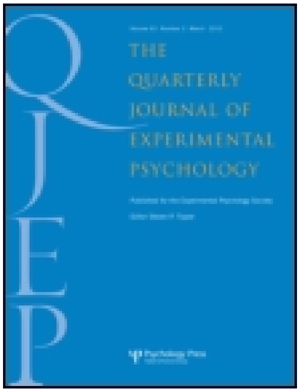

The Quarterly Journal of Experimental Psychology
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/pqje20

Evidence that a transient but cognitively demanding process underlies forward blocking

Pei-Pei Liu [a] & Christian C. Luhmann [a]

[a] Department of Psychology , Stony Brook University , Stony Brook , NY , USA

Accepted author version posted online: 09 Aug 2012.Published online: 14 Sep 2012.

THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY, 2013
Vol. 66, No. 4, 744–766, http://dx.doi.org/10.1080/17470218.2012.717952

Routledge
Taylor & Francis Group

Evidence that a transient but cognitively demanding process underlies forward blocking

Pei-Pei Liu, and Christian C. Luhmann

Department of Psychology, Stony Brook University, Stony Brook, NY, USA

Blocking is a learning phenomenon in which prior experience inhibits learning about novel cues. Though the phenomenon itself has been well documented, the details of blocking-related processes still remain contentious. Two experiments investigated whether participants were engaged in demanding cognitive processing during different portions of a standard blocking paradigm. Participants in Experiment 1 engaged in a simple secondary task while completing a standard blocking procedure. Results showed that performance on the secondary task was briefly diminished early in the second phase of the blocking paradigm, when the novel cue is first paired with the pretrained cue. Participants in Experiment 2 performed a difficult cognitive load task during either the early or the late portions of the second phase of blocking. The blocking effect was eliminated when learners were under load early in the second phase, but not when learners were under load late in the second phase. These results suggest that blocking relies on a cognitively demanding process with a distinct time course. Computational simulations illustrate how a model that includes top-down (i.e., cognitively demanding) attentional modulation can reproduce the observed behaviour. This suggests that purely associative processes are not sufficient to explain the observed behaviour. Implications for current accounts of blocking are discussed.

Keywords: Forward blocking; Cognitive load; Contingency learning; Causal learning.

When trying to attribute a cause to an event, learners often face a dilemma in which more than one candidate is present simultaneously. One way to deal with this problem is to rely on previous experiences in which those candidate causes were encountered in isolation. For example, suppose that one day you take a new vitamin along with your coffee and feel alert afterwards. You are very likely to attribute the alertness to coffee but not to the new vitamin. This is likely because of the sheer number of times in which you have felt alert after having coffee in the past. This phenomenon is known as blocking and has been intensively studied in the field of learning (for a review, see Pineno & Miller, 2007). In a blocking paradigm, participants first learn that the presence of a single event (Cue A) is always followed by a certain outcome. Subsequently, Cue A is paired with another novel event (Cue X), and this pair is followed by the same outcome. Despite the fact that Cue X was always followed by an outcome, learners do not come to strongly associate the

Correspondence should be addressed to Pei-Pei Liu, Department of Psychology, Stony Brook University, Stony Brook, NY 11794–2500, USA. E-mail: peiliu@notes.cc.sunysb.edu

two. Learning of the association between Cue X and the outcome is thus said to be blocked by Cue A.

Blocking has been observed in nonhuman animals (e.g., Beckers, Miller, De Houwer, & Urushihara, 2006; Kamin, 1969; Rescorla & Wagner, 1972) as well as in humans (e.g., Beckers, De Houwer, Pineno, & Miller, 2005; De Houwer & Beckers, 2003; Dickinson, Shanks, & Evenden, 1984; Kruschke & Blair, 2000). At least three accounts provide explanations for how blocking occurs: traditional associative accounts, attentional accounts, and a recently proposed inferential account. Classic associative models, like that proposed by Rescorla and Wagner (1972), were designed specifically to account for blocking effects. Attentional theories extend these traditional models by suggesting that attentional processes can help to account for blocking, as well as many other cue competition effects. Most recently, an inferential account has been proposed (De Houwer & Beckers, 2003), which argues that cue competition effects are the result of sophisticated, nonassociative reasoning.

