THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY
2010, 63 (1), 108–135

Psychology Press
Taylor & Francis Group

The effect of predictive history on the learning of sub-sequence contingencies

T. Beesley and M. E. Le Pelley
Cardiff University, Cardiff, UK

Two experiments demonstrated that the prior predictive history of a cue governs the extent to which that cue engages in sequence learning. Using a serial reaction time task, we manipulated the predictiveness of the stimulus locations (cues) with respect to the location of the stimulus on the next trial (outcome), such that half of the cues were good predictors of their outcomes, whilst the other half were poorer predictors. Following this, all cues were then paired with novel outcomes. Learning about those cues that were previously established as good predictors proceeded more rapidly than learning for those cues previously established as poor predictors. When the simple recurrent network is modified to include a variable associability parameter, the effects are easily modelled.

Keywords: Incidental sequence learning; Predictive history; Associability; Simple recurrent network.

A fundamental feature of human cognition is the ability to select the appropriate action for an event given a set of contextual cues. In the case of many complex real-world tasks such as driving vehicles, athletic activity, and playing musical instruments, these cues will take the form of a temporal series, often experienced in rapid succession. A learning system that is able to predict future events (on the basis of past events and actions) enables us to prepare future actions in advance and hence, potentially, to execute these actions more rapidly and accurately. In this article we examine one factor that might influence the rate of learning about a cue during a sequence learning task—namely, the predictive history of that cue.

Sequence learning has been studied extensively over the past 20 years using the Serial Reaction Time (SRT) task (Nissen & Bullemer, 1987). In this task participants respond to a target stimulus, which can appear at one of a set number of locations on the screen (usually 4 or 6), using corresponding response keys. Participants are commonly informed that the purpose of the task is to "investigate the effect of practice on motor control" and are instructed to respond as rapidly and as accurately as possible each time the target stimulus appears. Unbeknownst to participants, the movement of the target is governed by a consistent sequence. With practice, participants' reaction times tend to decrease, and, by using control trials in which

Correspondence should be addressed to T. Beesley, who is now at the Division of Psychology and Language Sciences, University College London, 26 Bedford Way, London WC1H OAP, UK. E-mail: t.beesley@ucl.ac.uk

We are grateful to Axel Cleeremans, Luis Jiménez, Andy Delamater and two anonymous reviewers for their helpful comments on an earlier draft of this manuscript. We would like to thank the Advanced Research Computing at Cardiff (ARCCA) team for their help in running the model simulations. This work was supported by Grant RES000230983 from the Economic and Social Research Council.

http://www.psypress.com/qjep DOI:10.1080/17470210902831767

the target makes unsequenced (often random) movements, at least part of this increase in speed can be attributed to learning of the underlying sequence. Interestingly, when participants are given a variety of verbal report and recognition memory tests, several studies have suggested that participants' explicit knowledge of the underlying sequence often seems incomplete (e.g., Destrebecqz & Cleeremans, 2001; Willingham, Nissen, & Bullemer, 1989), especially when a probabilistic method of sequence generation is employed (e.g., Jiménez, Méndez, & Cleeremans, 1996). The extent to which the knowledge can be described as implicit in this task has sparked a considerable amount of research and debate within the field (for a review, see Shanks, 2005). Whilst some have argued in favour of implicit learning (e.g., Reber, 1993), others have argued that the evidence for learning without awareness remains inconclusive (e.g., Wilkinson & Shanks, 2004). Furthermore, it has been suggested that the current methods for assessing conscious knowledge lack sufficient sensitivity to reveal dissociations between implicit and explicit tests (see Shanks & St. John, 1994). The current article is not intended to contribute directly to this debate; we avoid making strong claims as to whether participants can be classified as unaware in the tasks described below. Nevertheless, it is generally accepted that learning in the SRT task is incidental (see e.g., Cleeremans, Destrebecqz, & Boyer, 1998). That is, there is no directed instruction to participants to learn—participants are asked to respond to, rather than predict, the outcome on each trial, such that the task can be performed with perfect accuracy in the absence of any learning. Consequently there is no explicit requirement for participants to intentionally exploit knowledge of the situation.

The extent to which implicit learning is dependent on attentional resources has been studied extensively (for a review, see Shanks, Rowland, & Ranger, 2005). Studies typically use dual-task conditions, in which participants are given an attention-demanding secondary task (e.g., tone counting) in addition to the primary SRT task. Secondary tasks are presumed to provide sufficient working-memory load to limit the processing resources available for attentional learning. Hence, evidence that learning proceeds despite these resources being unavailable points to the presence of automatic learning mechanisms. Such results have been shown on several occasions (e.g., Cohen, Ivry, & Keele, 1990; Curran & Keele, 1993; Jiménez & Mendez, 1999); notably it seems the complexity of the sequences used in the primary task is a crucial factor in observing this effect (Shanks et al., 2005).

Although most studies examining the role of attention in sequence learning have focused on attention as a cognitive resource, a small but related literature has examined the extent to which learning in the SRT task is governed by selective attentional mechanisms. The term "selective attention" here refers to the orienting of attention towards the sequences within the task. One common method of examining selective attention is to present a secondary sequence in addition to the main SRT task. For example, Mayr (1996) arranged for the target stimulus to appear in a sequence of locations and also as a sequence of identities (different colours), with participants directed to respond only to the identity of the stimulus. Having not been instructed towards any sequenced movement of the stimulus, and despite this sequence having no facilitatory effect on the primary task of responding to target identity, participants were disrupted when the secondary location sequence was swapped to a control sequence. Findings such as these have suggested that learning can proceed when attention is not directly oriented towards the to-be-learnt material.

In this article we examine whether learning in a standard, single-sequence learning task proceeds uniformly, or whether attention-like selectional mechanisms also operate within the learning of a single sequence. In particular, we assess whether the predictive history of a cue governs the extent to which that cue engages the learning system and therefore the rate of learning about that cue. Why should certain cues in a sequence learning task be learnt about more readily than others? It seems reasonable to assume that this would happen when there is a discrepancy in the validity of parts

of a sequence. In other words, it might be the case that fewer processing resources will be devoted to learning about those parts of the sequence that are deemed less beneficial for sequence learning than to learning about those parts that are deemed more beneficial. This suggestion that the prior predictiveness of a cue modulates the rate of subsequent learning about that cue has been confirmed in both human contingency learning (hereafter HCL) and animal conditioning (see Le Pelley, 2004, for a review). The general procedure in these experiments is to pretrain a subset of cues as good predictors of outcomes, whilst simultaneously pretraining a subset of cues as poor predictors of outcomes. During a second stage, in which novel cue-outcome pairings are presented, new learning about good and poor predictors is assessed. The results of a number of recent studies of HCL indicate that learning about cues that were pretrained as good predictors proceeds at a faster rate than learning for cues pretrained as poor predictors (e.g., Kruschke & Blair, 2000; Le Pelley, Beesley, & Suret, 2007; Le Pelley & McLaren, 2003; Le Pelley, Suret, & Beesley, in press). Furthermore, evidence from eye-tracking studies suggests that participants devote less attention to cues that are poor or redundant predictors of the outcome that follows them than to cues that are good predictors of this outcome (Kruschke, Kappenman, & Hetrick, 2005; Wills, Lavric, Croft, & Hodgson, 2007; although see Hogarth, Dickinson, Austin, Brown, & Duka, 2008).

Much debate still surrounds the fundamental mechanisms governing HCL (see Shanks, 2007, for a review). While some authors have proposed that many aspects of contingency learning are best accommodated within an associative framework (Dickinson, Shanks, & Evenden, 1984; Le Pelley, Oakeshott, & McLaren, 2005a), others have questioned the importance of associative processes and have instead argued that contingency learning is invariably a product of higher order reasoning (Beckers, De Houwer, Pineño, & Miller, 2005; De Houwer, Vandorpe, & Beckers, 2005). Notably, these two opposing views are divided on the issue of the cognitive resources necessary for learning to proceed: Associative mechanisms are typically assumed to operate automatically (Squire, 1994; although see Lovibond & Shanks, 2002), whilst higher order reasoning is governed by controlled and effortful evaluation of the participants' beliefs (Lovibond, 2003). Given the incidental nature of learning in the SRT task, the current experiments will reveal whether the influence of predictive history on novel learning is contingent on directed instruction to learn. In the absence of such instruction, and given that the rapid mode of stimulus presentation used in the SRT task limits the available time in which participants can engage higher order reasoning processes (e.g., hypothesis testing), we would argue that the contribution of such controlled reasoning processes to learning is minimized in this task. To the extent that this is indeed the case, the results will also address the question of whether effects of predictive history are driven by controlled reasoning processes or automatic associative mechanisms.

In describing the manipulations made in our experiments we adopt the terminology of contingency learning studies, using the terms "cue" and "outcome" to describe successive movements in the sequence. For example, given the sequence 1234, in describing the first transition 1–2, we would refer to element 1 as the cue and element 2 as the outcome. For the next transition, 2–3, element 2 becomes the cue, and 3 is the outcome, and so on. In order to manipulate predictive history differentially amongst the cues in our task, we use probabilistic sequences, such that a subset of cues are relatively good predictors of their respective outcomes, whilst a different subset are less predictive of their respective outcomes. The experiments presented here feature designs that are analogous to those used in previous studies of human contingency learning: Pretraining establishes cues as good or poor predictors of outcomes, after which we assess the influence of this manipulation on novel learning about these cues.

EXPERIMENT 1

The aim of Experiment 1 was to manipulate the predictive history of certain elements of the

sequence, by pretraining a subset of elements as good predictors of the subsequent element. As this manipulation was based largely on previous work in both animal conditioning and human contingency learning, we decided to assign elements in the stimulus array to act as either "cue" or "outcome" positions in the task. Using a six-choice task, the four central positions (2–5) acted as cues, whilst the two outer positions (1 and 6) acted as outcomes—this distinction was not explicitly communicated to participants, however. The contingencies between cues and outcomes were manipulated such that two of the cues were good predictors of outcomes, and two were poor predictors of outcomes during a first stage of training. In a subsequent second stage, the two outer positions were removed, and participants were trained on a four-choice task in which all of the remaining positions (two of which had been good predictors in Stage 1 and two of which had been poor predictors) were now equally predictive of their respective outcomes. We could then compare the rate of learning about previously good and previously poor predictors during Stage 2. Given previous findings relating to manipulations of predictive history in human contingency learning, we expected that cues that had been pretrained as good predictors during Stage 1 would be learned about more rapidly in Stage 2 than those pretrained as poor predictors during Stage 1.

Method

Participants, apparatus, and stimuli

A total of 20 Cardiff University undergraduates, who had not participated in an SRT task before, participated for course credit or payment. Testing was conducted in a quiet room divided into two booths to allow 2 participants to be tested at the same time, using PCs with 17″ TFT monitors; participants sat approximately 80 cm from the monitor. The experiment was run using software written in Visual Basic. Reaction times (hereafter RTs) were recorded with Windows API functions QueryPerformanceCounter and QueryPerformanceFrequency for millisecond resolution. Responses were made with a standard keyboard using the keys X, C, V, B, N, M. Error signals were presented over headphones. The stimulus array consisted of six grey response circles (3 mm in diameter), evenly spaced 25 mm apart in a horizontal line across the middle of the screen. The target stimulus was a larger grey circle, 12 mm in diameter. On each trial the target stimulus would appear 20 mm below one of the response circles.

Sequence generation

For all participants the outcome locations were always in positions 1 and 6. The remaining positions (2–5) acted as cue elements. In describing the generation of the sequence we use as an example elements 2 and 3 to denote good predictors and elements 4 and 5 to denote poor predictors. Note, however, that this assignment is merely an example and that for each participant the four cue elements (two good predictors and two poor predictors) were randomly assigned to the locations 2–5 in the stimulus array.

