional cues or Stage 1 component cues. For the recombined component cues, one discrimination was composed of previously predictive component cues and the other was composed of previously non-predictive component cues. Participants completed the experiment following the same procedure used in Experiment 1.

Results

Stage 1 prediction accuracy. Looking at Figure 4A, prediction accuracy in Stage 1 was again consistently higher for the Component discriminations than for the Biconditional discriminations. A 2 x 4 repeated measures ANOVA with Discrimination (Component vs Biconditional) and Block (1-4) as factors confirmed this, with a significant main effect of Discrimination $F(1,51) = 68.28$, $p < .001$, $\eta_p^2 = .572$. The linear and quadratic trends across Block were both highly significant, $F(1,51) = 394.69$, $p < .001$, $\eta_p^2 = .886$, $F(1,49) = 8.48$, $p = .005$, $\eta_p^2 = .143$. The quadratic trend in block also interacted with Discrimination, $F(1,49) = 5.14$, $p = .028$, $\eta_p^2 = .092$.

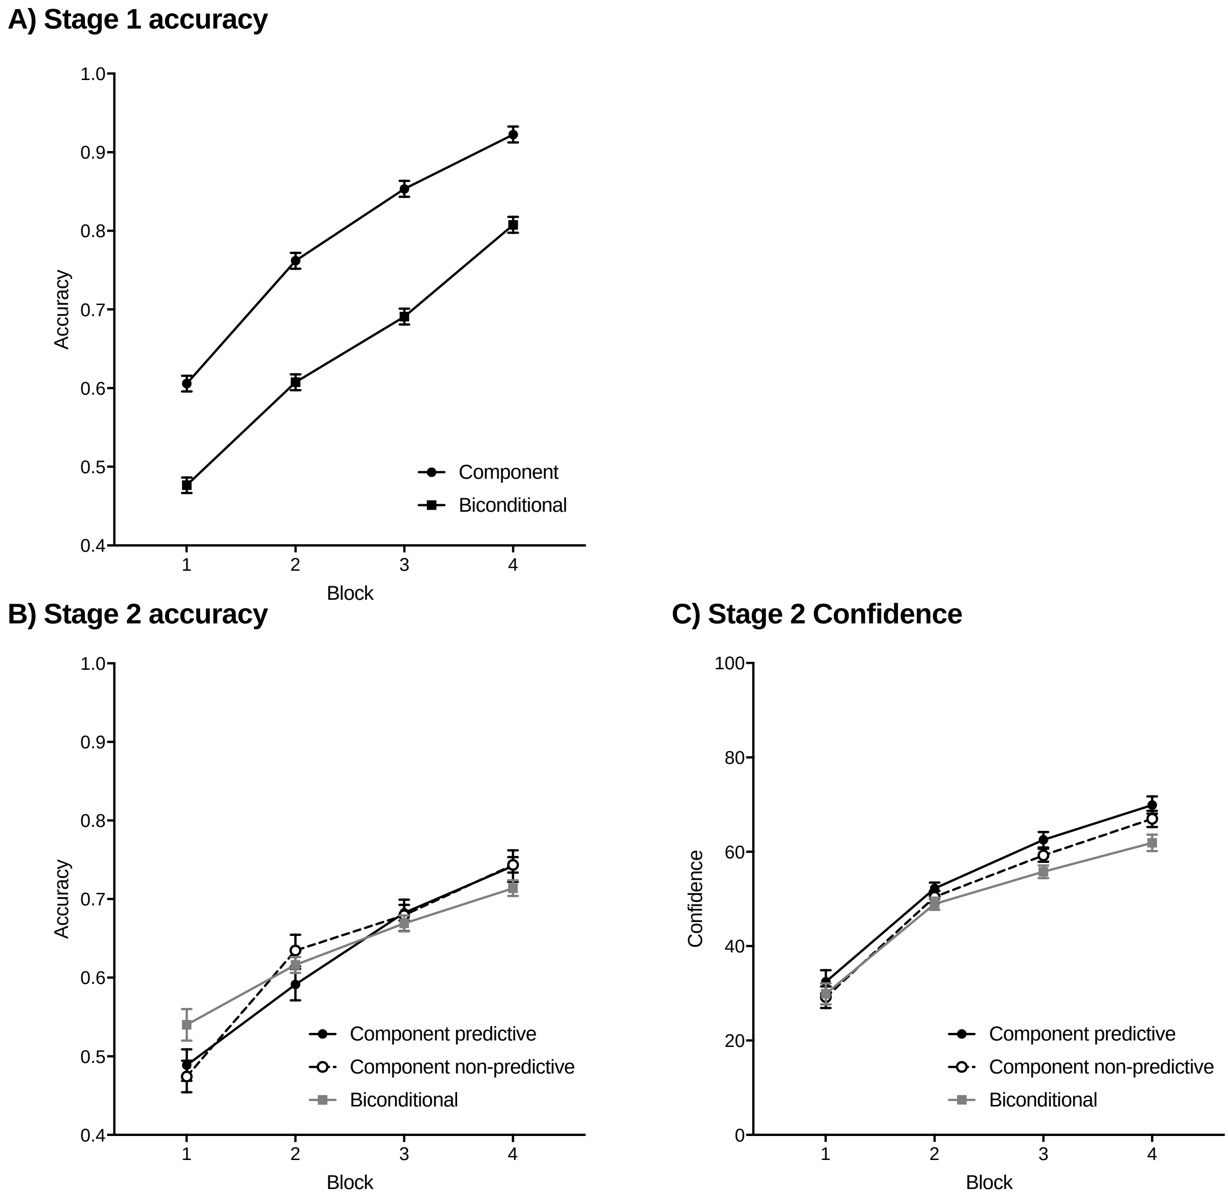

Figure 4. Stage 1 prediction accuracy (panel A), Stage 2 prediction accuracy (Panel B) and confidence (Panel C) in Experiment 2. Error bars indicate within-subjects error (Cousineau, 2005).

Stage 2 analysis. In Stage 2, there were three classes of discrimination defined in terms of the functional role that the cues had played in Stage 1 (predictive component cues, non-predictive component cues, and biconditional cues). Analyses were organised to test the hypothesis that a biconditional discrimination in Stage 2 would be easier to learn if based on cues that were previously involved in a biconditional discrimination in Stage 1, but also to test whether the individual cue associability changes led to differences in the difficulty of learning a biconditional discrimination. We constructed two orthogonal contrasts for Stage 2 discrimination, the first comparing performance on Stage 1 biconditional to the other two, the second comparing Stage 1 Predictive component to Stage 1 non-predictive component. These contrasts were applied to each of the Stage 2 dependent variables.