Traditional associative models

Because blocking was first observed in conditioning experiments using nonhuman animals (Kamin, 1969), theories of conditioning have been developed to explain this phenomenon. Dickinson et al. (1984) also argued that, because humans exhibit blocking, human contingency learning, like conditioning in animals, is associative in nature. According to one of the most influential associative accounts, the Rescorla–Wagner model (Rescorla & Wagner, 1972), learning only takes place when learners' expectations are violated. According to the Rescorla–Wagner model (RW), when first encountering a novel cue, the learner's predictions will be relatively inaccurate and lead to large prediction errors. These prediction errors will, in turn, induce large changes in the associative strength between the cue and the outcome. As the learner repeatedly encounters the same cue–outcome pairings, the predictions will become more accurate, and the changes in the associative strength will become smaller. In the blocking paradigm, the association between Cue A and the outcome gradually increases throughout Phase 1. At the end of Phase 1, the cue predicts the presence of the outcome with no prediction error. When Cue A is subsequently paired with Cue X in Phase 2, learners have already learned to fully expect the presence of the outcome based on Cue A alone. Thus, according to RW, there should be no prediction error even when the outcome is first observed following the AX pair. Because there is no prediction error, no learning about the cues take place at this point and the association between Cue X and the outcome does not increase even though Cue X is reliably followed by the outcome.

Attentional theories

Whereas the RW model suggests that blocking occurs because there is nothing to be learned, other associative models suggest that attention plays a critical role. Attentional theories of learning (e.g., Kruschke, 2001; Mackintosh, 1975a) argue that blocking occurs because learners actively ignore Cue X. More generally, these theories suggest that individuals learn to attend to cues that predict variations in reinforcement and ignore cues that do not predict variations in reinforcement. Ignored cues do not influence learners' predictions, and little (or nothing) is learned about ignored cues.

According to these accounts, as Cue A consistently predicts the presence of the outcome throughout Phase 1 of a blocking paradigm, the association between Cue A and the outcome increases. On the first trial of Phase 2, Cue X has an initial strength of zero and thus does not, on its own, predict the presence of the outcome. Cue A, in contrast, has accumulated significant associative strength during Phase 1 and thus predicts, on its own, the presence of the outcome. Given that the AX pair is followed by the outcome, Cue A provides a more accurate prediction than Cue X. Consequently, the attentional weight assigned to Cue X declines rapidly, which minimizes the influence of Cue X on learners' predictions. Because Cue X is quickly ignored, little is learned about it even though it is repeatedly paired with the outcome.

One unique prediction of the attentional account is that, since the blocking paradigm should cause learners to ignore Cue X, subsequent learning about Cue X should become more difficult. Kruschke and Blair (2000) designed a task to test this prediction. They first presented participants with a traditional blocking procedure in which Symptom A was repeatedly followed by the occurrence of a particular disease, and then the pair of symptoms, A and X, was repeatedly followed by the same disease. Subsequently, another novel cue was added to the AX pair, and the combination of these three cues was then followed by a new, novel disease. Their results showed that participants failed to associate Symptom X with the new, novel disease suggesting that the diminished attentional weight assigned to Symptom X persisted beyond the blocking paradigm per se. According to RW, no learning about the blocked cue happens during the blocking procedure but subsequent learning about the blocked cue should not be affected. Furthermore, Kruschke, Kappenman, and Hetrick (2005) used eye tracking to monitor participants' gaze throughout a traditional blocking task and found that those learners that guided their gaze away from Cue X also exhibited stronger blocking effects. Beesley and Le Pelley (2011) have also employed eye tracking and found that participants gaze less at the blocked cue, not only during blocking itself, but also during subsequent phases of learning. These findings are strong evidence in support of attentional theories in that the blocking procedure appears to cause learners to literally ignore Cue X. Again, RW is unable to account for such data.

Inferential account

Though associative accounts have long been used to explain blocking, other researchers (e.g., De Houwer & Beckers, 2003; Lovibond, 2003) have argued that blocking in humans may not involve associative processes at all. Instead, these researchers suggest that blocking may be a result of sophisticated inferential processes. According to the inferential account, learners often assume that if two cues each predict a certain outcome, then the outcome should be greater when both causes are present than when either appears on its own. As De Houwer and Beckers (2003) suggested, the logic people may apply during blocking is: "If Cue A causes the outcome to occur with a certain intensity and probability, and if adding Cue X does not change either the intensity or probability, this implies that Cue X is not a cause of the outcome." Based on this logic, learners in a blocking procedure can infer that Cue X has no influence because the AX pair is followed by the same outcome as the one that followed Cue A alone. An important theoretical difference between associative processing and inferential processing is that the former is thought to operate in an automatic manner and to not require cognitive resources, while the latter is thought to be analytic, relatively slow, and highly dependent on cognitive resources (Evans, 2003; Evans & Over, 1996; Kahneman, 2003; Sloman, 1996; Stanovich & West, 2000).