Stage 1. Table 1 shows the conditional probabilities for transitions used in Stage 1. The location in which the target appeared on a given trial (trial $N + 1$) was determined by the position in which it had appeared on the previous trial (trial N). For example, if the target appeared in position 2 on trial N, then on trial $N + 1$ there would be a probability of .9 of the target appearing in position 1 and a probability of .1 of the target appearing in position 6. Similarly, if the target appeared in position 3 on trial N, then on trial $N + 1$ there would be a probability of .9 of the target appearing in position 6, and a probability of .1 of the target appearing in position 1. Thus, position 2 was a good predictor of the target appearing in position 1 on the next trial, whilst position 3 was a good predictor of the target appearing in position 6. When the target appeared in position 4 or 5 on trial N, it would appear in positions 1 and 6 on trial $N + 1$ with equal probability of .5. Thus, positions 4 and 5 were relatively poor predictors of the location of the target on the next trial. After an outcome trial (positions 1 and 6) the location of

Table 1. *Conditional probabilities of stimulus transitions for Stages 1 and 2 of Experiment 1*

	Trial N (cue)											
	Stage 1						*Stage 2*					
Trial N+1 (outcome)	*1*	*2*	*3*	*4*	*5*	*6*	*1*	*2*	*3*	*4*	*5*	*6*
1		.9	.1	.5	.5							
2	.25					.25			.1	.1	.8	
3	.25					.25		.8		.1	.1	
4	.25					.25		.1	.8		.1	
5	.25					.25		.1	.1	.8		
6		.1	.9	.5	.5							

Note: Blank probability = 0. Positions 1 and 6 were removed from the task during Stage 2.

the next target was selected at random from the four cue positions (2–5), with the caveat that the sequence could feature the same pair of locations twice in a row (e.g., 121213), but not three times (e.g., 121212).

Stage 2. For the last two blocks of the experiment the two stimulus locations used as outcomes in Stage 1 (locations 1 and 6) were removed from the stimulus array, creating a four-choice task. Removing outcome positions 1 and 6 should reduce participants' tendencies to make these well-established outcome responses during Stage 2, hence reducing any potential response interference effects. Table 1 shows the conditional probabilities for the Stage 2 sequence. Each of the remaining four stimulus locations acted as both a cue and an outcome during Stage 2. In other words, the outcome ($N + 1$) for any given contingency would become the cue (N) for the next contingency (see Results for a more detailed discussion of this method of sequence generation). It is clear from Table 1 that the possible Stage 2 transitions were entirely different from those used in Stage 1. All cues now predicted one location with a probability of .8 and the two other locations with a probability of .1 each (repetitions were not permitted)—therefore in Stage 2 all cues were now equally valid predictors of outcomes. Nevertheless, we continue to use the terms "good predictors" and "poor predictors" to refer to the way in which these sets of cues had been pretrained in Stage 1.

Procedure

At the outset of the experiment, the instructions contained in the Appendix were read to participants. They were then asked to position their index fingers on keys V and B, their middle fingers on keys C and N, and their ring fingers on keys X and M.

The position of the target on the first six trials of all blocks was randomly determined. On each trial the target was displayed in one of the six positions and remained on the screen until participants made a response. When a response was made the appropriate response circle would turn red for 120 ms, after which the target stimulus would disappear, the response circle would reset to grey, and the program paused for a further 120 ms. The response–stimulus interval was therefore 240 ms. A response in an incorrect location produced a beep.

Stage 1 consisted of 10 blocks of 150 trials, with a rest break of 15 s between blocks. During the rest break between the 10th and 11th blocks, the following message was displayed: "PLEASE NOTE! From now on the circle will only appear in the middle 4 positions, using keys C, V, B and N." This was the first time participants were informed that the task would change in this way. The two outermost positions (1 and 6) were

removed from the stimulus array for the remainder of the task. Stage 2 comprised 2 blocks of 150 trials.

Results

The first six trials of each block were not analysed, since RTs on trials following a break tended to be longer than those during the remainder of the block. Mean error rates were 3% ($SD = 2.2$) and 3.6% ($SD = 2.5$) for Stages 1 and 2, respectively, and RTs on these trials were excluded from further analysis. Accuracy and RT data were also excluded for trials following an error, as were trials in which RTs were less than 100 ms (0.02% of trials in Stage 1; 0.14% in Stage 2) or greater than 1,000 ms (2.93% in Stage 1; 1.49% in Stage 2).

During Stage 1, three types of trial were of interest: high-probability outcomes following good predictors (GPH; occurring with a probability of .9); low-probability outcomes following good predictors (GPL; occurring with a probability of .1); and medium-probability outcomes following poor predictors (PPM; occurring with a probability of .5). For each of these trial types, the analysed trial was that occurring on trial $N + 1$ (the outcome). For example, consider the sequence 214631 and the conditional probabilities for these transitions in Table 1. The first transition, 2−1, is of type GPH since outcome 1 is predicted by cue 2 with a probability of .9. The extent to which participants have learnt that cue 2 predicts outcome 1, will be shown in the RT and accuracy to outcome 1, and it is therefore this trial (outcome 1) which contributes to the average for trial type GPH. The next transition, 1−4, is from outcome location 1 to cue 4, is not a trial of interest, and is not analysed. The next transition, 4−6, is of trial type PPM, since outcome 6 is predicted by cue 4 with a probability of .5. The response to outcome 6 will therefore contribute to the average for trial type PPM. Transition 6−3 is from outcome location 6 to cue 3 and is not analysed. Finally, transition 3−1 is of type GPL, as outcome 1 is predicted by cue 3 with a probability of .1. The response to outcome 1 will contribute to the average for trial type GPL.

Due to the low probability of an inconsistent outcome occurring after a good predictor cue in Stage 1, a participant would occasionally produce no data for the GPL trial type in a given block. There were 7 instances of missing RT data and 7 instances of missing accuracy data. In order to conduct a full analysis of variance (ANOVA), missing data were estimated using an average from the two blocks immediately before and after the block with missing data. When this was not possible (in Blocks 1 and 10) the nearest adjacent data were used.

Figure 1A shows RTs across the 10 blocks of Stage 1 for the three trial types. By the end of Stage 1, participants are fastest on outcomes that could be predicted with a high probability (GPH), slowest on those that could be predicted with lowest probability (GPL), and of intermediate speed on those with a medium probability (PPM). These data were subjected to repeated measures ANOVA, with factors of outcome probability (GPH, GPL, and PPM) and block. There was a significant effect of outcome probability, $F(2, 38) = 13.49$, $MSE = 6{,}237.38$, $p < .001$, and block, $F(9, 171) = 6.15$, $MSE = 2{,}495.25$, $p < .001$. The interaction between outcome probability and block was also significant, $F(18, 342) = 1.83$, $MSE = 1{,}198.13$, $p < .05$, which suggests that participants' sensitivity to the contingencies within the sequence increased with continued exposure during Stage 1. Pairwise comparisons between the three levels of the outcome predictability variable revealed significant differences between all three: GPH versus GPL trials, $F(1, 19) = 15.64$, $MSE = 21{,}516.03$, $p < .01$; PPM versus GPL trials, $F(1, 19) = 16.82$, $MSE = 4{,}828.50$, $p < .01$; GPH versus PPM trials, $F(1, 19) = 7.86$, $MSE = 11{,}079.76$, $p < .05$.

Accuracy data for Stage 1 are shown in Figure 1B. Although accuracy offers a less sensitive measure of learning in the SRT task, the same ordinal relationships as those that are observed in RTs emerge later on in Stage 1: Participants are most accurate on GPH trials, are least accurate on GPL trials, and show intermediate

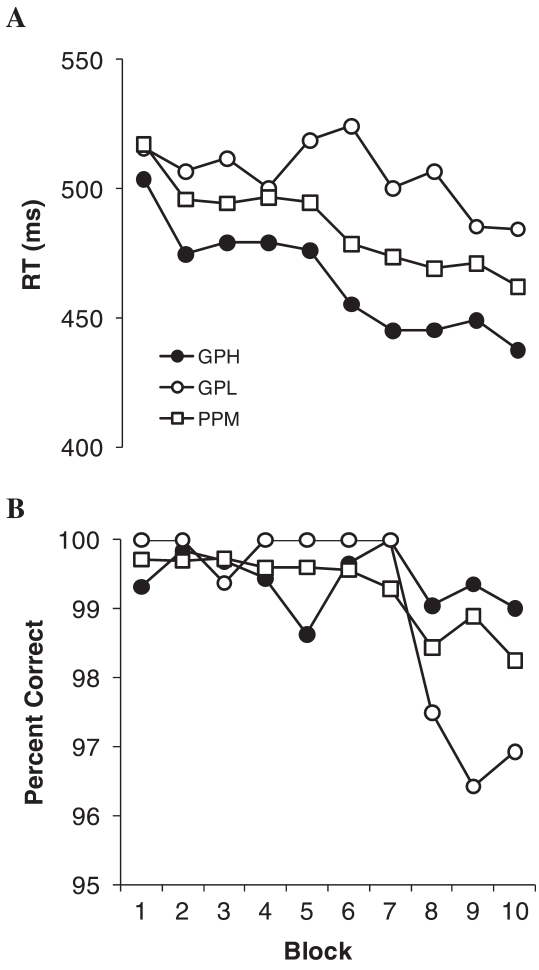

Figure 1. *(A) Reaction time (RT), and (B) accuracy data for Stage 1 of Experiment 1. GPH: responses to high-probability outcomes following good predictor cues; GPL: responses to low-probability outcomes following good predictor cues; PPM: responses to medium-probability outcomes following poor predictor cues.*

performance on PPM trials. These analyses were subjected to ANOVA, which revealed no effect of outcome probability, $F < 1$, but a significant effect of block, $F(9, 171) = 2.37$, $MSE = 14.08$, $p < .05$, indicating a general decline in accuracy towards the end of Stage 1. The interaction was not significant, $F < 1$.

During Stage 2, four trial types were of interest: high-probability outcomes following cues that were pretrained as good predictors in Stage 1 (GPH; occurring with a probability of .8); low-probability outcomes following cues that were pretrained as good predictors in Stage 1 (GPL; occurring with a probability of .2); high-probability outcomes following cues that were pretrained as poor predictors in Stage 1 (PPH; occurring with a probability of .8); low-probability outcomes following cues that were pretrained as poor predictors in Stage 1 (PPL; occurring with a probability of .2). Stage 2 analysis followed a similar method to that outlined above for Stage 1; however, in Stage 2 all trials contributed to a trial type average. Consider the sequence 23545 and the contingencies given in Table 1. The first transition, 2–3, is of trial type GPH, as cue 2 was pretrained as a good predictor in Stage 1, and outcome 3 now occurs with a probability of .8 after cue 2 (it is a high-probability transition in Stage 2). The next transition, 3–5, is of trial type GPL, since cue 3 was pretrained as a good predictor in Stage 1, and outcome 5 now occurs with a probability of .1 after cue 3 (it is a low-probability transition in Stage 2). The next transition, 5–4, is of trial type PPL, since cue 5 was pretrained as a poor predictor in Stage 1, and outcome 4 occurs with a probability of .1 after cue 5. The final transition, 4–5, is of trial type PPH, as cue 4 was pretrained as a poor predictor in Stage 1, and outcome 5 occurs with a probability of .8 after cue 4.