Stage 2 prediction accuracy. Examining Figure 4B, accuracy was initially higher for the biconditional discrimination comprising cues that had also served as biconditional cues in Stage 1, however this effect disappeared relatively quickly. There were significant linear and quadratic trends in Block, $F(1,51) = 132.45$, $p < .001$, $\eta_p^2 = .722$, and $F(1,51) = 10.43$, $p = .002$, $\eta_p^2 = .17$, respectively. The contrast comparing Stage 1 Biconditional cues to the Stage 1 Predictive and Non-predictive component cues significantly interacted with linear trend in block, $F(1,51) = 10.75$, $p = .002$, $\eta_p^2 = .174$, reflecting the pattern in Figure 4B in which Stage 1 Biconditional discrimination was initially learned about faster in Block 1 before the other discriminations caught up later in training. Indeed, follow-up block by block contrast analyses indicated that Stage 1 Biconditional accuracy was significantly higher in Block 1, $F(1,51) = 7.96$, $p = .007$, $\eta_p^2 = .135$, BF = 5.16 in favor of the alternative hypothesis that the conditions differ, but did not differ from the others in any other block, *F*s <= 1.53, *p*s >= .222, BFs > 3.22 in favor of the null. The second contrast comparing Stage 1 Predictive and Stage 1 Non-predictive did not interact with the linear trend in block, $F < 1$. When averaged across blocks, neither of the contrasts on Stage 1 Discrimination approached statistical significance, *F*s < 1. Taken together, these results suggest that there was an early advantage for learning the new biconditional discrimination if the cues had been involved in a biconditional discrimination in Stage 1, but that this advantage was transient, and not seen throughout the middle and later stages of training.

Stage 2 confidence. Confidence ratings are shown in Figure 4C and show a subtly different pattern. Confidence in the Stage 1 Biconditional discrimination was slightly lower than the other two, despite accuracy being initially better on these trials. Again, linear and quadratic trends across Block were significant, $F(1,51) = 99.37$, $p < .001$, $\eta_p^2 = .661$, and $F(1,51) = 22.41$, $p < .001$, $\eta_p^2 = .305$, respectively. This time, however, there was a significant contrast comparing Stage 1 Biconditional to the other two, $F(1,51) = 6.48$, $p = .014$, $\eta_p^2 = .113$, indicating that confidence in the Stage 2 discriminations containing biconditional cues was generally *lower*. The contrast comparing Stage 1 Predictive to Stage 1 Non-predictive cues was not significant, $F(1,51) = 2.61$, $p = .112$, $\eta_p^2 = .049$, BF = 1.97 in favor of the null.

The contrast comparing Stage 1 Biconditional to the other two discriminations interacted with linear trend in block, $F(1,51) = 6.30$, $p = .015$, $\eta_p^2 = .11$, this time indicating that the reduced confidence for Stage 1 Biconditional relative to the other discriminations actually increased across Stage 2 training. The contrast comparing Stage 1 Predictive and Stage 1 Non-predictive did not interact with block, $F < 1$.

Relationship with Stage 1 performance. We were again interested in whether the critical transfer effects observed in Stage 2 were related to overall performance in Stage 1, or the differences in performance on component and biconditional discriminations in Stage 1. Following the analyses conducted on Experiment 1, we ran two simple liner regressions with a Stage 2 difference score as the dependent variable regressed against the difference in mean accuracy for Stage 1 component and Stage 1 biconditional discriminations and overall mean accuracy for all Stage 1 discriminations. The critical transfer effect in Stage 2 was the contrast reflecting the difference between Stage 1 biconditional, and the mean of Stage 1 predictive and Stage 1 non-predictive conditions, which was strongest in Block 1. The Stage 2 difference score was therefore calculated as this difference in accuracy for Stage 1 biconditional and Stage 1 component cues in the first block of Stage 2. The difference in accuracy for component and biconditional discriminations in Stage 1 did not predict the magnitude of this difference score, $F(1,50) = 2.28$, $p = .137$, $R^2 = .044$. The effect of overall Stage 1 mean accuracy was marginal, though still falling short of statistical significance, $F(1,50) = 3.79$, $p = .057$, $R^2 = .071$, $\beta = .266$, providing tentative evidence that mastery of the Stage 1 discriminations may be positively related to the magnitude of the transfer effect observed at the beginning of Stage 2. Figure 5 displays the Stage 2 data from Experiment 2 split into two groups based on Stage 1 accuracy (median = 72.66%); a Low Stage 1 Performance group who performed at or worse than the median (n=28) and a High Stage 1 Performance group who performed above the median (n=24). The critical interaction effect was strongly significant in the High Stage 1 Performance group, $F(1,23) = 21.88$, $p < .001$, $\eta_p^2 = .488$, and did not approach significance in the Low Stage 1 Performance group, $F < 1$.