Several studies have found evidence in support of the inferential account. For example, Beckers et al. (2005) provided learners with pretraining, which demonstrated that additivity did not hold before participants completed a blocking procedure. In the additive pretraining, two food items each caused an allergic reaction individually, and they caused a stronger allergic reaction when paired together. In the subadditive condition, two food items each caused the same allergic reaction when paired together as they did individually. All participants then completed a standard blocking paradigm. Participants receiving the additive pretraining could apply the logic mentioned above and infer that Cue X does not influence the outcome. Participants receiving the subadditive pretraining, because the assumption of additivity was violated, could not use the rule mentioned above to conclude that Cue X had no influence and thus could not infer that Cue X was not a cause of the outcome. Results indicated that participants in the subadditive condition exhibited weaker blocking than participants in the additive condition. The evidence that people take additivity information into account suggests that blocking requires the assumption of additivity as predicted by the inferential account.

Additional evidence for the inferential account comes from work on the maximality effect. Beckers et al. (2005) manipulated whether the outcome used in a blocking paradigm was of the maximal intensity or not. When the allergic reaction that follows Cue A itself and follows Cue A and X together is of maximal intensity, learners cannot utilize the above rule to reason about Cue X. They cannot be sure whether Cue X has failed to affect the allergic reaction or whether Cue X has actually affected the allergic reaction that was already at ceiling, and thus could not have been made any stronger. This uncertainty would naturally weaken the blocking effect. On the other hand, when the intensity of the allergic reaction is not at ceiling (submaximal), learners can be certain that Cue X does not cause an allergic reaction on its own. Consistent with these predictions, participants in the maximal condition showed weaker blocking effects than those in the submaximal condition. Such evidence suggests that violating the premises of the inferential rule can modulate the traditional blocking effect. These findings are again uniquely consistent with the inferential account.

Another line of work supporting the inferential account of blocking involves the role of cognitive resources. As mentioned earlier, associative processing is assumed to operate automatically and to not require cognitive resources. Inferential processing, on the other hand, is assumed to operate relatively slowly and to be strongly dependent on cognitive resources. Thus, if blocking is the result of solely associative processing, then it should not be affected by the availability of cognitive resources. De Houwer and Beckers (2003) conducted a study in which participants had to perform a concurrent secondary task while completing an otherwise standard blocking paradigm. In the difficult condition, high-pitched and low-pitched tones were presented at random intervals. Participants had to press one of two corresponding buttons depending on which tone was presented. In the easy condition, there was only one tone, and it was presented at a fixed interval (e.g., every 1.2 seconds), and participants only had to press a single key when they heard the tone. De Houwer and Beckers found that blocking was weaker for participants in the difficult condition than for those in the easy condition.

Interestingly, De Houwer and Beckers (2003) reported that the effect of the secondary task was only observed when participants had to perform the difficult secondary task both during the learning sequence and as they made their postlearning judgements. When learners performed the secondary task during the learning sequence but not during the judgement phase, difficulty had no reliable effect on blocking. The authors suggest that this discrepancy may be evidence that the inferential processes responsible for blocking, "can take place both during the learning phase and during the test phase" (De Houwer & Beckers, 2003, p. 355). Thus, when there was no opportunity for learners to engage in the predicted inferential processing, the blocking effect was modulated by the difficulty of the secondary task. However, when the secondary task was only required during the learning phase, inferential processing may have been postponed until later (i.e., until the test phase), which would have allowed learners to exhibit blocking effects. The evidence that blocking depends on the availability of cognitive resources is consistent with the idea that blocking results, at least partly, from inferential processing and cannot be accounted for by either traditional associative models or the more elaborated attentional models.

Investigating the processes underlying blocking

Each of the theories reviewed above provides a different explanation for blocking. However, distinguishing between these different accounts is not always straightforward. For example, standard experimental approaches have learners complete a traditional blocking paradigm, under typical conditions, and evaluate the final product of learning (e.g., standard postlearning judgements). However, in such scenarios, all of the theories reviewed above make exactly the same prediction; they each predict the classic blocking effect. Because of this common behavioural prediction, traditional methods allow only weak inferences about the nature of the underlying learning processes. Nonetheless, research has largely focused on postlearning judgements to evaluate theories of blocking (e.g., Beckers et al., 2005; Chapman & Robbins, 1990; Dickinson et al., 1984; Kruschke & Blair, 2000; Le Pelley, Oakeshott, & McLaren, 2005; Mitchell, Lovibond, Minard, & Lavis, 2006; Shanks, 1985; Waldmann & Holyoak, 1992; Wasserman & Berglan, 1998). Here, we briefly note two facets of blocking that remain elusive.