RT data for the two blocks of Stage 2 are shown in Figure 2A. Overall, learning progressed rapidly for both good and poor predictor contingencies, as indicated by the difference in RTs to high- and low-probability outcomes. These data were subjected to repeated measures ANOVA, with factors of prior predictiveness (i.e., the predictiveness of the cues in Stage 1—good vs. poor), outcome probability (high vs. low), and block. This revealed a significant main effect of outcome probability, $F(1, 19) = 61.62$, $MSE = 3{,}658.25$, $p < .001$, a marginally significant effect of block, $F(1, 19) = 3.27$, $MSE = 961.21$, $p = .086$, but no main effect of prior predictiveness, $F(1, 19) = 1.28$, $MSE = 1{,}119.14$, $p = .27$. The interaction of prior predictiveness and block was significant, $F(1, 19) = 5.24$, $MSE = 314.23$, $p < .05$, which reflects the different pattern of RTs to

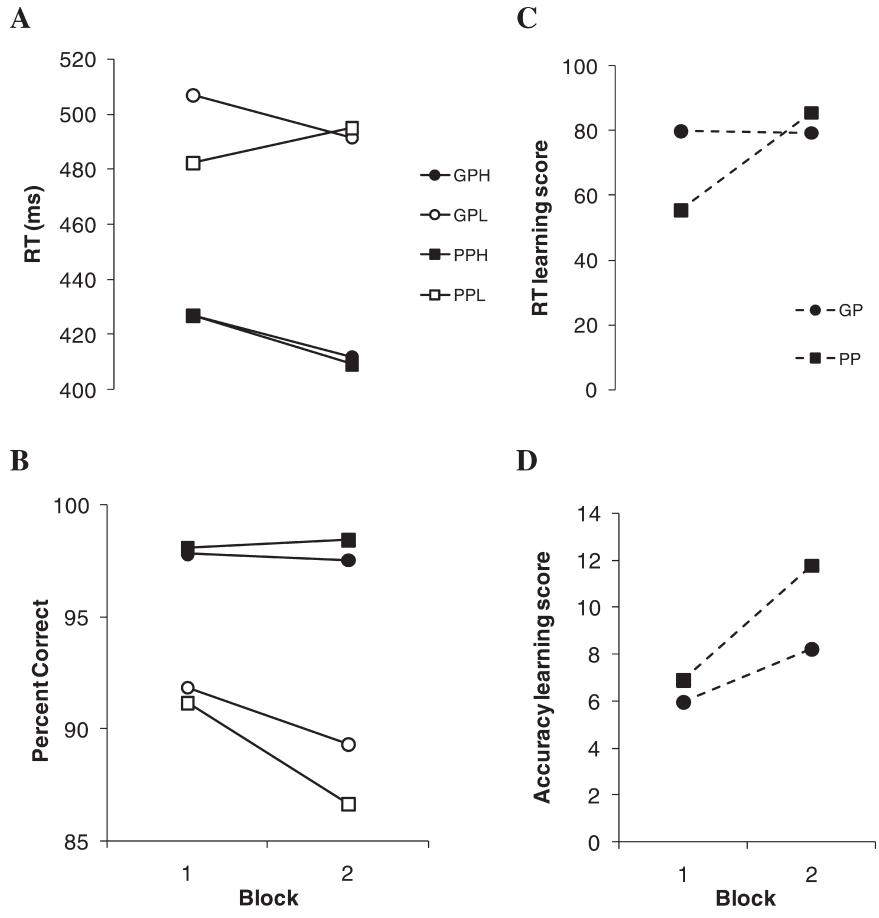

Figure 2. *(A) Reaction time (RT), and (B) accuracy data for Stage 2 of Experiment 1, with data plotted as learning scores—the difference between responses to high- and low-probability outcomes—in Panels (C) and (D), respectively. GPH: responses to high-probability outcomes following cues pretrained as good predictors in Stage 1; GPL: responses to low-probability outcomes following cues pretrained as good predictors in Stage 1; PPH: responses to high-probability outcomes following cues pretrained as poor predictors in Stage 1; PPL: responses to low-probability outcomes following cues pretrained as poor predictors in Stage 1.*

low-probability outcomes for good and poor predictors. No other interactions were significant: prior predictiveness by outcome probability, $F(1, 19) = 1.71$, $MSE = 484.48$, $p = .21$; outcome probability by block, $F(1, 19) = 2.43$, $MSE = 909.98$, $p = .14$; nor the three-way interaction, $F(1, 19) = 2.84$, $MSE = 821.11$, $p = .11$.

In this task, learning about cue-outcome relationships can be assessed by subtracting RTs to high-probability outcomes following a particular cue position from RTs to low-probability outcomes following the same cue position. This yields a "learning score", which reflects how much participants have learnt about high-probability outcomes relative to low-probability outcomes, and provides a means of comparing how much has been learnt about good predictors and poor predictors in Stage 2, unconfounded from response interference effects. Given the differential pretraining of good and poor predictors in Stage 1, we might expect proactive interference to have a larger detrimental effect on good predictor contingencies than on poor predictor contingencies in Stage 2. Consider, as an example, a participant

whose mapping of elements 1 to 6 (shown in Table 1) directly translates onto the stimulus locations 1 to 6 on the screen (e.g., element 1 maps to key X, element 2 to key C, and so on). Suppose that during Stage 1 this participant has learnt that cue 2 is very likely to be followed by outcome 1—that is, this participant has learnt to prepare the response of pressing the X key on the trial following cue 2. In Stage 2 all cue-outcome contingencies are changed, such that cue location 2 is now usually followed by location 3. This participant's previously learnt tendency to prepare an X response following location 2 is therefore now inappropriate and might be expected to interfere with appropriate responding on this trial (pressing the V key). This can be compared to previously poor predictors, for which participants will not have developed such strong response tendencies during Stage 1, and hence for which there will be less interference during Stage 2. Consequently, during Stage 2, response interference could potentially mask any advantage that might exist for learning the new high-probability responses for those cues previously trained as good predictors. However, any proactive interference affecting responding during Stage 2 will have an equivalent effect on responses made to both high- and low-probability outcomes. Taking the difference between these two trial types as our measure of learning for the Stage 2 contingencies therefore allows us to subtract out any influence of response interference, such that any difference observed will reflect a difference in the rate of learning about good and poor predictor cues.

While the nonsignificant interaction of prior predictiveness with outcome probability in the ANOVA reported above indicates that this difference between RTs to high- and low-probability outcomes does not differ significantly between good and poor predictors when assessed across the whole of Stage 2 training, a finer grained analysis based on learning scores reveals a significant, yet short-lived, influence of prior predictiveness. Figure 2C plots the difference between RTs on high- and low-probability outcomes (i.e., the learning score) for both good and poor predictors. The learning score data were subjected to ANOVA with factors of cue and block. There was no main effect of cue, $F(1, 19) = 1.70$, $MSE = 969.37$, $p = .21$, nor block, $F(1, 19) = 2.43$, $MSE = 1{,}821.01$, $p = .14$, and no interaction, $F(1, 19) = 2.84$, $MSE = 1{,}642.74$, $p = .11$. Preplanned tests of simple effects revealed a significant effect of cue in Block 1, $F(1, 19) = 5.10$, $MSE = 1{,}163.91$, $p < .05$, but not in Block 2, $F < 1$. Consistent with our experimental hypothesis, the significant difference in Block 1 indicates that learning about good predictor cues was initially at an advantage during Stage 2 learning.

Accuracy data for Stage 2 are shown in Figure 2B. Again, the data show that acquisition progresses rapidly in Stage 2, with accuracy on high-probability outcomes greater than that on low-probability outcomes. These data were again subjected to ANOVA with factors of prior predictiveness, outcome probability, and block. This revealed a significant effect of outcome probability, $F(1, 19) = 24.18$, $MSE = 111.59$, $p < .001$, but no effect of prior predictiveness, $F < 1$, nor block, $F(1, 19) = 2.41$, $MSE = 49.93$, $p = .14$. There were no significant interactions between any of the factors: outcome probability by prior predictiveness, $F(1, 19) = 1.18$, $MSE = 42.36$, $p = .29$; outcome probability by block, $F(1, 19) = 2.45$, $MSE = 51.52$, $p = .13$; prior predictiveness by block, $F < 1$; nor the three-way interaction, $F < 1$. Figure 2D shows learning scores for accuracy data (accuracy on high-probability outcomes minus accuracy on low-probability outcomes). These data were subjected to ANOVA, which revealed no main effect of cue, $F(1, 19) = 1.21$, $MSE = 85.64$, $p = .29$, nor of block, $F(1, 19) = 2.41$, $MSE = 103.17$, $p = .14$, and no interaction, $F < 1$. Preplanned tests of simple effects found no significant effect of cue in Block 1 or Block 2, maximum $F(1, 19) = 1.50$, $p = .24$.

Discussion

Experiment 1 examined whether incidentally pretraining cues as either good or poor predictors of outcomes has an effect on the subsequent rate of learning about these cues. By the end of Stage 1, the pattern of participants' responses reflected the

different outcome probabilities within the sequence: Our manipulation of predictiveness was successful. In Stage 2 the contingencies between the elements changed to entirely novel transitions that were not pretrained during Stage 1. All cues were now equally predictive of their respective outcomes. Nevertheless, those cues that were previously good predictors in Stage 1 were learnt about faster during the initial part of Stage 2 than were those that were previously poor predictors. By the end of Stage 2, learning about poor predictor cues was as great as learning about good predictors. This is presumably because learning of the good predictors approached asymptote by the end of Block 1, allowing learning about poor predictors to "catch up" during Block 2. The results of Experiment 1 indicate that predictive history influences the rate at which cues are learnt about in sequence learning, and, in line with previous work in HCL (e.g., Le Pelley & McLaren, 2003), our findings suggest that pretraining cues as good predictors facilitates the acquisition of novel outcome associations for these cues relative to those pretrained as poorer predictors.

Although Experiment 1 provided some evidence for an effect of prior predictiveness on new learning in the SRT task, this evidence was restricted to an analysis on the short-lived effect in Block 1. Experiment 2 was therefore conducted as a replication of this novel effect, using a design that aimed to increase the chances of observing a greater effect of prior predictiveness on Stage 2 learning.

EXPERIMENT 2

The design of Experiment 1 was analogous to previous studies examining the influence of predictive history in animal conditioning and HCL (e.g., Le Pelley & McLaren, 2003), in that certain stimuli acted as cues, whilst others acted as outcomes. It is unclear, however, whether the effect of predictiveness observed in Experiment 1 is limited to this arrangement of designated cue and outcome positions. In Experiment 2 we sought to replicate the effect of predictive history using a six-choice SRT task in both stages of the experiment, with stimulus locations acting as both "cues" and "outcomes" throughout.

Using a six-choice SRT task in both stages of the experiment also allowed us to address another factor that might potentially have influenced the results of Experiment 1—namely, the change in context between Stages 1 and 2. That is, the removal of the two "outcome" positions following Stage 1 of Experiment 1 presumably made it clear to participants that the structure of the task, and the movements of the target, would be different in Stage 2. A similar argument applies to most of the previous studies of predictive history effects in HCL (e.g., Bonardi, Graham, Hall, & Mitchell, 2005; Le Pelley et al., 2007; Le Pelley & McLaren, 2003), in which there is an explicit change in context between the first phase of the experiment in which predictiveness is established and the second phase in which the impact of this predictive history on novel learning is assessed. In contrast, in many studies of learned predictiveness effects in animals the same (or very similar) cues and outcomes occur throughout the experiment (e.g., Holland, 1984; Mackintosh, 1969, 1973). This raises the possibility that effects of predictive history observed in human learning rely on, or are in some way influenced by, the change of context occurring before the critical learning phase. Perhaps, for example, this change in task signals that the cues are now to be involved in different relationships and therefore leads participants to generalize their previous learning about the predictiveness of the different cues in a way that would not occur if such explicit evidence of a change were not provided. In order to test this suggestion, in Experiment 2 there was no change in context between the two phases of the experiment. That is, both stages involved a six-choice SRT task with the same stimuli, and therefore participants were given no indication that the structure of the task had changed in any way.

While it is at least theoretically possible that a change in context might be required to generate an effect of predictive history on novel learning, it seems unlikely that this would be the case. Studies that have investigated this issue

systematically in both animals and humans typically find that a change in context will weaken, rather than enhance, the effect of predictive history (e.g., Lovibond, Preston, & Mackintosh, 1984; Nelson & Sanjuan, 2006). Such findings have intuitive plausibility—the greater the difference between the two phases of an experiment, the less likely participants might be to transfer what they have learnt during the first phase to what they are about to learn in the second. To the extent that this applies to the current learning preparation, we would expect, if anything, the removal of an explicit context change in Experiment 2 to enhance the influence of predictive history on novel learning (all other things being equal).

Method

Participants, apparatus, and stimuli

A new sample of 16 Cardiff University undergraduates participated for course credit or payment. All apparatus and stimuli were identical to those used in Experiment 1.

Sequence generation

All locations were used as cue and outcome elements. In describing the generation of the sequence we use as an example elements 1, 2, and 3 to denote good predictors and elements 4, 5, and 6 to denote poor predictors. Note, however, that for each participant all six cue elements (three good predictors and three poor predictors) were randomly assigned to the locations 1–6 of the stimulus array.

Stage 1. Table 2 shows the conditional probabilities for the sequence used in Stage 1. Positions 1, 2, and 3 were good predictors of their respective outcomes, as all of these positions predicted the location of the target on the following trial with relatively high probability (.9). In contrast, positions 4, 5, and 6 were poor predictors as the position of the target on the following trial could occur in one of two positions with equal probability (.5).