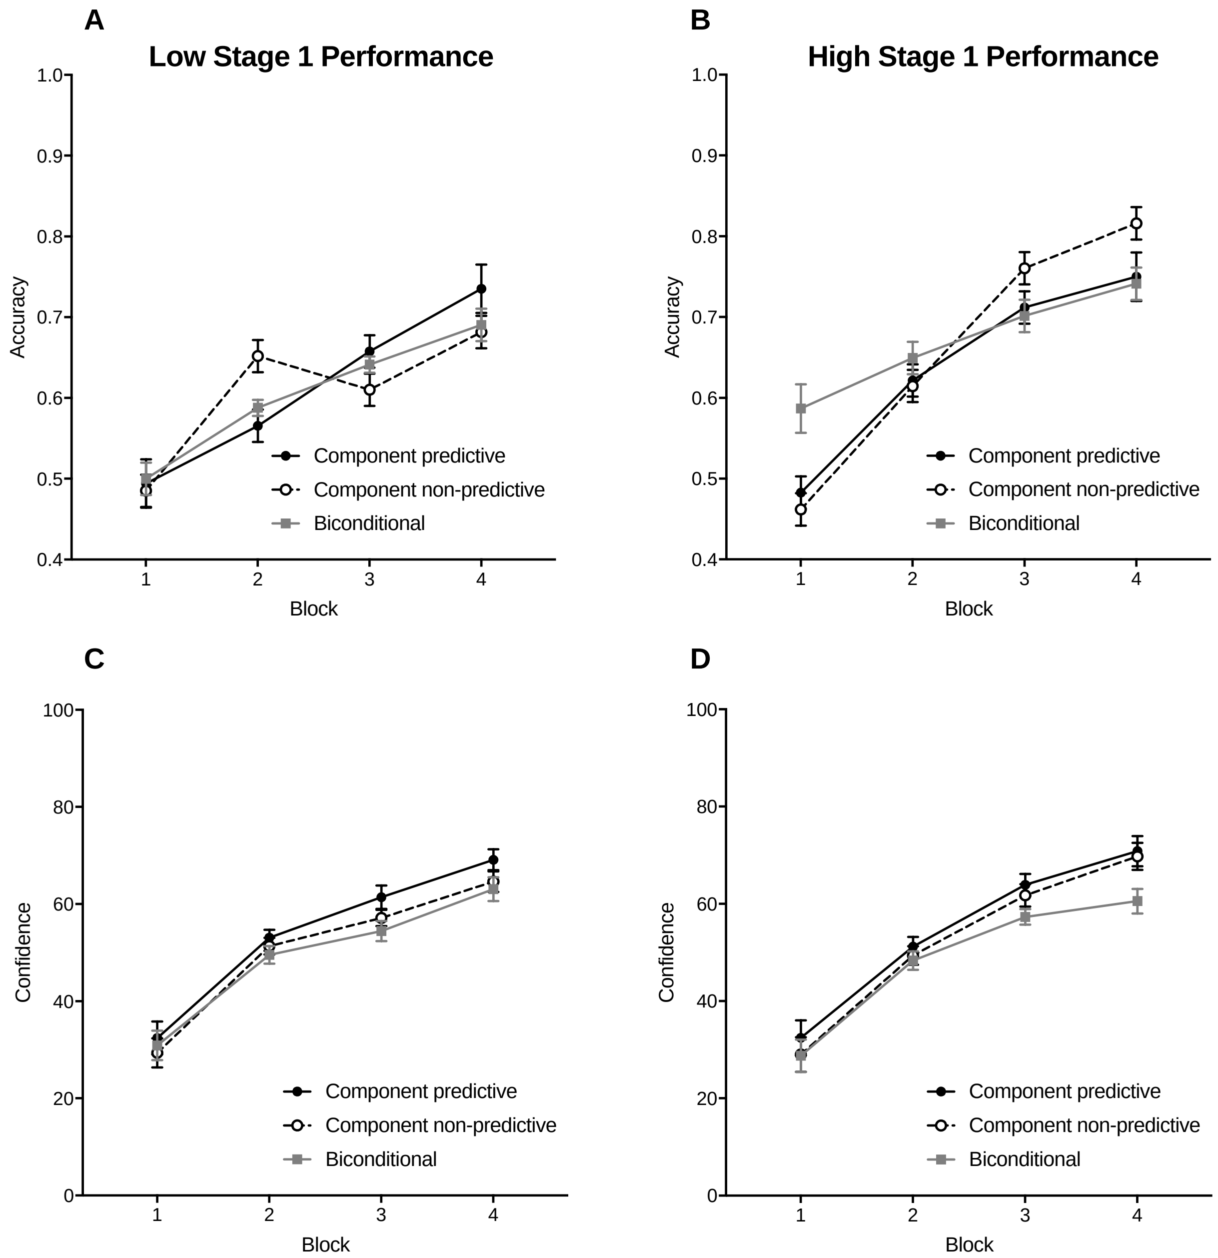

Figure 5. Stage 2 results for Experiment 2 split according to overall performance (mean prediction accuracy) in Stage 1. Error bars indicate within-subjects error (Cousineau, 2005).

Discussion

Experiment 2 revealed a cue-specific positive transfer effect between biconditional discriminations used in the first and second training stage. Performance on a biconditional discrimination in Stage 2 was enhanced when the cues were involved in a biconditional discrimination in Stage 1, relative to biconditional discriminations involving previously predictive or previously non-predictive cues that were involved in linearly solvable component discriminations in Stage 1. To the best of our knowledge, this is the only demonstration of a within-participant (and thus completely cue-specific) structural transfer effect of this nature. The positive transfer effect was transient, disappearing early in training, and was not accompanied by stronger confidence in the Stage 1 Biconditional discrimination relative to the other two.

In contrast, the Stage 2 biconditional discriminations comprising predictive component cues from Stage 1 were learned at the same rate as discriminations comprising non-predictive component cues from Stage 1. Thus, if these discriminations suffered from negative transfer then they did so at the same rate. If these cues acquired different properties as a consequence of their roles in Stage 1, for instance enhanced associability for predictive over non-predictive cues, then this had a neglible effect overall on participants' ability to learn about the cues in a biconditional discrimination.

Like other relational learning effects, it appears that the key structural transfer result from Experiment 2 was produced entirely by a subset of participants, namely the stronger learners who performed relatively well during Stage 1. We will return to the possible cause and implication of these results in the General Discussion. First, however, Experiment 3 provides a complement to this experiment, in which cues were transferred to linearly solvable component discriminations in Stage 2.

Experiment 3

Experiment 3 provided a similar test to Experiment 2 but this time involved recombining the cues for use in component discriminations rather than biconditional. This design differed from that used in Experiment 1 in that rather than putting cues from different Stage 1 discriminations in direct competition for learning, they were grouped according to their functional role in Stage 1, just as they had been in Experiment 2. This is shown in Table 6. We used four possible recombinations of the Stage 1 predictive component cues. This is because the consistent associations of these cues with particular Stage 1 outcomes (i.e. O1 and O2) can create the potential for outcome equivalence effects (e.g. see Le Pelley & McLaren, 2003).

Table 6

Stage 1 and Stage 2 training contingencies used in Experiment 3.