First, there is little in the way of empirical data addressing the time course of blocking-related processes. This is somewhat surprising because there is good reason to believe that blocking is a dynamic phenomenon, gradually evolving over the course of the learning sequence. As just one example, RW suggests that learners' beliefs about Cue A stabilize by the end of Phase 1 and then remain unchanged throughout the rest of the learning sequence. Additionally, this model predicts that learners' beliefs about Cue X are never updated because the outcome is fully predicted based on Cue A alone. On the contrary, attentional theories suggest that, because learners need to first learn that Cue X is redundant, blocking should gradually develop as Phase 2 progresses. Predictions at this level of granularity have rarely been tested within the literature on nonhuman learning (though some attempts have been made in the animal learning literature, e.g., Mackintosh, 1975b). Two recent exceptions to this tendency (Beesley & Le Pelley, 2011; Wills, Lavric, Croft, & Hodgson, 2007) utilized eye tracking to investigate patterns of overt attention over the course of learning. Consistent with earlier findings (Kruschke et al., 2005), both studies report diminished attention directed toward the redundant Cue X during Phase 2. Interestingly, neither study observed actual shifts in attention; Cue X was essentially ignored from the beginning of Phase 2. This is inconsistent with attentional theories and with Mackintosh's (1975b) finding that animals exhibit no blocking at the beginning (i.e., on the first trial) of Phase 2. Wills et al. (2007) aggregated their results across blocks of trials, which may explain the curious lack of attentional dynamics. That is, participants may have gradually learned to ignore Cue X during the first block itself. Beesley and Le Pelley (2011) did not average over trials in this manner, suggesting that overt attentional measures such as eye tracking do not conform to the predictions of standard attentional learning models. For these reasons, the time course of these attentional shifts remains somewhat ambiguous. In addition, it remains unclear exactly how critical these attentional biases are to blocking. As Beesley and Le Pelley (2011) stated, eye-tracking data suggest that, "the observed bias in learning is *associated with* a bias in attention, but they do not allow us to assess the stronger claim of attentional theories, that biases in learning [are] *caused by* biases in attention" (p. 119, emphasis in original).

Second, the relationship between top-down cognitive processing and cue competition effects such as blocking remains unclear. The inferential rules posited by the inferential account are explicitly assumed to be deliberative. However, the account itself does not specify when such analysis should take place, and data supporting this account (e.g., Beckers et al., 2005; De Houwer & Beckers, 2002a, 2002b, 2003) have failed to address the time course of these cognitively demanding processes. On the other hand, attentional accounts' predictions about deliberative processing are also unclear. At the heart of these models are associative processes that are typically assumed to be automatic (Evans, 2003; Evans & Over, 1996; Kahneman, 2003; Sloman, 1996; Stanovich & West, 2000). However, less is known about the nature of the processes underlying the reported attentional biases. Indeed, Wills et al. (2007) noted that whether attentional biases, "are the top-down result of high-level reasoning processes or the result of the lower-level, automatic processes that are sometimes assumed to be implied by associative accounts, is an important topic for future research" (p. 853).

The current study includes two experiments that are designed to provide complimentary, converging evidence about both the nature and the temporal dynamics of blocking-related processes. In Experiment 1, we attempt to measure the cognitive demands that learners face at various points during a standard blocking paradigm. In Experiment 2, we move beyond simply observing these processes and attempt to disrupt deliberative processing at different points during the learning sequence in order to assess the potential causal role of these processes.

EXPERIMENT 1

In the first experiment, we utilized a concurrent tone discrimination task to monitor the time course of blocking-related processing. The periodic tone discrimination task that we employed was relatively easy and was designed to measure the depth at which learners were processing the main learning task. Participants processing the learning task deeply should take longer to respond to the secondary tone task. Participants processing the learning task more shallowly, on the other hand, should respond to the tones more quickly. For example, Posner and Boies (1971) used this measure to evaluate the processing mechanisms operating as participants judged whether two consecutively presented letters were identical. Responses to auditory probes were faster during the first letter than during the second letter, suggesting that the comparative processing required during the presentation of the second letter consumed more cognitive resources than the initial encoding of the first letter. Liu and Luhmann (2012) also have demonstrated that response times to simple tone discrimination were strongly related to the expectations in a contingency learning task. In one experiment, participants were presented with a sequence of mostly homogeneous trials intermixed with a small number of contradictory trials. For instance, the homogenous trials in one condition presented two events that tended to exhibit strong, positive covariation (i.e., when one event was present, the other was also present; when one event was absent, the other was also absent). The contradictory trials in this condition exhibited the opposite pattern (i.e., one event was present while the other was absent). Participants' responses to the tones were slower immediately after encountering unexpected outcomes relative to when encountering the expected outcomes. This evidence suggests that participants were engaged in deeper, more demanding cognitive processing when they encountered unexpected covariation information.