The possible outcomes that could follow each of the good predictor cues were always themselves poor predictor cues; each poor predictor cue location acted as a high-probability outcome (.9) for one good predictor and a low-probability outcome (.1) for a different good predictor. Similarly, the possible outcomes that could follow each of the poor predictor cues were themselves always good predictors; each good predictor cue location acted as a medium-probability outcome (.5) for two different poor predictors. This allowed us to effectively double the number of trials of each contingency as compared to Experiment 1, in which half the trials involved transitions from outcome positions (1 and 6) to cue positions (2–5), which were not analysed.

Table 2. *Conditional probabilities of stimulus transitions for Stages 1 and 2 of Experiment 2*

	Trial N (cue)											
	Stage 1						*Stage 2*					
Trial N + 1 (outcome)	*1*	*2*	*3*	*4*	*5*	*6*	*1*	*2*	*3*	*4*	*5*	*6*
1					.5	.5		.1	.8	.1		
2				.5		.5	.8		.1		.1	
3				.5	.5		.1	.1				.8
4	.9		.1					.8			.1	.1
5	.1	.9							.1	.8		.1
6		.1	.9				.1			.1	.8	

Note: Blank probability = 0.

Stage 2. Table 2 shows the conditional probabilities for the sequence used in Stage 2. Comparing the conditional probabilities for Stages 1 and 2 in Table 2 it is clear that all Stage 2 transitions were different from those used in Stage 1. As in Stage 2 of Experiment 1, all of the locations now predicted one location with a probability of .8 and two other locations with a probability of .1 each.

Procedure

The procedure was identical to that used in Experiment 1, with the exception that as Stage 2 also used a six-choice task, the notice displayed during the rest break following Block 10 in Experiment 1 was omitted. Participants did not receive any indication that the task would change in any way during the experiment. In order to assess whether the change between stages was particularly salient, at the end of the experiment participants were informed that the movement of the target had been sequenced and were asked: "If you did notice a pattern, did you feel this pattern changed during the experiment? If so, at what point did it change and in what way?" Participants typed their answers into a text box on the screen.

Results

Trials were excluded on the same basis as in Experiment 1. One participant produced a mean RT of 1,230 ms (median RT of 1,094 ms). Given that the majority of this participant's data would have been excluded on our RT criterion, this participant was excluded from all analyses presented here. For the remaining participants, mean error rates were 3.1% ($SD = 1.7$) and 4.6% ($SD = 2.8$) for Stages 1 and 2, respectively. As in Experiment 1, trials were excluded if RTs were less than 100 ms (0.12% of trials in Stage 1; 0.11% in Stage 2) or greater than 1,000 ms (3.70% in Stage 1; 3.29% in Stage 2).

Trials of interest in Stage 1 were high-probability outcomes following good predictor cues (GPH), low-probability outcomes following good predictor cues (GPL), and medium-probability outcomes following poor predictor cues (PPM). Since more trials now contributed to each variable, there were no missing data in the current experiment. Figure 3A shows the RT data for Stage 1. Participants' responses reflected the differing probabilities of the three outcomes: fastest for GPH trials, slowest for GPL trials, and of intermediate speed for PPM trials. These data were subjected to repeated measures ANOVA with factors of outcome probability (high, medium, and

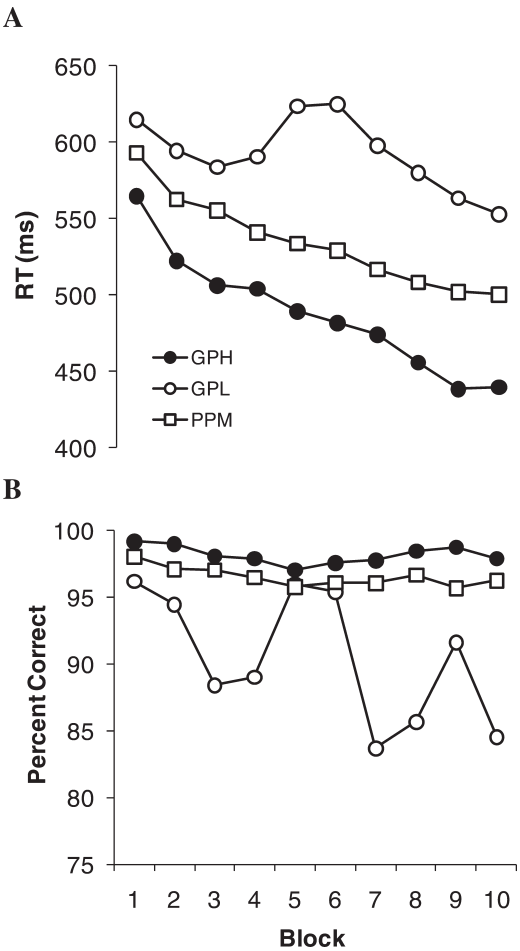

Figure 3. *(A) Reaction time (RT), and (B) accuracy data for Stage 1 of Experiment 2. GPH: responses to high-probability outcomes following good predictor cues; GPL: responses to low-probability outcomes following good predictor cues; PPM: responses to medium-probability outcomes following poor predictor cues.*

low) and block, which revealed significant effects of outcome probability, $F(2, 28) = 23.10$, $MSE = 17{,}947.07$, $p < .001$, and block, $F(9, 126) = 17.98$, $MSE = 1{,}996.76$, $p < .001$, and a significant interaction, $F(18, 252) = 4.15$, $MSE = 1{,}063.49$, $p < .001$. Pairwise comparisons between the three levels of the outcome probability variable revealed significant differences in RT between all three: GPH versus GPL trials, $F(1, 14) = 29.45$, $MSE = 56{,}068.28$, $p < .001$; PPM versus GPL trials, $F(1, 14) = 23.13$, $MSE = 22{,}095.97$, $p < .001$; GPH versus PPM trials, $F(1, 14) = 11.01$, $MSE = 29{,}518.18$, $p < .01$.

Figure 3B shows the accuracy data for Stage 1. In line with the findings in RTs, participants were most accurate on GPH trials, least accurate on GPL trials, and of intermediate accuracy for PPM trials. ANOVA revealed a significant effect of outcome probability, $F(2, 28) = 16.77$, $MSE = 145.86$, $p < .001$, and block, $F(9, 126) = 3.33$, $MSE = 42.08$, $p < .01$, and a significant interaction, $F(18, 252) = 2.89$, $MSE = 40.37$, $p < .001$. Pairwise comparisons between the three levels of the outcome probability variable revealed significant differences in response accuracy between all three: GPH versus GPL trials, $F(1, 14) = 19.61$, $MSE = 450.27$, $p < .01$; PPM versus GPL trials, $F(1, 14) = 14.25$, $MSE = 381.66$, $p < .01$; GPH versus PPM trials, $F(1, 14) = 9.45$, $MSE = 43.21$, $p < .01$.

It is clear from both the RT and accuracy data that participants were sensitive to the varying levels of predictiveness of the cues. Although Experiment 2 had fewer participants than Experiment 1, the significance levels achieved in most comparisons were greater. This is almost certainly due to an increase in the number of cue-outcome pairings, resulting in greater training of the Stage 1 contingencies, as well as a reduction in the variance due to an increase in the number of sampled trials per data point.

The RT data for Stage 2, shown in Figure 4A, were analysed as in Experiment 1. As previously, trial types GPH and GPL refer to high- and low-probability outcomes, respectively, following cues that were pretrained as good predictors during Stage 1. Similarly, trial types PPH and PPL refer to high- and low-probability outcomes, respectively, following cues that were pretrained as poor predictors during Stage 1. These data were subjected to repeated measures ANOVA with factors of prior predictiveness (good vs. poor), outcome probability (high vs. low), and block. The main effect of outcome probability was significant, $F(1, 14) = 10.45$, $MSE = 2{,}628.55$, $p < .01$, indicating faster RTs to high- than to low-probability outcomes. There was no main effect of prior predictiveness, $F < 1$, nor of block, $F(1, 14) = 2.13$, $MSE = 1{,}795.69$, $p = .17$. Importantly, there was a significant interaction between prior predictiveness and outcome probability, $F(1, 14) = 5.60$, $MSE = 2{,}269.39$, $p < .05$, which indicates that the difference in RT between high- and low-probability outcomes was greater for the good predictor contingencies than for the poor predictor contingencies. The interaction between outcome probability and block was significant, $F(1, 14) = 15.86$, $MSE = 351.70$, $p < .01$, which indicates that overall learning was greater in Block 2 than in Block 1. The interaction between prior predictiveness and block was not significant, $F < 1$, nor was the three-way interaction, $F(1, 14) = 1.29$, $MSE = 549.83$, $p = .28$.

Figure 4C plots RT data as learning scores (RTs on low-probability trials minus RTs on high-probability trials) for good and poor predictor contingencies. These data were subjected to ANOVA with factors of prior predictiveness and block. This revealed a significant effect of prior predictiveness, $F(1, 14) = 5.60$, $MSE = 4{,}538.79$, $p < .05$, and of block, $F(1, 14) = 15.85$, $MSE = 703.90$, $p < .01$, but no significant interaction, $F(1, 14) = 1.29$, $MSE = 1{,}099.66$, $p = .28$. Planned tests of simple effects revealed a significant effect of cue both in Block 1, $F(1, 14) = 4.82$, $MSE = 4{,}030.36$, $p < .05$, and in Block 2, $F(1, 14) = 4.60$, $MSE = 1{,}608.09$, $p < .05$.

Figure 4B shows accuracy data for Stage 2. These data were again subjected to ANOVA with factors of prior predictiveness, outcome probability, and block. This revealed a main effect of outcome probability, $F(1, 14) = 5.38$, $MSE = 57.75$, $p < .05$, which indicates that

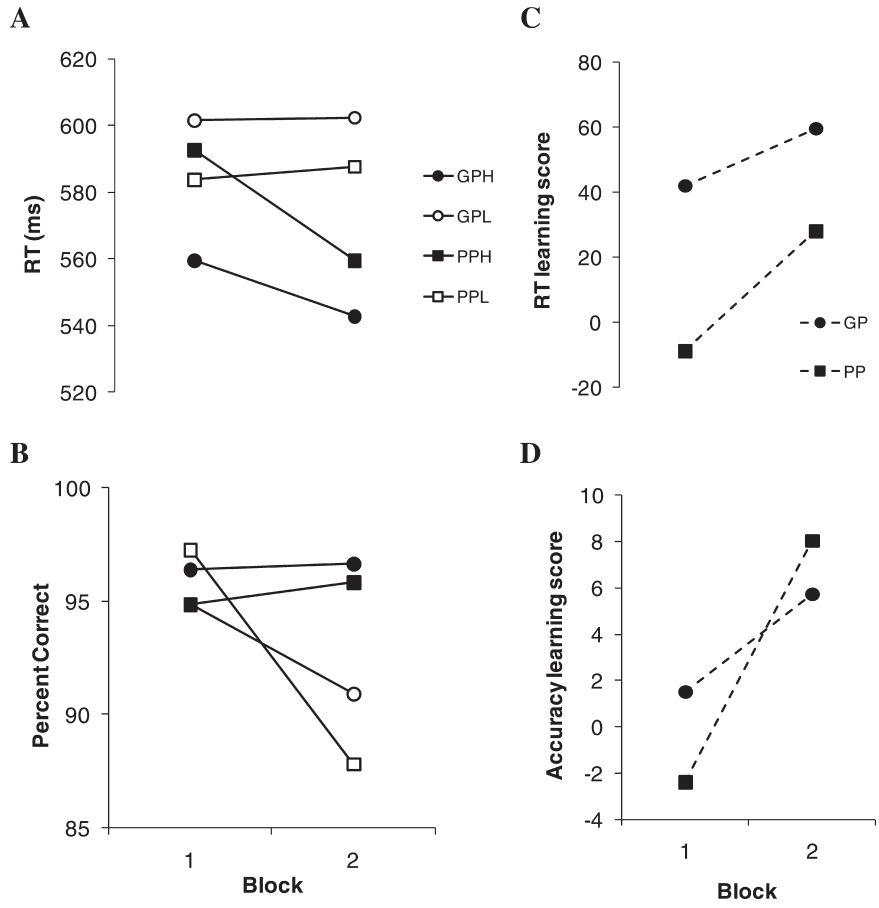

Figure 4. *(A) Reaction time (RT), and (B) accuracy data for Stage 2 of Experiment 2, with data plotted as learning scores—the difference between responses to high- and low-probability outcomes—in Panels (C) and (D), respectively. GPH: responses to high-probability outcomes following cues pretrained as good predictors in Stage 1; GPL: responses to low-probability outcomes following cues pretrained as good predictors in Stage 1; PPH: responses to high-probability outcomes following cues pretrained as poor predictors in Stage 1; PPL: responses to low-probability outcomes following cues pretrained as poor predictors in Stage 1.*

participants were more accurate on high-probability outcomes than on low-probability outcomes. There was also a main effect of block, $F(1, 14) = 9.47$, $MSE = 29.56$, $p < .01$, indicating that accuracy decreased from Block 1 to Block 2, driven largely by a decrease in accuracy on low-probability outcomes. There was a significant interaction between outcome probability and block, $F(1, 14) = 6.49$, $MSE = 61.87$, $p < .05$, indicating stronger evidence for learning in Block 2 than in Block 1. Prior predictiveness did not exert a significant main effect or interact with any other variable, maximum $F(1, 14) = 1.51$, $p = .24$.