Stage 1		Stage 2				
Component	Biconditional	Component to component				Biconditional to component
		(four possible recombinations for the predictive components)				
AW – O1	JN – O1	AC – O3	AC – O4	AB – O3	AB – O4	JK – O3
AX – O1	JO – O2	AD – O3	AD – O3	AD – O3	AD – O3	JM – O3
BW – O2	KN – O2	BC – O4	BC – O4	CB – O4	CB – O4	LK – O4
BX – O2	KO – O1	BD – O4	BD – O3	CD – O4	CD – O3	LM – O4
CY – O1	LP – O1				WX – O3	NO – O4
CZ – O1	LQ – O2				WZ – O3	NQ – O4
DY – O2	MP – O2				YX – O4	PO – O3
DZ – O2	MQ – O1				YZ – O4	PQ – O3

Note: Letters refer to individual cues, A-D: predictive components, W-Z: non-predictive components, J-Q: equally relevant stimuli from biconditional discrimination. O1, O2, O3, and O4 refer to four neutral outcomes.

Method

Participants. A total of 68 undergraduate students (29 from the University of Marburg and 39 from the University of Sydney) participated in the experiment. 16 participants failed to satisfy the Stage 1 learning criteria used in the previous experiments (60% accuracy in the final quarter for both component and biconditional discriminations). Testing continued until each of four counterbalancing conditions contained 13 participants who satisfied the learning criteria. All analyses were run on the resulting sample of 52 participants (38 female, mean age = 21.0 years).

Stimuli, Apparatus, Design and Procedure. Experiment 3 used the same stimuli and apparatus as Experiments 1 and 2. The design was the same with the exception that cues and outcomes were again grouped differently in Stage 2. Following Table 6, four component discriminations were constructed by recombining cues in the same groups of four cues used in Experiment 2. Thus one discrimination was composed of previously predictive component cues, one composed of previously non-predictive component cues, and two composed of previously biconditional cues. Four different recombinations of the predictive cues were used, counterbalanced between participants.[2] The procedure was the same as that used in Experiments 1 and 2.

Results

Stage 1 prediction accuracy. Looking at Figure 6A, accuracy in Stage 1 was again consistently higher for the Component discriminations than for the Biconditional discriminations, with a significant main effect of Stage 1 Discrimination, $F(1,51) = 86.07$, $p < .001$, $\eta_p^2 = .628$. There was also a significant linear and quadratic trends in Block, $F(1,51) = 486.88$, $p < .001$, $\eta_p^2 = .905$, $F(1,51) = 19.32$, $p < .001$, $\eta_p^2 = .275$, respectively. The linear and quadratic trends in block both interacted with Stage 1 Discrimination, $F(1,51) = 12.62$, $p = .001$, $\eta_p^2 = .198$, $F(1,51) = 4.53$, $p = .038$, $\eta_p^2 = .082$, respectively.

Stage 2 analysis. The analyses of Stage 2 data in Experiment 3 were conducted in an identical fashion to those in Experiment 2.

[2] These counterbalancing conditions fall into functionally equivalent pairs, two in which predictors of different outcomes in Stage 1 again predict different outcomes in Stage 2, and two in which predictors of the same outcome in Stage 1 predict different outcomes in Stage 2. The former pair have a higher potential for outcome equivalence effects. Therefore, it is worth noting that none of the effects of interest in Stage 2 learning interacted with counterbalancing condition when it was added as an additional between-subjects variable (either as a factor with 4 levels or a factor with two levels, grouping for this functional equivalence).