In the current study, the tone discrimination task will allow us to better evaluate explanations of blocking. For example, recall that traditional associative accounts (e.g., RW) suggest that blocking occurs because the outcome that follows the AX pair is fully expected based on prior learning about Cue A. Thus, the AX+ trials in Phase 2 of a blocking paradigm should be treated exactly as the A+ trials at the end of Phase 1. Therefore, if blocking is solely a result of the associative processes described by RW, we should not see any obvious change in participants' responses as they progress from Phase 1 to Phase 2.

In contrast, attentional theories suggest that blocking results from shifts in attentional allocation. These models predict that blocking should not be observed until individuals learn that Cue X is redundant and shift attention away from the cue. If these attentional shifts rely on both prediction error and more sophisticated cognitive processing (e.g., executive control), then we should expect to see participants' responses be slower during Phase 2. However, if the predicted shifts in attention happen automatically or relatively quickly, then Phase 2 may require no more processing than Phase 1. Alternatively, the inferential account suggests that learners will engage in inferential analysis during the learning task and should thus be slower to respond to the tones when the analytic process is initiated. The first opportunity to engage in the inferential analysis in a blocking paradigm is early in Phase 2 (when the novel Cue X is first encountered). However, as mentioned earlier, the theory does not make concrete predictions about when during learning participants should perform their inferential analysis.

Method

Participants

Thirty-two undergraduate students at Stony Brook University participated in the experiment for partial course credit.

Materials

Participants were asked to learn about how different medications were related to an increase in body temperature. Pictures of medications, each of which was one of 12 different colours, were used to represent the cues/causes. Two pictures of a thermometer were used to represent the outcomes. The temperature on the thermometer was either 98.6 or 100.2 degrees F, representing normal and increased body temperatures, respectively. Medications of different colours were randomly assigned to each role (e.g., Cue A, Cue X, etc.) for each participant.

Table 1 shows the types of trials used in two phases of the learning sequence. In Phase 1, Cue A was always followed by the presence of the outcome, and Cue Z was always followed by the absence of the outcome. In Phase 2, Cue A was paired with Cue X and followed by the outcome. Another pair of cues (C and D, positive controls) was followed by the outcome and was used as a control condition (neither C nor D was ever presented alone). There was also a pair of cues (Z and Y, negative controls) which was followed by the absence of the outcome. By including the negative controls we were able to lower the base rate of the outcome. Specifically, in the current design, the outcome was present on exactly half of the trials in Phase 1 and on exactly two thirds of the trials in Phase 2. Without these negative controls, participants might have reasonably concluded that the outcome was likely to follow any arbitrary cue and would be less likely to exhibit blocking (Sobel, Tenenbaum, & Gopnik, 2004). Each trial type was repeated 12 times.

For the secondary task, tones of three different frequencies were used. The high tone had a frequency of 3,520 Hz, the medium tone had a frequency of 880 Hz, and the low tone had a frequency of 220 Hz. Each tone lasted for 50 ms. Only one third of the trials included a tone. For each participant, one of the tones was randomly selected and was played during four trials of each type, and no tones were played on the other eight trials of each type. The exact trials on which a tone was presented were randomly determined for each participant. For example, the first probed A+ trial might be the first trial of A+ for one participant, but the third trial of A+ for another.

Table 1. *Design of the learning task in Experiment 1*

Condition	*Phase 1: Elemental Training*	*Phase 2: Compound Training*
Blocking	A +	AX +
Positive control		CD +
Negative control	Z–	ZY–

Procedure

Participants were seated in front of a computer and were told that their task was to learn about how different medications relate to increases in body temperature. To do so, they were told that they would be provided with a set of hypothetical medical records, each of which included information about what medications that patient had taken and whether the patient's body temperature had increased or not. After brief instructions, participants were given an opportunity to familiarize themselves with the three tones. During this practice, participants were presented with the three different tones and responded by pressing corresponding keys on the keyboard. This practice continued until participants were able to discriminate between the three tones. Participants then completed a brief sequence of sample learning trials to familiarize them with the task and to allow them practice with making the postlearning judgements about the medications.

On a typical trial (without a tone presented), the medications were presented on the left side of the screen for 1000 ms. The image of the thermometer was then presented on the right side of the screen, beside the medications, for an additional 1,750–2,250 ms. The screen was then cleared, and the trial was over. Participants did not have to make any responses on these trials. Trials on which a tone was to be presented were identical except that a tone would start 750–1,250 ms after the onset of the outcome information. Participants had up to 3 seconds to respond to the tones. If they failed to respond within 3 s, the trial ended, and the experiment automatically moved on to the next trial.