Figure 4D shows learning scores for accuracy (accuracy on high-probability outcomes minus accuracy on low-probability outcomes). These data were subjected to ANOVA, which revealed no main effect of cue, $F < 1$, a significant main effect of block, $F(1, 14) = 6.47$, $MSE = 123.72$, $p < .05$, and no interaction, $F(1, 14) = 1.51$, $MSE = 96.04$, $p = .24$. Tests of simple effects

found no significant effect of cue in either block, maximum $F(1, 14) = 2.07$, $p = .17$.

Of the 15 participants, only 5 reported that they felt the sequence changed in some way during the experiment. Of these participants, 2 indicated that they felt the sequence changed halfway through the experiment, 1 that it changed every block, 1 that it changed before the last four blocks, and 1 that it "changed back to random", but they did not specify when. These data suggest that the change in the sequence structure for the last two blocks of the experiment went unnoticed by many of, if not all, the participants.

Discussion

The results of Experiment 2 are similar to those of Experiment 1: The rate at which sequence learning proceeded for a given cue was dependent on the predictive history of that cue. The removal in Experiment 2 of any explicit change in context between Stages 1 and 2 indicates that such a change is not necessary for an effect of predictiveness to be observed. Indeed, verbal report data suggest that the change in the task structure between Stages 1 and 2 went unnoticed by most, if not all, participants. Whilst it might be argued that participants would not have kept a running count of the block number, since the questions appeared immediately after Stage 2, it seems unlikely that participants would not have been able to report that the task changed during the last two blocks if this change had been noticed.

Unlike in Experiment 1, the facilitation in Stage 2 learning for those cues trained as good predictors in Stage 1 was evident for the duration of Stage 2. As suggested earlier, it is possible that the longer lived influence of prior predictiveness in Experiment 2 was a consequence of the lack of an explicit context change as compared to Experiment 1 (cf. Lovibond et al., 1984; Nelson & Sanjuan, 2006) leading to greater transfer between the two stages in Experiment 2. An alternative possibility, however, makes reference to the fact that in Experiment 2 all of the positions in the stimulus array acted as both cues and outcomes. Consequently, participants experienced twice as many presentations of both the good and poor predictor contingencies in Stage 1 as they did in the Stage 1 procedure used in Experiment 1 (see Method). It is therefore possible that this difference resulted in the apparently greater effect of predictiveness on Stage 2 learning observed in Experiment 2. However, our current results do not allow us to decide between these alternatives.

Although we observed greater learning about good predictor cues over poor predictor cues in both Experiments 1 and 2, comparison of Figure 2A with Figure 4A reveals that the pattern of data is slightly different between the two demonstrations of this effect. In Experiment 1, the effect was driven entirely by slower RTs to low-probability outcomes following good predictor cues than to those following poor predictor cues during Block 1 of Stage 2. However, in Experiment 2 the effect was driven largely by faster RTs to high-probability outcomes following good predictor cues than following poor predictor cues. The most likely explanation for this difference in the pattern of results between experiments is a floor effect in RTs to high-probability outcomes in Experiment 1. In Stage 2 of Experiment 1 the outcome positions were removed from the screen, which we had hoped would reduce the impact of direct interference from Stage 1 associations on learning of, and responding to, the new cue-outcome contingencies of Stage 2. This seems to have been successful: Figures 1A and 2A show that RTs to high-probability outcomes at the end of Stage 1 (438 ms) are similar to RTs to the new high-probability outcomes in Stage 2 (427 ms). In addition, reducing the number of possible target locations will reduce response competition and hence tend to produce faster RTs. Although we see a small decrease in RTs to high-probability outcomes across Stage 2 training, it is likely that floor effects are masking any potential RT benefit for high-probability outcomes following good predictor cues. In other words, while it is likely that RTs for a four-choice SRT task can improve beyond 427 ms, we would suggest that the

likelihood of observing differences between high-probability outcomes at this level is impaired by floor effects.

Participants in Experiment 2 received twice as many cue-outcome pairings in Stage 1 as did those in Experiment 1, resulting in more robust learning of the Stage 1 contingencies. It is likely that this would have led to greater interference on Stage 2 responding, while the continued use of a six-choice SRT task in Stage 2 will also maintain a high level of response competition. Consequently we would expect to see a general slowing of RTs at the outset of Stage 2, and this was indeed observed in the results of Experiment 2—there is a clear increase in RT to high-probability outcomes between the final block of Stage 1 (Figure 3A; 440 ms) and the first block of Stage 2 (Figure 4A; 560 ms). Given that Stage 2 RTs have moved away from floor levels, this allows scope to observe an advantage for good predictor cues over poor predictor cues in RTs to high-probability outcomes.

The data from Experiments 1 and 2 are consistent with recent findings in human contingency learning studies employing analogous designs (e.g., Le Pelley et al., 2007; Le Pelley & McLaren, 2003; Le Pelley, Oakeshott, Wills, & McLaren, 2005b). Le Pelley (2004) has suggested that "learned predictiveness" effects such as these can be accommodated by the associative model proposed by Mackintosh (1975; see also Kruschke, 2001). Within this model, learning about each cue is modulated by a cue-specific associability parameter (sometimes referred to as an attentional parameter), which influences the extent to which a cue is able to engage the learning process. On each trial, the associability of a cue changes as a function of that cue's predictiveness, with consistent predictors maintaining a higher associability than inconsistent predictors.

It is easy to see how this model could account for the results of Experiments 1 and 2. During Stage 1 those cues trained as good predictors will maintain high associability, whilst the associability of the poor predictors will fall. Consequently learning about those stimuli that were previously trained as good predictors will proceed more rapidly during Stage 2, with the difference in learning rate dependent on the extent to which the associability of these cues has diverged during Stage 1.

In the next section we examine the extent to which the simple recurrent network (SRN; Elman, 1990)—an associative model of sequence learning—can accommodate the results of Experiment 2. In particular, following the models of Mackintosh (1975) and Le Pelley (2004), we examine whether modifying the SRN to include cue-specific learning rate parameters improves its ability to capture the influence of prior predictiveness on sequence learning.

Simple recurrent network simulation of Experiment 2

The SRN (Elman, 1990) has been shown to be an accurate model of human sequence learning and is able to capture robust effects in the artificial grammar learning and SRT paradigms (e.g., Cleeremans, 1993; Cleeremans & McClelland, 1991). The SRN is a multilayer connectionist network that is trained with the back-propagation algorithm (Rumelhart, Hinton, & Williams, 1986) to minimize the error between "network-produced" and "target" output patterns. In the case of sequence learning, on each trial the network is presented with a pattern across the input units: One input unit is "turned on" to represent the position of the target stimulus on trial N. The pattern of activation across the output units is taken as the network's prediction as to the position of the target stimulus on trial $N + 1$, which is evaluated against the actual position on trial $N + 1$. The characteristic that sets the SRN apart from other multilayer connectionist models is the recurrent loop, which allows the network to copy its internal representation (the current activation pattern across the hidden units) onto a set of "context" units, which are then fed back into the network at the next time step. In doing so, the model uses not only the information from the present trial in order to learn, but

an integrated representation of several previous trials.[1]

Our network functioned in much the same way as that described by Elman (1990): a back-propagation multilayer network with a recurrent loop. On each trial, the activation of one input unit was set to 1, and all other input units were set to 0. The activation was then fed forward to the hidden units by calculating the sum of all products of input activations and their respective connection strengths with each hidden unit. Thus, for hidden unit h:

$$in_h = B_h + \sum_{i=1}^{I} w_{hi} \cdot a_i \quad (1)$$

where in_h is the input for hidden unit h; B_h is the bias associated with hidden unit h; w_{hi} is the weight of the connection between hidden unit h and input unit i; a_i is the activation of input unit i; and I is the total number of input units. This input is then transformed into an activation value for hidden unit h, by the activation function given in Rumelhart et al. (1986):

$$a_h = \frac{1}{1 + e^{-in_h}} \quad (2)$$

The input to, and activation of, the output units is calculated in much the same way as that of the hidden units, such that for output unit o:

$$in_o = \sum_{h=1}^{H} w_{oh} \cdot a_h \quad (3)$$

$$a_o = \frac{1}{1 + e^{-in_o}} \quad (4)$$

where in_o is the input for output unit o; w_{oh} is the weighted connection between output unit o and hidden unit h; a_h is the activation of hidden unit h; and H is the total number of hidden units in the network.

When applied to sequence learning, the target output on each trial is the next element in the sequence. The accuracy of the model in selecting the next element was calculated as the activation of the target output unit divided by the total activation of all output units, commonly referred to as the Luce choice ratio (hereafter LCR; Luce, 1959).

Target values for "active" and "inactive" stimulus positions were set at .9 and .1, respectively. Although only one target stimulus position (output unit) is "active" at any one time, values of 1 and 0 cannot be reached without infinitely large weights and so effectively cannot be achieved (Rumelhart et al. 1986). As the model learns the sequence it will get better at predicting the appropriate output for each input, and hence LCR values should approach 1 (corresponding to high activation of the target output unit and low activation of the nontarget outputs). We might expect that LCR will be inversely related to RT; hence figures showing SRN performance are plotted using (1 − LCR).

Following each response made by the network, the error of each output unit is back-propagated through the network to update the weights

[1] The context loop in the SRN allows the model to learn sequences containing cue-outcome contingencies that span several intervening elements (see Cleeremans, 1993), but since the sequences used in the current experiments are created from exclusively first-order transitions, prima facie this functionality of the model might seem redundant. However, with respect to sequence learning, the SRN has received more attention than any other model and therefore seems the most appropriate model to apply to these data. Moreover, including context units will only provide a model with greater flexibility, thus allowing for a better assessment of the ability of a model that does not allow variable cue processing to predict our empirical data. In fact, our parameter search included parameter sets with very low learning rates (e.g., .01) for context–hidden unit connections. Thus situations in which the possible contribution of the context units to learning is minimized (i.e., situations in which the SRN will behave in a manner similar to a standard back-propagation network; Rumelhart, Hinton, & Williams, 1986) form a subset of our simulation data. Finally, although second-order information is no more useful than first-order information for the learning of these sequences, this is not to say that second-order information is not present in these sequences. For instance, given the transitions in Table 2, sequences such as 1–2–4 (GPH transition followed by GPH transition) will be more common than sequences such as 5–2–4 (PPL transition followed by GPH transition).

between each layer of units. Error terms for output and hidden units were calculated as follows:

$$\delta_o = (t_o - a_o) \cdot (1 - a_o) \cdot a_o \quad (5)$$

$$\delta_h = \left(\sum_{o=1}^{o=O} \delta_o \cdot w_{oh} \right) \cdot (1 - a_h) \cdot a_h \quad (6)$$

where δ_o and δ_h refer to the error on output unit o and hidden unit h, respectively. These errors were then used to update the weights and biases in the network using the generalized delta rule (Rumelhart et al., 1986):

$$\Delta W_{oh} = LR_{oh} \cdot \delta_o \cdot a_h \quad (7)$$

$$\Delta W_{hi} = LR_{hi} \cdot \delta_h \cdot a_i \quad (8)$$

$$\Delta W_{hc} = LR_{hc} \cdot \delta_h \cdot a_c \quad (9)$$

$$\Delta B_h = LR_b \cdot \delta_h \quad (10)$$

where LR denotes a learning rate parameter. Note that in our implementation of the model, independent learning rate parameters were used for the three sets of weights and for the hidden unit biases (LR_{oh}, LR_{hi}, LR_{hc}, and LR_b). Given that the aim of these simulations was to establish whether a "standard" SRN model (that is, an SRN that does not contain an additional associability parameter) could feasibly capture the effects of predictive history observed in Experiments 1 and 2, we wished to specify as few constraints on the model as possible. Hence, if it is possible, by whatever means, for this model to account for the effects of predictive history observed empirically, this approach of providing the greatest possible flexibility gives the best chance of detecting this ability.