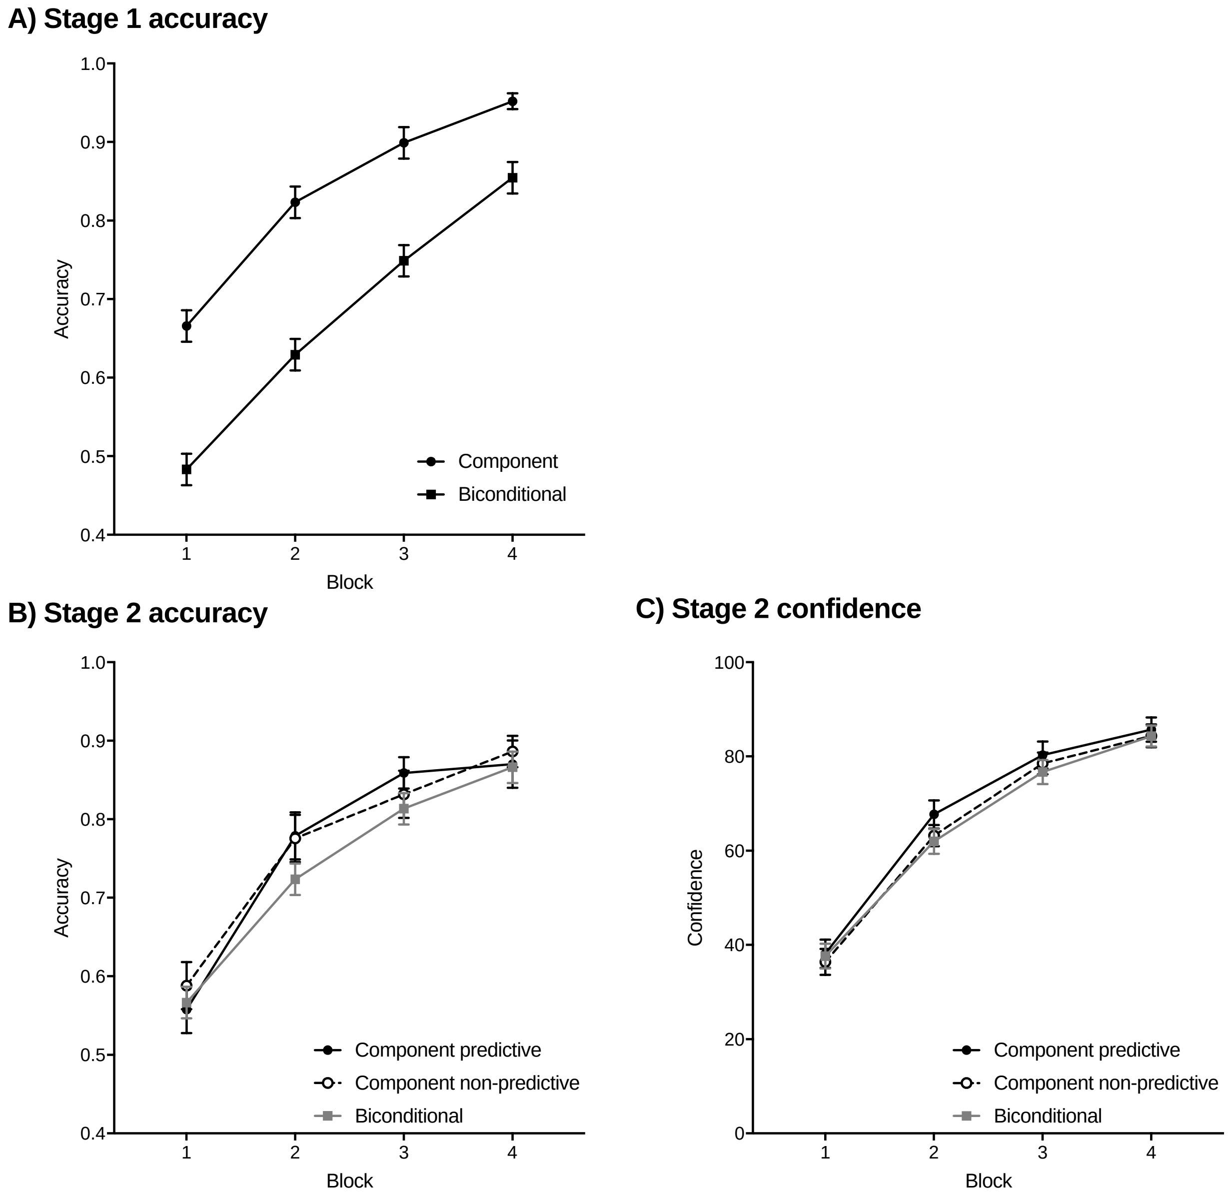

Figure 6. Stage 1 prediction accuracy (Panel A), Stage 2 prediction accuracy (Panel B) and confidence (Panel C) in Experiment 3. Error bars indicate within-subjects error (Cousineau, 2005).

Stage 2 prediction accuracy. Examining Figure 6B, Stage 2 accuracy for the different Stage 1 Discrimination conditions began and ended at roughly the same levels of performance, however there appeared to be at least some reduction in performance for the Stage 1 Biconditional condition relative to the other two through the middle of Stage 2, which may indicate a slower rate of acquisition. There were significant linear and quadratic trends in Block, $F(1,51) = 197.05$, $p < .001$, $\eta_p^2 = .794$, and $F(1,51) = 27.04$, $p < .001$, $\eta_p^2 = .346$, respectively. The contrast comparing

Stage 1 Biconditional cues to the Stage 1 Predictive and Non-predictive component cues significantly interacted with quadratic trend in block, $F(1,51) = 4.832$, $p = .033$, $\eta_p^2 = .087$, reflecting the pattern in Figure 6B in which performance on Stage 1 Biconditional discrimination was weaker in the middle blocks of Stage 2 but not the beginning or end. Follow-up contrast analyses taking an average of performance over the middle two blocks indicated that Stage 1 Biconditional accuracy was significantly lower, $F(1,51) = 5.52$, $p = .023$, $\eta_p^2 = .098$, BF = 1.84 in favor of the alternative hypothesis, but did not differ from the others in the first and last blocks, $F < 1$, BF = 5.47 in favor of the null. The second contrast comparing Stage 1 Predictive and Stage 1 Non-predictive also produced a marginal but non-significant interaction with this quadratic trend in block, $F(1,51) = 3.03$, $p = .088$, $\eta_p^2 = .056$. Neither of the Stage 1 Discrimination contrasts interacted with linear trend in block, Fs < 1. When averaged across blocks, the contrast comparing Stage 1 Biconditional to the other two was marginal but non-significant, $F(1,51) = 3.18$, $p = .08$, $\eta_p^2 = .059$, while the second contrast comparing Stage 1 Predictive and Stage 1 Non-predictive did not approach significance, $F < 1$. Taken together, these results suggest that there was at least some disadvantage for learning the new linearly solvable component discrimination in Stage 2 if the cues had been involved in a biconditional discrimination in Stage 1. This effect was transient, and only seen through the middle stages of training.

Stage 2 confidence. Confidence ratings are shown in Figure 6C and appear to be very similar for the different Stage 1 Discrimination conditions. There were significant linear and quadratic trends for Block, $F(1,51) = 152.67$, $p < .001$, $\eta_p^2 = .750$ and $F(1,51) = 49.40$, $p < .001$, $\eta_p^2 = .492$ respectively. However, these did not interact with either of the Stage 1 Discrimination contrasts, largest $F(1,51) = 2.54$, $p = .117$, $\eta_p^2 = .047$ for the interaction between Contrast 1 (Stage 1 Biconditional vs others) and quadratic trend in Block. The Stage 1 Discrimination contrasts averaging across blocks were also non-significant, larger $F(1,51) = 2.59$, $p = .113$, $\eta_p^2 = .048$ for Contrast 2 (Stage 1 Predictive vs Stage 1 Non-predictive).

Relationship with Stage 1 performance. As with the previous two experiments, we ran two simple liner regressions of a Stage 2 difference score based on the critical transfer effect regressed against the difference in mean accuracy for Stage 1 component and Stage 1 biconditional discriminations and overall mean accuracy for all Stage 1 discriminations. For Experiment 3, we calculated the difference score for the critical Stage 2 transfer effect using accuracy in the middle two blocks of Stage 2 and took the difference between the Stage 1 biconditional condition and the mean of the Stage 1 predictive and Stage 1 non-predictive conditions (i.e. again selecting the contrast that revealed strongest evidence of a transfer effect). Difference in accuracy for component and biconditional discriminations in Stage 1 did not predict the magnitude of the difference score, $F(1,50) < .01$, $R^2 < .001$, nor did overall mean accuracy, $F(1,50) = 1.38$, $p = .246$, $R^2 = .027$. Therefore, similar to Experiment 1, it appears that the associability advantage observed here for previously predictive cues over biconditional cues is not closely related to Stage 1 performance.

Discussion

As in Experiment 2, we found evidence that cues involved in a Biconditional discrimination were treated differently during Stage 2 learning compared to cues involved in a linearly solvable component discrimination. This time, accuracy for the Stage 1 Biconditional condition was slightly worse during the middle two blocks of Stage 2 training. This negative transfer effect suggests at least some (albeit fairly transient) transfer of learning the structural properties of the discriminations in Stage 1.

General Discussion

This study confirmed that stimuli which have previously served as predictive cues in a linearly solvable component discrimination appear to gain greater associability than stimuli which have previously served as relevant cues in a biconditional discrimination. However, we also found novel evidence that knowledge about the relational structure of these learning tasks influences subsequent learning, in a way that might explain some of this difference in cue associability. As far as we know, this result is novel in that it demonstrates cue-specific learning of structural relations in a within-subjects design. Other studies concerning relational discovery in this form of associative learning task have examined the propensity for such effects to be learned in abstract form and thus transferrable to other stimuli. We will first summarise the key findings before discussing their implications.