After completing the trials, participants judged which of two given medications was more likely to cause an increase in body temperature. Participants made the judgements on an unmarked visual analogue scale, with each of the two medications associated with one end of the scale. Participants utilized the left and right arrow keys to move a "cursor" along the scale to indicate their judgements. Thus, moving the cursor all the way to one end of the scale would indicate absolute certainty that the associated medication would be more likely to cause increased body temperature. Placing the cursor near the centre of the scale would indicate more uncertainty about which medication is more likely to cause increased body temperature. The most important judgements were the comparisons between Cue X and each of the positive control cues (C and D). If participants were to show blocking, they should judge Cue X to be less likely to cause an increased body temperature than the positive controls. The stronger the blocking effect, the more participants should move the cursor towards the end of the scale representing the positive controls (for similar measures of blocking, see Kruschke & Blair, 2000; Kruschke et al., 2005). Participants also had to select between Cue A and the negative control cues. We included these comparisons so that the critical comparisons would be less obvious to participants and also to ensure that participants paid attention to all of the cues. In order to increase statistical power, each participant completed two runs of the learning task, with different sets of medications in each run.

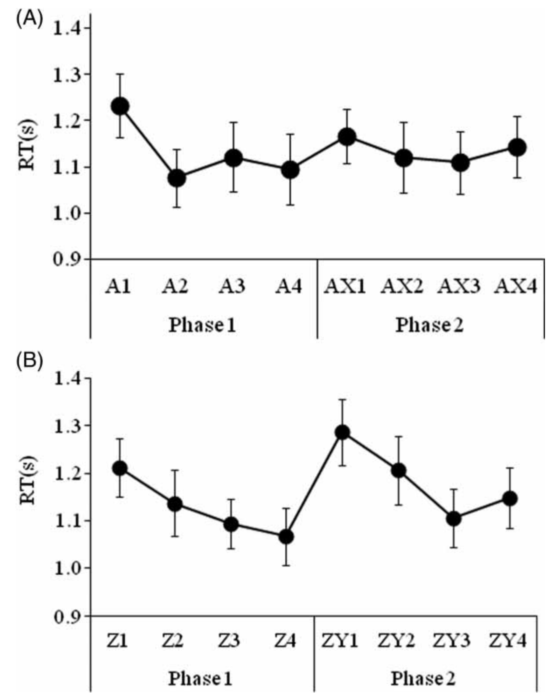

Figure 1. *Response times (RTs) to the secondary tone discrimination task across the learning sequence in Experiment 1. A: Response times for blocking-related trials (A + /AX+) of interest. The RT on the first A+ trial in Phase 1 was slower than the RT on the second A+ trial, probably due to practice effects. B: Response times for control trials (Z–/ZY–). The most obvious effect is that the RT on the last Z– trial in Phase 1 was slower than the RT on the first ZY– trial in Phase 2.*

Results

To evaluate the magnitude of the blocking effect, we turned to participants' comparisons between Cue X and the positive control cues (C and D). If participants believed that the control cues were more likely to cause increases in body temperature than Cue X (e.g., the classic blocking effect), they should have moved the cursor to the side of the scale representing the control cues. We coded the response scale as ranging from –1 to 1, with 1 representing the control cues and –1 representing Cue X. The average judgment was .18 (SD = .49). That is, participants judged Cue X to be marginally less likely than the positive control cues to cause an increased body temperature, $t(31) = 2.04$, $p = .06$.

We next analysed participants' responses to the tones during the learning sequence. Figure 1 shows the response times (RTs) to the tones over the course of the trial sequence, both for the critical A + /AX+ blocking trials and for the Z–/ZY– negative control trials. Recall that 4 of the 12 trials of each type were probed with a tone. In Figure 1, A1 refers to the first probed A+ trial, A2 refers to the second A+ trial that was probed, and so on. A

2 (trial type: A + /AX+ vs. Z–/ZY–) × 2 (order: 1st run vs. 2nd run) × 8 (trial) three-way repeated measure analysis of variance (ANOVA) was conducted. Lower bound correction was applied when necessitated by violations of sphericity in this and all subsequent analyses. Results suggested that neither the main effect nor the interactions with order were significant (ps > .3). To increase statistical power, we thus collapsed across order in subsequent analyses. We conducted a 2 (trial type: A + /AX+ vs. Z–/ZY–) × 8 (trial) two-way repeated measures ANOVA, which revealed a main effect of trial, $F(7, 210) = 3.14$, $p < .01$. The main effect of trial type and the interaction between factors were not significant (Fs < 1). Post hoc analyses showed that the first probed A+ trial in Phase 1 elicited slower RTs than the second A+ trial, $t(31) = 3.78$, $p < .01$. No changes in RT were observed at the Phase 1–Phase 2 transition, consistent with the predictions of RW. Interestingly, post hoc tests revealed that RTs on the last Z– trial in Phase 1 were faster than RTs on the first ZY– trial in Phase 2, $t(31) = 2.93$, $p < .01$. There were no other differences among the responses across time in the two trial types.