The model was examined using the sequences from Experiment 2 as this produced a significant predictiveness effect across the whole of Stage 2. In searching for parameters that would produce the observed pattern of data we varied learning rate values (which could be set at .01, .05, .1, .3, .5, or .7) and the number of hidden units (3, 15, 30, or 50). Given that each of the four independent learning rates within the network could take one of six values, the simulation space produced 5,184 parameter sets. For each of these sets, the model was trained using the exact sequences given to the 15 experimental participants of Experiment 2, in 100 separate simulated experiments, resulting in 1,500 simulated participants for each parameter set. For each simulated participant, all weights within the network were initialized with random values between −.5 and .5.

The model's performance was evaluated by examining the average performance across the 100 simulated experiments for each parameter set. For each set we first assessed whether learning of Stage 1 contingencies had been successful. As in the human data, the following trials were of interest: high-probability outcomes following good predictors (GPH); low-probability outcomes following good predictors (GPL); medium-probability outcomes following poor predictors (PPM). If the network produced the ordinal relationship shown in the empirical data (i.e., higher LCR for GPH than for PPM, and higher LCR for PPM than for GPL), then it was considered to have successfully learnt Stage 1. In total 237 of the 5,184 parameter sets failed to meet this criterion and were excluded from further analysis.

For each of the remaining 4,947 sets, learning scores (LCR on high-probability outcomes minus LCR on low-probability outcomes) for good and poor predictor cues were compared across the two blocks of Stage 2. This revealed that 1,702 parameter sets mirrored our empirical data, showing greater learning about good predictor cues than poor predictor cues.

Although these results seem to suggest that the SRN is potentially able to produce the desired direction of effect, further analysis tested the ability of the model to produce a veridical match to the pattern of data observed in Stage 2. In the empirical data of Experiment 2 (see Figure 4A), RTs for responses following poor predictor cues were within the range of those following good predictor cues. That is, participants

were fastest to GPH trials, were slowest on GPL trials, and were of intermediate speed on PPH and PPL trials. However, only some of these differences between the four trial types reached significance. RTs on GPH outcomes were significantly faster than those on PPH, PPL, and GPL outcomes, whilst RTs on PPH outcomes were significantly faster than those on GPL outcomes, all $t(14)$s > 2.62, $ps < .05$. Therefore a parameter set was assessed as having produced the correct ordinal prediction overall if it made the correct ordinal predictions with respect to these differences considered individually (i.e., GPH < PPH, GPH < PPL, GPH < GPL, and PPH < GPL).

For each parameter set a comparison of the Stage 2 learning scores for good and poor predictor cues was made using a *t* test on the data from the 15 participants in each simulated experiment. The average of the *t*-values from the 100 simulated experiments provided a measure of the "robustness" of the effect produced by each parameter set. Table 3 shows, for the standard SRN, the proportion of parameter sets showing the general predictiveness effect of greater learning about good predictors than poor predictors and the proportion showing the correct ordinal predictions, as a function of mean *t*-value. It is clear that, for the standard SRN, very few parameter sets produce a "robust" effect of greater learning about good predictors than poor predictors. The critical *t*-value for a paired test with 15 participants is 2.145; of those parameter sets that do generate a predictiveness effect with $t > 2$, almost none are able to reproduce the ordinal patterns seen in the empirical data. In fact, of the 50 parameter sets producing the largest Stage 2 effects of prior predictiveness (those producing $t > 2$), 47 produced higher LCR values on GPL trials than on PPH trials. In other words, in the most successful simulations with the standard SRN, the ordinal relationships between the trial types produced a main effect of cue, with poorer performance to outcomes following poor predictor cues than to outcomes following good predictor cues (an effect that was not observed in either experiment). Figure 5 shows the results of a simulation with the SRN using parameters typical of those producing the largest effects of prior predictiveness.

Learned predictiveness effects analogous to those observed in Experiments 1 and 2 have previously been taken as evidence of the operation of cue-processing mechanisms, which modulate the associability of cues (see Le Pelley, 2004). It seems likely, therefore, that an SRN that incorporates cue-processing mechanisms would be better able to capture the pattern of results observed in our empirical data.

Since only input-to-hidden weights are directly connected to cues (input units) within the network, in order for a cue's associability to directly modulate the amount of learning that accrues to it we allow associability to modulate

Table 3. *Stage 2 simulation results for the standard SRN and the alpha SRN broken down by the robustness of the produced effect*

	Proportion of parameter sets matching empirical predictiveness effect		*Proportion matching empirical ordinal pattern of data*	
Mean t-value	*Standard*	*Alpha*	*Standard*	*Alpha*
> 1	8.0	18.8	4.3	15.2
> 2	1.0	4.8	0.1	4.3
> 3	0	0.9	0	0.9

Note: SRN = simple recurrent network. Proportions are expressed as the percentage of the total number of parameter sets producing the Stage 1 ordinal pattern of data observed in Experiment 2 (SRN: 4,947; alpha-SRN: 23,752). *t*-values are the result of a comparison of the Stage 2 learning scores (Luce choice ratio for high-probability outcomes minus Luce choice ratio for low-probability outcomes) for good and poor predictor cues, across the 15 participants of each simulated experiment. Mean *t*-values are an average of the 100 simulated experiments for each parameter set.

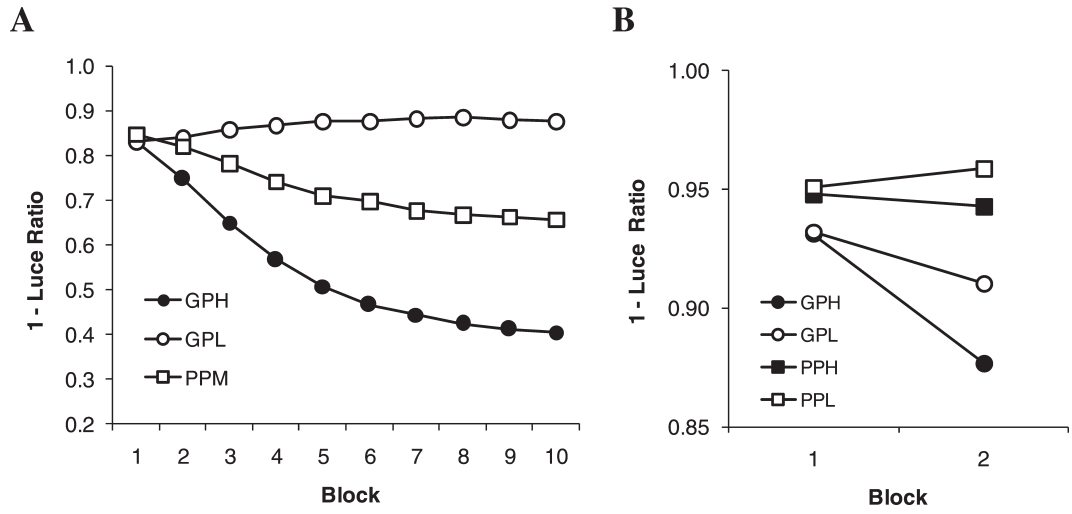

Figure 5. *Example data from a standard simple recurrent network (SRN) simulation of Experiment 2, with parameters $LR_{hi} = .3$; $LR_{ob} = .5$; $LR_{bc} = .01$; $LR_b = .3$; 50 hidden units. (A) (1 − Luce choice ratio) across Stage 1 training. GPH: responses to high-probability outcomes following good predictor cues; GPL: responses to low-probability outcomes following good predictor cues; PPM: responses to medium-probability outcomes following poor predictor cues. (B) (1 − Luce choice ratio) across Stage 2 training. GPH: responses to high-probability outcomes following cues pretrained as good predictors in Stage 1; GPL: responses to low-probability outcomes following cues pretrained as good predictors in Stage 1; PPH: responses to high-probability outcomes following cues pretrained as poor predictors in Stage 1; PPL: responses to low-probability outcomes following cues pretrained as poor predictors in Stage 1.*

weight changes only at this level. The learning rule for weight change between input and hidden units in the standard SRN (Equation 8) was modified to incorporate associability as follows:

$$\Delta W_{hi} = \alpha_i \cdot LR_{hi} \cdot \delta_h \cdot a_i \qquad (11)$$

where α_i is the associability of cue i. In previous models incorporating cue-processing mechanisms (e.g., Kruschke, 2001; Mackintosh, 1975), the change in a cue's associability is determined by the extent to which that cue predicts the outcome compared to the predictive value of all other cues present on that trial—that is, the predictiveness of a given cue is assessed *relative to* that of other simultaneously presented cues. However, since only one cue is present on screen at any one time in the SRT task, changes in associability cannot be made on the basis of direct cue comparison in the current model. Consequently, we based changes in associability on the *absolute* predictiveness of a cue—the model would increase a cue's associability if that cue predicted the correct target on the next trial and would decrease associability if that cue predicted an incorrect outcome. In line with previous cue-processing mechanisms (Le Pelley, 2004), the magnitude of associability change on a trial was governed by the prediction error of the model, in this case using the LCR:

$$\begin{aligned} &\text{If correct then } \alpha_i = \alpha_i + (LCR)^4 \\ &\text{If incorrect then } \alpha_i = \alpha_i - (LCR)^4 \end{aligned} \qquad (12)$$

Values of α were allowed to vary between lower and upper limits of .1 and 1. Raising the LCR to the fourth power results in small changes in α on each learning cycle, ensuring a gradual approach to these limits. In addition to controlling parameter values as for the standard SRN simulations, in the "alpha SRN" simulations we also examined various starting values of α (.1, .2, .3, .4, and .5), yielding a parameter space of 25,920 sets. All

other procedural and analytical aspects of these simulations were as for the standard SRN.

Of the 25,920 parameter sets, 2,168 failed to produce the Stage 1 ordinal pattern of results and were excluded from further analysis. Of the remaining 23,752 parameter sets, a total of 12,144 produced the observed Stage 2 predictiveness effect of greater learning about good predictors than poor predictors. Table 3 shows the proportion of parameter sets showing the Stage 2 predictiveness effect and the correct ordinal predictions as a function of the robustness of effect, for the alpha SRN. The alpha SRN produces over four times as many robust predictiveness effects (4.8% of the parameter sets) as the standard SRN, and in contrast to the standard SRN the parameter sets yielding the most robust predictiveness effects almost invariably also produce the correct ordinal predictions.

Figure 6 shows detailed results for one particular parameter set using the alpha SRN. Figure 6A shows that Stage 1 learning in this simulation proceeds as observed empirically; Figure 6B shows how α values for good and poor predictors diverge over the course of this training. By the end of Stage 1 the mean α for good predictor cues has reached an asymptotic level of 1, as these cues consistently predict the outcome location on the next trial. In contrast, poor predictor cues are equally likely to be followed by one of two possible outcomes. Consequently the model will make the correct prediction on 50% of trials at best. As a result, α values for these cues will be subject to fluctuating positive and negative adjustments and hence will rise less rapidly.[2] Figure 6C shows Stage 2 performance for the alpha SRN, which is able to capture the correct ordinal pattern of results (cf. Figures 4A).

Although the alpha SRN performs considerably better than the standard SRN, the model fails to predict greater Stage 2 learning about good predictors than about poor predictors for almost half of the parameter sets tested. Further analysis revealed that these failures are largely a consequence of undifferentiated α values following Stage 1 training: There was a significant positive correlation between the difference in α at the end of Stage 1 (α for good predictor cues minus α for poor predictor cues) and the size of the Stage 2 effect, $r(23, 752) = .57$, $p < .001$. This correlation is shown in Figure 7. This pattern is particularly striking if one compares those parameter sets showing an effect in the direction of greater Stage 2 learning about good predictors than about poor predictors (the 4,467 sets with $t > 1$) with those showing an effect in the opposite direction (the 983 sets with $t < -1$). For the former, the average difference between the mean α values for good predictor and poor predictor cues at the end of Stage 1 was .34, while for the latter it was only .05. Thus the simulations that failed to show the observed pattern of Stage 2 learning tended to be those in which α values for good predictors failed to rise above those of poor predictors during Stage 1.