In each of three experiments, participants were initially trained with the same Stage 1 learning tasks, comprising component discriminations, for which one set of cues was predictive of trial outcome and another set of cues was non-predictive, as well as biconditional discriminations for which trial outcome was signalled by the particular configurations of cues (all cues are relevant but none are predictive in isolation). In Experiment 1, we arranged new Stage 2 component discriminations composed of a cue that had previously served as a predictive component in Stage 1 and a cue that had previously served as a biconditional cue in Stage 1. Acquisition proceeded faster when the Stage 2 predictive components had also been predictive components in Stage 1 compared to the Stage 2 discrimination in which predictive components had been biconditional cues in Stage 1. In contrast, when Stage 2 components were composed of a cue that had previously served as a non-predictive component in Stage 1 and a cue that had previously served as a biconditional cue in Stage 1, learning appeared to be unaffected by the prior functional role of the cues. These discriminations were learned at an equal rate regardless of whether the predictive components were Stage 1 biconditional cues or Stage 1 non-predictive components.

In Experiment 2, Stage 1 training was followed by several biconditional discriminations in Stage 2, each discrimination organised using cues that were solely from one functional class from Stage 1 (i.e. all predictive components, all non-predictive components, or all biconditional cues). In Stage 2, acquisition occurred at the same rate for a biconditional discrimination composed of Stage 1 predictive cues and a biconditional discrimination composed of Stage 1 non-predictive cues. However, relative to these discriminations, we observed facilitation in Stage 2 for biconditional discriminations composed of Stage 1 biconditional cues. This transfer effect occurred in the initial blocks of Stage 2 training and post hoc analyses suggested that those individuals with higher accuracy in Stage 1 were primarily responsible for producing it. Experiment 3 used a conceptually very similar design except that participants received linearly solvable component discriminations in Stage 2 rather than biconditional discriminations. Like Experiment 2, in Stage 2, a component discrimination comprising previously predictive component cues was learned at the same rate as a component discrimination comprising previously non-predictive component cues. However, relative to these conditions, acquisition was slightly impaired when the component discrimination was created by recombining cues from the Stage 1 biconditional discriminations. This subtle transfer effect manifested in the middle of Stage 2 training and was not related to Stage 1 performance. Thus although it is complementary to the transfer effect observed in Experiment 2, its characteristics were different.

What do these results mean for changes in cue associability and their transfer to the learning of new associative relationships? The conflicting observations that motivated this study were 1) that learning a biconditional discrimination results in relatively low cue associability in subsequent learning, equivalent to that afforded to non-predictive component cues and less than that of predictive component cues (Livesey et al., 2011), and 2) that learning a biconditional discrimination results in relatively high associability of the relevant stimulus dimension, equivalent to that afforded to the predictive dimension in a component discrimination and more than that of a non-predictive dimension (Uengoer & Lachnit, 2012). There were substantial differences in the way these two prior studies tested for associability changes, Livesey et al. using a further test phase to separate the different types of cues that competed for association in Stage 2, Uengoer and Lachnit using relative difficulty of different Stage 2 discriminations to draw inferences about associability. The results of Experiment 1 were consistent with the Livesey et al. finding despite using relative difficulty of Stage 2 as the key indicator of associability and thus this experiment confirms that the discrepancy in the results is not due to these different test methods. This is important because the test phase used by Livesey et al. assessed learning about individual cues by testing generalisation to new compounds, and thus cue-specific differences in generalization rather than associability could have produced an apparent deficit in learning about cues previously used in a biconditional discrimination.

Although cue associability appeared to be relatively low for previously non-predictive component cues and previous biconditional cues, the reasons for this low associability may not be the same for both classes of cue. Experiments 2 and 3 suggest that other properties of the task such as structural aspects of the discriminations are also learned and may affect subsequent learning about the cues involved in those discriminations. The experiments suggest that, relative to cues that are predictive or non-predictive in a straightforward way, a cue that has been involved in a biconditional discrimination may be easier to learn about in subsequent biconditional discriminations and harder to learn about in linearly solvable component discriminations. While these effects were relatively fleeting compared to the consistent difference in cue associability observed in Experiment 1, they provide evidence suggestive of processing changes that are qualitatively different from those described by Mackintosh (1975) and other theories based on learned predictiveness principles.

Recent literature has focused on the effects of uncertainty as a causal factor in generating enhanced attention for partially- or even non-predictive cues (e.g. Beesley et al., 2015; Griffiths, Johnson & Mitchell, 2011). The rationale, often invoked in conjunction with learning models in which attention to cues is proportional to prediction-error (Pearce & Hall, 1980), is that explorative learning processes should direct the individual's attention to cue-outcome relationships that still need to be learned rather than simply relying on exploiting those that are already learned well. Although the cues presented in the biconditional discrimination do not convey objective uncertainty about the outcome (the outcome is completely predictable based on the configuration of the cues), performance was less than perfect and was worse on biconditional discrimination than on component discriminations in Stage 1. To the extent that prediction accuracy is proportional to the prediction error experienced by participants during Stage 1, participants would have experienced greater prediction error on biconditional discrimination trials than on the component discrimination trials. Could this difference have caused relative impairment for Stage 1 biconditional cues when used in Stage 2? For a number of reasons, we do not favour an explanation of these results in these terms. First, confidence ratings at the beginning of Stage 2 did not reveal any differences across any experiments. Second, in Experiments 1 and 3, any effect caused by reduced certainty about Stage 1 biconditional cues that transfers to Stage 2 would have to have an effect on attention that runs *opposite* to the conventional prediction of such models.

Two other proposed changes in cue processing are relevant to the results from these experiments. First, relational learning of the structural properties of a discrimination is well documented for other nonlinear problems, particularly positive and negative patterning. As noted, Shanks and Darby (1998) and several others since have found that a subset of generally high-performing participants identified an abstract "opposites" rule that efficiently describes the structural relations involved in patterning discriminations and generalised this rule to new predictive cues in a way that dramatically changed their predictions on test. It is possible that similar structural characteristics are learned by some individuals when completing a biconditional discrimination. Humans and other animals typically find biconditional discriminations more difficult to acquire than patterning discriminations (Harris, et al., 2008; Harris & Livesey, 2008). The relations between cues in a biconditional discrimination are more complex and difficult to describe via a simple heuristic like the opposites rule (see Baetu et al., 2018). Thus it is debatable whether learning the relations involved in the structure of the biconditional task is as useful as in patterning, where it may be advantageous for overcoming the deleterious effects of summation on negative patterning, in particular (e.g. see Livesey, Thorwart & Harris, 2011; Thorwart et al., 2017).