To examine whether there was any particular pattern of tone responses associated with the classic blocking effect, we first classified participants into blockers and nonblockers based on whether they judged the positive control cues or Cue X to be more likely to cause increases in body temperature. According to this criterion, we identified 21 blockers and 11 nonblockers within our sample. Looking at Figure 2, all learners appear to exhibit an initial speed-up in which they were slower to respond to the first tone on an A+ trial than to the second tone on an A+ trial. It seems likely that this change can be attributed to simple practice effects. More interestingly, blockers and nonblockers seemed to exhibit different patterns as they transitioned from Phase 1 to Phase 2. Blockers seemed to slow down between the last A+ trial in Phase 1 to the first AX+ trial in Phase 2 (what we refer to as a transitional cost) and speed up between the first to the second AX+ trial in Phase 2 (what we refer to as a transitional recovery). In contrast, nonblockers did not seem to exhibit either of these changes during the Phase 1–Phase 2 transition.

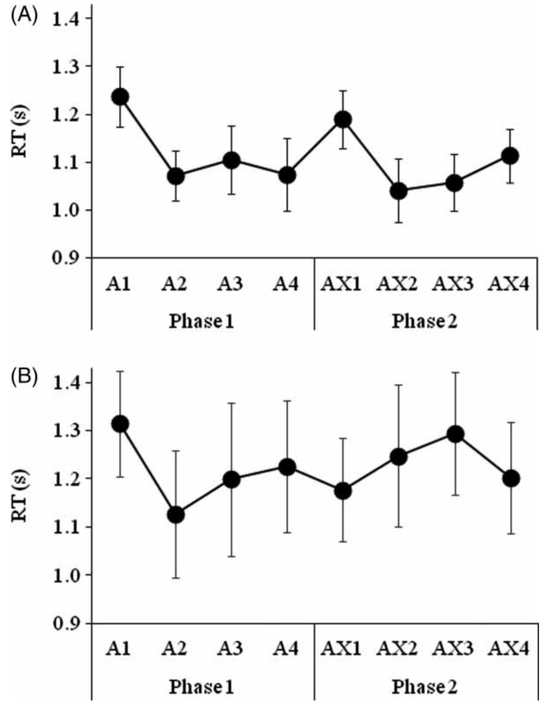

Figure 2. *Response times (RTs) to the secondary tone discrimination task during A+ and AX+ trials for blockers and nonblockers in Experiment 1. A: The blockers' responses were slower on the first AX+ trial in Phase 2 than on the last A+ trial in Phase 1. Their responses on the second AX+ trial were also faster than those on the first AX+ trial. B: Nonblockers showed no such changes as they transitioned from Phase 1 to Phase 2.*

To more rigorously assess this pattern, we evaluated the magnitude of the two different RT effects. First, to quantify the transitional cost, we subtracted each participant's RT on the last A+ trial from their RT on the first AX+ trial. Second, to quantify the transitional recovery, we subtracted each participant's RT on the first AX+ trial from their RT on the second AX+ trial (see Figure 3A). We first conducted a 2 (transitional cost vs. transitional recovery) × 2 (order: 1st run vs. 2nd run) × 2 (subgroup: blocker vs. nonblocker) three-way mixed ANOVA. The results suggested that neither the main effect nor any interactions with order was significant, ps > .1. To increase statistical power, we again collapsed across order.