In summary, simulation work demonstrated that the SRN is able to produce the effects of prior predictiveness observed empirically, albeit in only a small number of parameter sets. However, the success of the model was improved considerably when the model incorporated cue-specific associability parameters to capture the predictive histories of cues in the task. Further work will be required to refine these mechanisms, in order to understand how best to accommodate the influence of prior predictiveness on learning in multilayer networks.

GENERAL DISCUSSION

Two experiments examined the extent to which the predictive history of a cue can modulate the rate at

[2] Since poor predictor cues consistently predict two different outcomes during Stage 1, we would expect the associability of these cues to decrease as often as it increases. However, whilst this is true, since LCR values will be greater for correct than incorrect predictions, positive changes in associability will always be greater than negative changes, and hence we would expect to observe gradual increases in the associabilities of these cues during this stage.

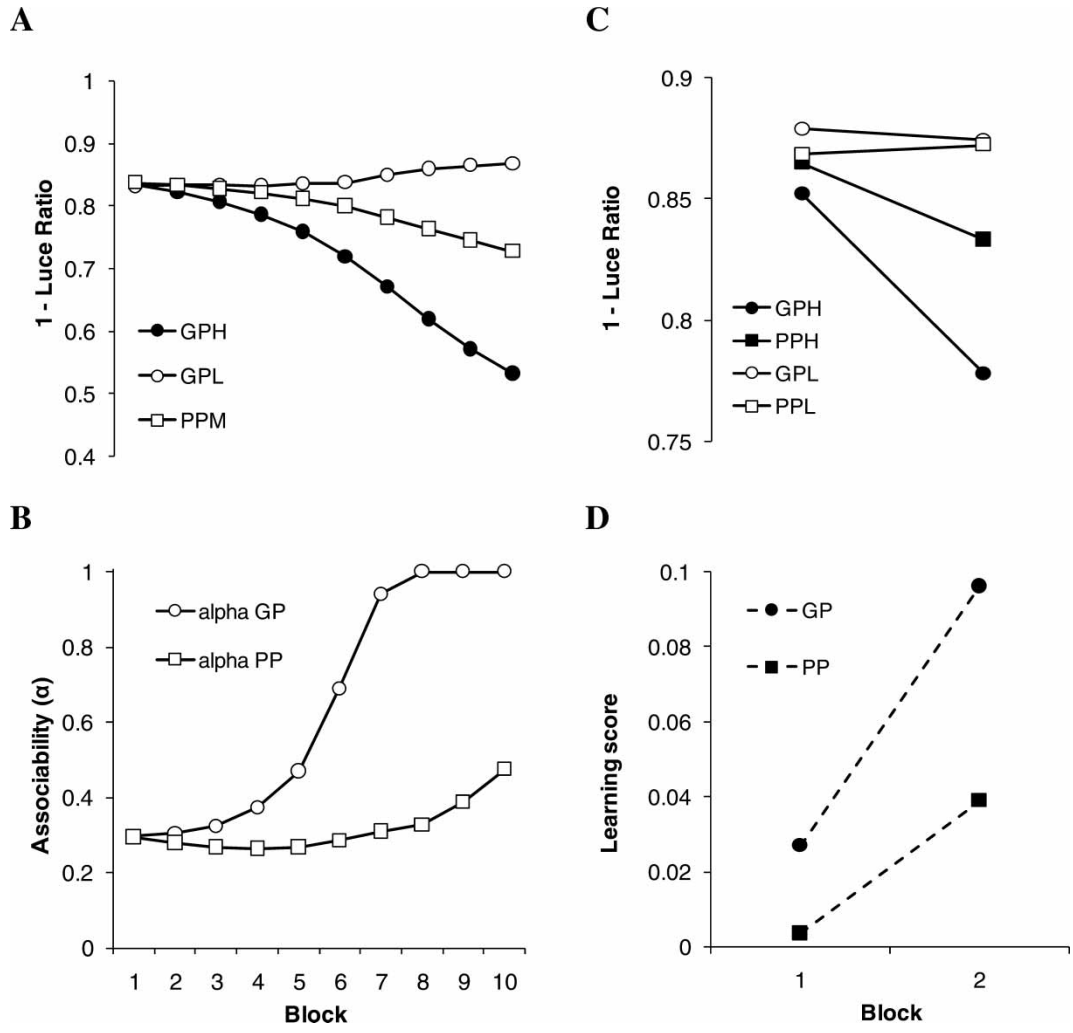

Figure 6. *Example data from an alpha simple recurrent network (SRN) simulation of Experiment 2, with parameters $LR_{hi} = .7$; $LR_{oh} = .3$; $LR_{bc} = .1$; $LR_{b} = .1$; starting value of $\alpha = .3$; 15 hidden units. (A) (1 − Luce choice ratio) across Stage 1 training. GPH: responses to high-probability outcomes following good predictor cues; GPL: responses to low-probability outcomes following good predictor cues; PPM: responses to medium-probability outcomes following poor predictor cues. (B) Alpha values across Stage 1 training. GP: good predictor cues; PP: poor predictor cues. (C) (1 − Luce choice ratio) across Stage 2 training. GPH: responses to high-probability outcomes following cues pretrained as good predictors in Stage 1; GPL: responses to low-probability outcomes following cues pretrained as good predictors in Stage 1; PPH: responses to high-probability outcomes following cues pretrained as poor predictors in Stage 1; PPL: responses to low-probability outcomes following cues pretrained as poor predictors in Stage 1. (D) The Stage 2 data expressed as learning scores.*

which that cue is subsequently learnt about during a sequence learning task. In both experiments the predictiveness of a subset of cues was manipulated during an initial pretraining phase, after which new sequenced transitions involving these cues were created, and learning about these new transitions was examined. Analysis of learning at the subsequence level revealed that those cues that had previously been established as good predictors were learnt about more readily than those cues previously established as poor predictors. This effect was found in Experiment 1, where each stimulus

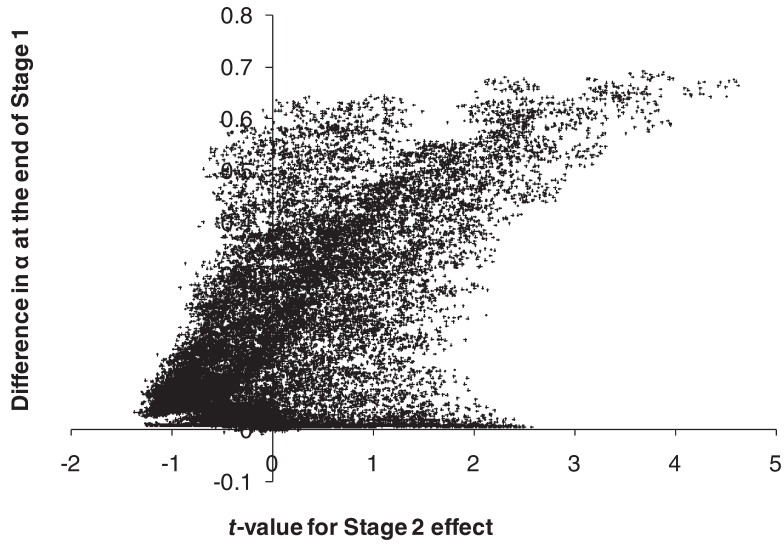

Figure 7. *The magnitude of the Stage 2 effect as a function of the difference in α for good and poor predictor cues at the end of Stage 1, for simulations with the alpha simple recurrent network (SRN). Positive values on the α difference scale reflect a greater average α for good predictor cues over the average for poor predictor cues. Positive t-values reflect greater learning during Stage 2 about cues pretrained as good predictors over cues pretrained as poor predictors.*

location was designated as either a cue or an outcome during pretraining, and in Experiment 2, where all stimulus locations functioned as both cues and outcomes throughout. The effect of predictive history on new learning is rather short-lived in the context of a sequence learning experiment. However, it is worth noting that within a single Stage 2 block, each cue-outcome pairing was presented, on average, 30 times in Experiment 1 and 20 times in Experiment 2 (this difference was due to the former being a four-choice task, while the latter was a six-choice task). The amount of Stage 2 training in each experiment is therefore far greater than that used in analogous HCL demonstrations of these associability effects—for example, Le Pelley and McLaren (2003) gave four presentations of each cue-outcome pairing in Stage 2.

We might consider the transitions pertaining to poor predictor cues in Stage 1 of each experiment to be "ambiguous", in the sense that each cue was followed by two likely outcomes, whilst the transitions pertaining to good predictor cues were somewhat "unique", in the sense that there was a high probability of one outcome occurring after these cues. Cohen et al. (1990; see also Curran & Keele, 1993) found that learning sequences of ambiguous transitions required attentional processing, whilst learning sequences of unique transitions could be achieved under conditions of attentional load. The current results extend these findings by indicating that differences in the ambiguity of the transitions following cues can themselves lead to a change in the attentional processing that participants devote to those cues. Furthermore these differences in attentional processing can differentially influence subsequent learning about those cues even when all cues are now involved in equally ambiguous transitions.

Computational simulation revealed that while the simple recurrent network (Elman, 1990) was, under certain conditions, able to produce a predictiveness effect consistent with the empirical data, such predictions were typically not robust. Moreover, for those few sets of parameters that did produce robust predictiveness effects, the network almost invariably failed to reproduce the ordinal patterns observed in our data. However,

a modification of this model that included a cue-specific associability parameter produced a far higher proportion of simulations showing robust effects in line with our empirical findings, and the vast majority of these parameter sets also produced ordinal predictions matching our data. It is of course possible that other modifications of the standard SRN not based on associability would also allow the resultant model to explain our findings. However, given the conceptual similarity of our studies of predictive history on cue learning in the SRT task to earlier investigations of cue-processing effects in human and animal learning (e.g., Le Pelley & McLaren, 2003; Mackintosh, 1973; see Le Pelley, 2004, for a review), we believe that an account based on differences in cue processing (as implemented by associability) deserves consideration.

A key question in human sequence learning is the extent to which learning engages selective attentional mechanisms. Previous work has focused on whether learning of concurrent sequential information can occur despite these cues having no necessary benefit for the learning of the primary task (e.g., Jimenez & Mendez, 1999; Mayr, 1996; Rowland & Shanks, 2006). The data presented here support the suggestion that selectional cue-processing mechanisms can also act at a "within-sequence" level, having a differential effect on learning about different elements of the same sequence, depending on the predictive history of those elements.

Several previous studies of HCL in more traditional, nonsequential paradigms have also claimed to demonstrate cue-processing effects in learning (e.g., Bonardi et al., 2005; Kruschke & Blair, 2000; Le Pelley et al., 2007; Le Pelley & McLaren, 2003). The current experiments extend this existing research by demonstrating associability effects in a preparation in which participants receive no directed instruction to learn—they are asked to respond to, rather than predict, the outcome on each trial. That is, the instructions involved in this task provide participants with no reason or incentive to engage in strategic, hypothesis-testing processes, or to intentionally exploit any knowledge of the cue-outcome contingencies that they may possess. Furthermore, the SRT paradigm involves a rapid mode of stimulus presentation and response, with average reaction times generally well below 600 ms, limiting the time available for participants to engage in higher order reasoning processes. We would argue that the observation of associability effects under such conditions is at least consistent with the idea that these effects have a more automatic, lower level basis, although more research is required to verify whether this is, indeed, the case.