Nevertheless, it is still possible that such learning occurs and that it could affect future processing of those cues in new learning tasks. It is worth noting that the evidence for these relational transfer effects comes from tasks where there are no contingencies that directly conflict with the relations in question. The current experiments differ in that the relations are specific to a subset of discriminations and, in fact, only a cue-specific transfer of such learning would result in differential performance in Stage 2 in Experiments 2 and 3.

We have focused much of our discussion on how learning the relational structures present in Stage 1 might alter subsequent learning of those cues initially involved in biconditional discriminations. Of course, the component discrimination also follows a relational structure that could potentially be learned and transferred to new situations. It is worth considering whether all of the transfer effects observed across these experiments could be explained in terms of learning relational structure. We observed a considerably stronger transfer effect in Experiment 1 than in Experiment 3, which suggests that individuals do not learn a completely abstract relation such as "in any given pair, one cue is predictive and the other is irrelevant" that is retrieved when presented with a cue from a component discrimination. This would generate transfer effects in Experiment 3 but not Experiment 1, where a component cue is present in all Stage 2 discriminations and thus the trial types in *every* discrimination could, in principle, retrieve the abstract relation. Alternatively, participants might retrieve cue-specific relational information on observing a previously predictive component that suggests that the *other* presented cue is irrelevant, regardless of its prior functional role. That is, observing cue A retrieves information that the learner should attend to A and ignore whatever else is presented. This would be more consistent with the transfer effect observed in Experiment 1, and potentially also the weaker pattern observed in Experiment 3, since information of this kind might cancel itself out if both of the cues presented in stage 2 combinations were previously predictive. But it is worth noting that this account is difficult to distinguish from the conventional feature-driven explanation that learned predictiveness enhances attention paid to predictive stimuli, potentially at the cost of learning about other features presented at the same time.

The second proposed change in cue processing that may result from learning complex discriminations is a change from elemental to configural processing. For instance, in reviewing a range of evidence from animal and human learning, Melchers, Shanks, and Lachnit (2008) argued that learning could shift dynamically between relying on elemental and configural cue representations, and that these shifts might be motivated by the functional role that cues served in prior learning or the subject's experience with configural (i.e. not linearly solvable) learning problems in similar contexts. Some criticism of this proposal was raised essentially because elemental models of associative learning do learn nonlinear problems (e.g. see Harris, 2006; Harris & Livesey, 2010; Thorwart, Livesey & Harris, 2012), and thus it is difficult to claim that experience with these problems should encourage configural processing, or that improvement on these problems constitutes evidence for configural processing (Livesey & Harris, 2008). Nevertheless a theoretically more neutral version of this hypothesis would put it that experience with nonlinear problems might change the extent to which learning about a cue among one set of stimuli generalises to the same cue among other sets of stimuli. This learning could also be stimulus-specific, as would be necessary to explain the current results. Thus, perhaps experience learning about a cue in a configural problem affects processing of that cue such that it is treated in a more configural way in future learning episodes.

Either of these two changes in the processing of cues and cue-outcome relationships could be relevant to the transfer effects observed in Experiments 2 and 3. In either case, it is interesting that the effects observed in Experiment 2, where a positive transfer effect was observed for Stage 1 biconditional cues, were confined to high performing individuals whereas the effect observed in Experiment 3, where a negative transfer effect was observed for Stage 1 biconditional cues was general to high- and low-performing participants. The former result is consistent with previous studies linking rule transfer to fast acquisition of patterning discriminations (Shanks & Darby, 1998). It is not clear why these transfer effects should only be linked to Stage 1 performance in Experiment 2 and not Experiment 3 but we suggest that using relational knowledge to improve performance on a relatively difficult task (like the Stage 2 biconditional discriminations in Experiment 2) is cognitively more challenging and requires deliberate effort on the part of the participant. This assumes that the structural transfer effect we see in Experiment 2 is under high cognitive control, which is consistent with literature on the Shanks-Darby patterning task that suggests rule transfer is associated with working memory, cognitive reflection and other deliberative processes (Don et al., 2015; 2016; Wills et al., 2011a; 2011b). In contrast, the cue associability changes associated with the learned predictiveness effect do not appear to load on cognitive resources in the same way (Pinto, Vogel & Núñez, 2017). Although Mitchell et al. (2012) found that the learned predictiveness in one experiment was completely reversible by instructions, others have reported persistent resistance to instructed reversal (Don et al., 2015; Shone et al., 2015). Baetu et al. (2018) found that improvements in performance completing the second of two successive biconditional discriminations were predicted by measures of fluid ability. Rather than concluding that this effect was based on rule-learning, they argued that this demonstrates flexible use of configural processing amongst individuals with higher cognitive capacities. Thus it is possible, for instance, that fast learners in our task engage more flexible stimulus representations in ways that are not related to rule learning *per se*. The learning tasks that we present here involve relatively complex designs, for instance, using more cue-outcome contingencies than most other studies we are aware of that examine the learning of biconditional discriminations. We can only speculate as to whether engaging in flexible stimulus processing to represent configurations of cues ought to require additional cognitive resources but, if this were the case, then we would expect individuals with stronger learning capacity to show the most pronounced transfer effects.

There are, of course, other important differences in the methodology used by this study and by Uengoer and Lachnit (2012). These experiments (and those reported by Livesey et al., 2011) were run completely within-subjects whereas Uengoer and Lachnit used a between subjects design. Unegoer and Lachnit also used features organised along coherent stimulus dimensions such that whole dimensions were relevant or irrelevant to the discrimination. The fact that these properties of the experiment make an important difference to the outcome highlights the need to better understand the underlying processes that lead to associability changes as the conventional theoretical approaches offered by formal models of attention-based associative learning do not speak to these issues directly.

References

Baetu, I., Burns, N. R., Yu, E., & Baker, A. G. (2018). Fluid Abilities and Rule Learning: Patterning and Biconditional Discriminations. *Journal of Intelligence, 6*, 7; doi:10.3390/jintelligence6010007.