Learned predictiveness in multiple- versus single-cue designs

In all previous demonstrations of learned predictiveness effects in human contingency learning, manipulations of predictiveness have involved the use of compound stimuli, comprising two or more individual cues, in the pretraining stage (e.g., Bonardi et al., 2005; Griffiths & Le Pelley, 2009; Le Pelley et al., 2007; Le Pelley & McLaren, 2003; Whitney & White, 1993). Typically one of the cues of the compound is arranged to be an accurate predictor of the outcome, whilst another is arranged to be a poor predictor. This provides the opportunity for stimulus comparison—that is, a relative assessment of which of the available cues is the best predictor of the outcome. Previous models designed to produce cue-processing effects state that associability is explicitly based on such a comparison process (Krushcke, 2001; Mackintosh, 1975), acting to select between several alternative candidate cues.[3]

The importance of stimulus comparison is a central part of Le Pelley's (2004) hybrid model of associability effects. He noted that, in studies of animal learning, pretraining involving multiple simultaneously presented cues tends to lead to faster learning about good predictors than about

[3] Even in animal studies of learned irrelevance that ostensibly only involve a single conditioned stimulus (e.g., Mackintosh, 1973), the standard analysis assumes a comparison of predictiveness between this conditioned stimulus and the experimental context, with the latter operating essentially as an additional, simultaneously presented cue.

poor predictors (as anticipated by the model of Mackintosh, 1975). In contrast, pretraining of animals with single cues (in which case there can be no direct stimulus comparison to determine the more predictive stimulus) tends to give faster learning about poor predictors than about good predictors (Hall & Pearce, 1979; Swan & Pearce, 1988; Wilson, Boumphrey, & Pearce, 1992). This finding is more consistent with the model of associability offered by Pearce and Hall (1980), which in some sense can be considered as the opposite of the approach offered by Mackintosh.

The SRT task provides a clear example of a single-cue learning paradigm: Only a single stimulus is presented on each trial, and participants are required to respond to that stimulus regardless of its predictive status. We can therefore be certain that participants were engaged in the processing of, and the active responding to, the stimulus on each trial. Hence, drawing a parallel with studies of animal conditioning (and on the basis of Le Pelley's, 2004, hybrid model) we might expect to find faster learning about poor predictors than about good predictors in Stage 2. The fact that the opposite result was observed, with faster learning about good predictors than poor predictors, presents a clear discrepancy between learned predictiveness effects in animals and humans (and by extension suggests that the hybrid model might be best confined to accounting for associability effects in animal, rather than human, learning). The demonstration of a learned predictiveness effect in a single-cue task such as the SRT suggests that in these situations changes in cue processing might be driven more by the absolute predictiveness of the cue (i.e., the extent to which a cue is a good predictor of outcomes) than its relative predictiveness (i.e., the extent to which a cue is a better predictor than other available cues). Consistent with this suggestion, in our alpha SRN simulations associability was determined by the absolute predictiveness of each cue considered individually, as defined by the LCR.

While our experiments represent, to the best of our knowledge, the first evidence of learned predictiveness effects favouring learning of good predictors over poor predictors in a design involving presentation of only a single cue on each trial, we are not the first to suggest a discrepancy between learned predictiveness effects in humans and animals. Latent inhibition refers to the phenomenon of retarded conditioning to a previously nonreinforced stimulus observed in studies of animal conditioning (see Lubow, 1989). It is an effect that occurs in single-cue learning procedures and follows naturally from the Pearce and Hall (1980) model of associability. While reports of latent inhibition are ubiquitous in animal conditioning, analogous procedures in humans typically do not yield similar effects—observing latent inhibition in humans typically requires the use of a masking task during preexposure, which permits explanation of the resulting effect in terms of alternative processes (Graham & McLaren, 1998). More generally this represents another discrepancy that is consistent with (although does not prove) the suggestion that single-cue learning procedures do not engage the same associability processes in humans and animals.

Associability and attention

In the current tasks, we measured changes in associability by examining the effect of prior predictiveness on the rate of subsequent learning about cues. This raises the question of how best to characterize this change in associability. One possibility is that the associability of a cue is simply a learning rate—that is, that the effect of prior predictiveness is restricted to modulating only the rate of learning about a cue. An alternative possibility is that the observed changes in associability result from changes in attention to cues (see e.g., Kruschke, 2001). That is, participants might pay more attention to cues that have been consistent predictors in the past than to those that have been inconsistent. We need only assume then that attention influences the rate of learning about a cue, in order to explain the influence of prior predictiveness on learning rate observed empirically. Importantly, on this attentional account the influence of prior predictiveness

need not be restricted to learning rate: If attention to the cues has changed, then this might also be expected to influence other aspects of the processing of these cues.

Given that the current experiments measure only the rate of learning about cues, our data do not allow us to decide between these alternatives. Future experiments will address this issue by providing a more direct measure of visual attention during this task, using eye-tracking equipment. If the advantage for good predictor cues does indeed reflect greater attention to these cues, then we should observe a bias in eye gaze towards good predictor cue locations.

CONCLUSION

In conclusion, this paper offers, to the best of our knowledge, the first experimental evidence that the readiness with which a cue will develop a mapping to a novel outcome can be modulated by the predictive history of that cue in an incidental sequence learning task. Computational simulations demonstrate that this pattern of results can be accommodated by the associative framework of a modified SRN, in which cue-specific learning rate parameters modulate associative sequence learning. It remains possible that other modifications to multilayer networks will also allow them to account for the influence of prior predictiveness on cue processing in incidental sequence learning. Future work will investigate this possibility and will subsequently investigate how such an influence is best characterized.

Original manuscript received 17 July 2008
Accepted revision received 3 February 2009
First published online 18 June 2009

REFERENCES

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

Bonardi, C., Graham, S., Hall, G., & Mitchell, C. (2005). Acquired distinctiveness and equivalence in human discrimination learning: Evidence for an attentional process. *Psychonomic Bulletin & Review*, *12*, 88–92.

Cleeremans, A. (1993). *Mechanisms of implicit learning: Connectionist models of sequence processing.* Cambridge, MA: MIT Press.

Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: News from the front. *Trends in Cognitive Sciences*, *2*, 406–416.

Cleeremans, A., & McClelland, J.L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*, 235–253.

Cohen, A., Ivry, R. I., & Keele, S. W. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 17–30.

Curran, T., & Keele, S. W. (1993). Attentional and nonattentional forms of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 189–202.

De Houwer, J., Vandorpe, S., & Beckers, T. (2005). On the role of controlled cognitive processes in human associative learning. In A. J. Wills (Ed.), *New directions in human associative learning* (pp. 41–63). Mahwah, NJ: Lawrence Erlbaum Associates.

Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence from the process dissociation procedure. *Psychonomic Bulletin & Review*, *8*, 343–350.

Dickinson, A., Shanks, D., & Evenden, J. (1984). Judgments of act–outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology*, *36A*, 29–50.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179–211.

Graham, S., & McLaren, I. P. L. (1998). Retardation in human discrimination learning as a consequence of pre-exposure: Latent inhibition or negative priming? *Quarterly Journal of Experimental Psychology*, *51B*, 155–177.

Griffiths, O., & Le Pelley, M. E. (2009). Attentional changes in blocking are not a consequence of lateral inhibition. *Learning & Behavior*, *37*, 27–41.

Hall, G., & Pearce, J. M. (1979). Latent inhibition of a CS during CS–US pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, *3*, 31–42.

Hogarth, L., Dickinson, A., Austin, A., Brown, C., & Duka, T. (2008). Attention and expectation in human predictive learning: The role of uncertainty. *Quarterly Journal of Experimental Psychology*, *61*, 1658–1668.

Holland, P. C. (1984). Unblocking in Pavlovian appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 476–497.

Jiménez, L., & Méndez, C. (1999). Which attention is needed for implicit sequence learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 236–259.

Jiménez, L., Méndez, C., & Cleeremans, A. (1996). Comparing direct and indirect measures of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 948–969.

Kruschke, J. K. (2001). Towards a unified model of attention in associative learning. *Journal of Mathematical Psychology*, *45*, 812–863.

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

Kruschke, J. K., Kappenman, E. S., & Hetrick, W. P. (2005). Eye gaze and individual differences consistent with learned attention in associative blocking and highlighting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 830–845.

Le Pelley, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *Quarterly Journal of Experimental Psychology*, *57B*, 193–243.

Le Pelley, M. E., Beesley, T., & Suret, M. B. (2007). Blocking of human causal learning involves learned changes in stimulus processing. *Quarterly Journal of Experimental Psychology*, *60*, 1468–1476.

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

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

Le Pelley, M. E., Oakeshott, S. M., Wills, A. J., & McLaren, I. P. L. (2005b). The outcome-specificity of learned predictiveness effects: Parallels between human causal learning and animal conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *31*, 226–236.

Le Pelley, M.E., Suret, M. B., & Beesley, T. (in press). Learned predictiveness effects in humans: A function of learning, performance or both? *Journal of Experimental Psychology: Animal Behavior Processes*.

Lovibond, P. F. (2003). Causal beliefs and conditioned responses: Retrospective revaluation induced by experience and by instruction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 97–106.

Lovibond, P. F., Preston, G. C., & Mackintosh, N. J. (1984). Context specificity of conditioning and latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 360–375.

Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*, 3–26.

Lubow, R. E. (1989). *Latent inhibition and conditioned attention theory.* Cambridge, UK: Cambridge University Press.

Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.

Mackintosh, N. J. (1969). Further analysis of the overtraining reversal effect. *Journal of Comparative and Physiological Psychology*, *67*, 1–18.

Mackintosh, N. J. (1973). Stimulus selection: Learning to ignore stimuli that predict no change in reinforcement. In R. A. Hinde & J. S. Hinde (Eds.), *Constraints on learning* (pp. 75–96). London: Academic Press.

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

Mayr, U. (1996). Spatial attention and implicit sequence learning: Evidence for independent learning of spatial and nonspatial sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 350–364.

Nelson, J. B., & Sanjuan, M. C. (2006). A context-specific latent inhibition effect in a human conditioned suppression task. *Quarterly Journal of Experimental Psychology*, *59*, 1003–1020.

Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*, 1–32.

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

Reber, A. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. New York: Oxford University Press.

Rescorla, R. A. (1999). Within-subject partial reinforcement extinction effect in autoshaping. *Quarterly Journal of Experimental Psychology*, *52B*, 75–87.

Rowland, L. A., & Shanks, D. R. (2006). Attention modulates the learning of multiple contingencies. *Psychonomic Bulletin & Review*, *13*, 643–648.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1: Foundations*. Cambridge, MA: MIT Press.

Shanks, D. R. (2005). Implicit learning. In K. Lamberts & R. Goldstone (Eds.), *Handbook of cognition* (pp. 202–220) . London: Sage.

Shanks, D. R. (2007). Associationism and cognition: Human contingency learning at 25. *Quarterly Journal of Experimental Psychology*, *60*, 291–309.

Shanks, D. R., Rowland, L. A., & Ranger, M. R. (2005). Attentional load and implicit sequence learning. *Psychological Research*, *69*, 369–382.

Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioural and Brain Sciences*, *17*, 367–447.

Squire, L. R. (1994). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 203–231). Cambridge, MA: MIT Press.

Swan, J. A., & Pearce, J. M. (1988). The orienting response as an index of stimulus associability in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, *4*, 292–301.

Whitney, L., & White, K. G. (1993). Dimensional shift and the transfer of attention. *Quarterly Journal of Experimental Psychology*, *46B*, 225–252.

Wilkinson, L., & Shanks, D. R. (2004). Intentional control and implicit sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 354–369.

Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 1047–1060.

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

Wilson, P. N., Boumphrey, P., & Pearce, J. M. (1992). Restoration of the orienting response to a light by a change in its predictive accuracy. *Quarterly Journal of Experimental Psychology*, *44B*, 17–36.

APPENDIX

Instructions preceding Experiment 1

"The aim of this study is to examine the effect of practice on motor control. In this task you are required to follow a grey circle as it moves between six positions on the screen. The six positions are situated across the middle of the screen and are represented by six smaller grey circles. Each of the six positions corresponds to a key on the keyboard, these keys are X, C, V, B, N and M, along the bottom row of the keyboard. X is used to respond to the far left position, M the far right, and the others for the positions in-between. Each time the larger circle changes position you are required to press the appropriate key to identify its new location. Once you have pressed a key, the circle representing that position will turn red to indicate where you have responded. Should you respond incorrectly you will hear a beep in the headphones. After you have responded the larger grey circle will disappear and reappear in a new location. In summary, your task is to follow the circle as it moves between the six positions. Each time the stimulus moves we would like you to respond as fast and as accurately as you can. In particular, we want you to avoid making errors in this task. The participant who performs the best over the course of the experiment will win a £10 prize. The experiment is split into 12 blocks, each of which lasts for approximately 3 minutes. At the end of each block you will be given a rest of 15 seconds before the next block starts. During this break a countdown will appear to show you when the next block of trials will start."