Beesley, T., Nguyen, K. P., Pearson, D., & Le Pelley, M. E. (2015). Uncertainty and predictiveness determine attention to cues during human associative learning. *Quarterly Journal of Experimental Psychology, 68*, 2175-2199.

Cobos, P. L., Gutiérrez-Cobo, M. J., Morís, J., & Luque, D. (2017). Dependent measure and time constraints modulate the competition between conflicting feature-based and rule-based generalization processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 43,* 515-.

Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson’s method. *Tutorials in quantitative methods for psychology*, *1*, 42-45.

Don, H. J., Goldwater, M. B., Otto, R., Livesey, E. J. (2016). Rule abstraction, model-based choice and cognitive reflection. *Psychonomic Bulletin & Review, 23,* 1615–1623.

Don, H. J. & Livesey, E. J. (2015). Resistance to instructed reversal of the learned predictiveness effect. *Quarterly Journal of Experimental Psychology, 68,* 1327-1347.

George, D. N., & Pearce, J. M. (1999). Acquired distinctiveness is controlled by stimulus relevance not correlation with reward. *Journal of Experimental Psychology: Animal Behavior Processes, 25*, 363–373.

Goldwater, M. B., Don, H. J., Krusche, M., & Livesey, E. J. (2018). Relational Discovery in Category Learning. *Journal of Experimental Psychology: General, 147*, 1-35.

Griffiths, O., Johnson, A. M., & Mitchell, C. J. (2011). Negative transfer in human associative learning. *Psychological Science*, 22(9), 1198-1204.

Harris, J. A. & Livesey, E. J. (2010). An Attention-Modulated Associative Network. *Learning & Behavior, 38,* 1-26.

Harris, J. A., Livesey, E. J., Gharaei, S., & Westbrook, R. F. (2008). Negative patterning is easier than a biconditional discrimination. *Journal of Experimental Psychology: Animal Behavior Processes, 34*, 494-500.

Harris, J. A. & Livesey, E. J. (2008). Comparing Patterning and Biconditional Discriminations in Humans. *Journal of Experimental Psychology: Animal Behavior Processes, 34*, 144-154.

Kattner, F. (2015). Transfer of absolute and relative predictiveness in human contingency learning. *Learning & Behavior, 43*, 32–43.

Kruschke, J. K. (1996). Dimensional relevance shifts in category learning. *Connection Science, 8*, 225–247.

Le Pelley, M. E., & McLaren, I. P. L. (2003). Learned associability and associative change in human causal learning. *The Quarterly Journal of Experimental Psychology, 56B*, 68–79.

Le Pelley, M. E., Mitchell, C. J., Beesley, T., George, D. N., & Wills, A. J. (2016). Attention and associative learning in humans: An integrative review. *Psychological Bulletin, 142*, 1111-1140.

Le Pelley, M. E., Turnbull, M. N., Reimers, S. J., & Knipe, R. L. (2010). Learned predictiveness effects following single-cue training in humans. *Learning and Behavior, 38*, 126–144.

Livesey, E. J., & Harris, J. A. (2008). What are flexible representations? Commentary on Melchers, Shanks and Lachnit. *Behavioral Processes, 77*, 437–439.

Livesey, E. J., Thorwart, A., & Harris, J. A. (2011). Comparing positive and negative patterning in human learning. *Quarterly Journal of Experimental Psychology, 64*, 2316–2333.

Livesey, E. J., Thorwart, A., De Fina, N. L., & Harris, J. A. (2011). Comparing learned predictiveness effects within and across compound discriminations. *Journal of Experimental Psychology: Animal Behavior Processes, 37,* 446-465.

Livesey, E. J. & McLaren, I. P. L. (2007). Elemental Associability Changes in Human Discrimination Learning. *Journal of Experimental Psychology: Animal Behavior Processes, 33*, 148-159.

Lochmann, T., & Wills, A. J. (2003). Predictive history in an allergy prediction task. In F. Schmalhofer, R. M. Young, & G. Katz (Eds.), Proceedings of EuroCogSci 03 (pp.217–222). Mahwah, NJ: Erlbaum.

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

Melchers, K. G., Shanks, D. R., & Lachnit, H. (2008). Stimulus coding in human associative learning: flexible representations of parts and wholes. *Behavioral Processes, 77,* 413–427.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87,* 532-552.

Pinto, J. A., Vogel, E. H., & Núñez, D. E. (2017). The effect of cognitive load and outcome congruency on the learned predictiveness effect in human predictive learning. *Psicologica: International Journal of Methodology and Experimental Psychology, 38*(2), 257-281.

Saavedra, M. A. (1975). Pavlovian compound conditioning in the rabbit. *Learning and Motivation, 6,* 314–326.

Shanks, D. R., & Darby, R. J. (1998). Feature- and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes, 24,* 405–415.

Shone, L. T., Harris, I. M., & Livesey, E. J. (2015). Automaticity and Cognitive Control in the Learned Predictiveness Effect. *Journal of Experimental Psychology: Animal Learning and Cognition, 41*, 18-31.

Thorwart, A., Uengoer, M., Livesey, E. J. & Harris, J. A. (2017). Summation effects in human learning: evidence from patterning discriminations in goal-tracking. *Quarterly Journal of Experimental Psychology, 70,* 1366–1379.

Thorwart, A., Livesey, E. J. & Harris, J. A. (2012). Normalisation between stimulus elements in a model of Pavlovian conditioning: Showjumping on an elemental horse. *Learning & Behavior, 40*, 334-346.

Uengoer, M., & Lachnit, H. (2012). Modulation of attention in discrimination learning: The roles of stimulus relevance and stimulus-outcome correlation. *Learning & Behavior, 40,* 117–127.

Wills, A. J., Barrasin, T. J., & McLaren, I. P. L. (2011). Working memory capacity and generalization in predictive learning. In L. Carlson, C. Hölscher, & T. Shipley (Eds.), *Proceedings of the 33rd annual conference of the Cognitive Science Society* (pp. 3205–3210). Austin, TX: Cognitive Science Society.

Wills, A. J., Graham, S., Koh, Z., McLaren, I. P. L., & Rolland, M. D. (2011). Effects of concurrent load on feature- and rule-based generalization in human contingency learning. *Journal of Experimental Psychology: Animal Behavior Processes, 37*, 308–316.