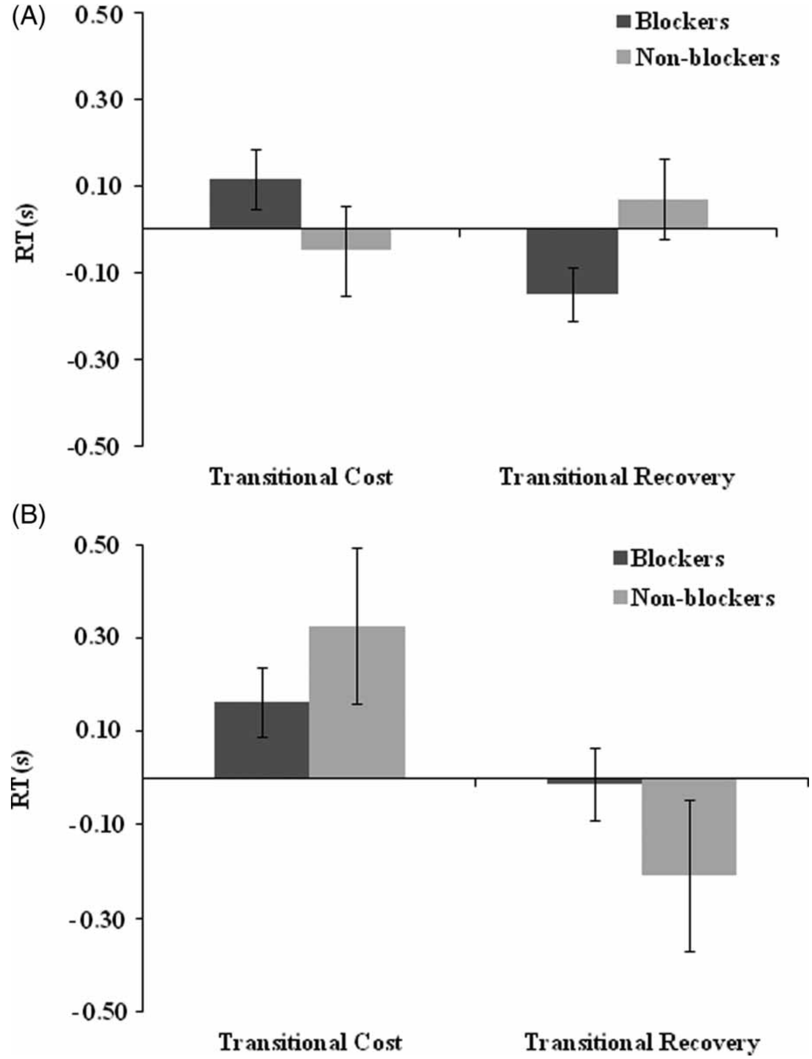

Figure 3. *Response time (RT) effects observed near the Phase 1–Phase 2 transition in Experiment 1. A: Blockers were slower to respond during the first AX+ trial than during the last A+ trial (transitional cost) and faster to respond during the second than during the first AX+ trial (transitional recovery). Nonblockers did not show either of these effects. B: RTs of both blockers and nonblockers were slower during the last Z– trial than during the first ZY– trial. No difference was observed between the first and second ZY– trials. No group differences were observed.*

We conducted a 2 (transitional cost vs. transitional recovery) × 2 (subgroup: blocker vs. nonblocker) mixed ANOVA with repeated measures on the former factor. Neither main effect was significant, *p*s > .4, but the interaction between the two factors was, $F(1, 30) = 4.52$, $p < .05$, confirming the above observation that the two subgroups exhibited different RT patterns as they entered Phase 2 from Phase 1. Post hoc comparisons showed that blockers exhibited a significant transitional cost, one-sample *t* test against zero, $t(20) = 2.13$, $p < .05$, and a significant transitional recovery, $t(20) = 2.34$, $p < .05$. In contrast, nonblockers showed neither the transitional cost ($t < 1$) nor the transitional recovery ($t < 1$).

One potential explanation for these RT effects is that they reflect stimulus novelty rather than anything related to learning itself. For example, perhaps blockers were responding to the novelty of Cue X introduced at the beginning of Phase 2, and perhaps nonblockers were simply not engaged in the task enough to notice this novelty. If the novelty explanation is correct, then we should expect blockers and nonblockers to show similar novelty-related effects in their responses during other trials as well. To explore this possibility, we analysed the pattern of RTs exhibited during the negative control trials. The transition from Z– in Phase 1 to ZY– in Phase 2 represents the same change in novelty as the transition from A+ to AX+. Thus, if the RT differences exhibited by the subgroups were simply due to novelty of Cue X, we should expect the same differences for Z– and ZY–. The RT effects for the Z–/ZY– trials of the two subgroups are depicted in Figure 3B. Just as above, a 2 (RT effect: transitional cost vs. transitional recovery) × 2 (order: 1st run vs. 2nd run) × 2 (subgroup: blocker vs. nonblocker) three-way mixed ANOVA was conducted. The results suggested no significant main effect for order, $F(1, 28) = 2.00$, $p = .168$. Neither of the two-way interactions with order was significant, Fs < 1. The three-way interaction was significant, $F(1, 28) = 7.37$, $p < .05$. Analysing the two runs separately suggested that this three-way interaction was driven by the fact that the two-way interaction between RT effect and subgroup was significant in the first run, $F(1, 30) = 8.57$, $p < .01$, but not in the second run, $F(1, 28) = 1.04$, $p > .3$. However, the pattern observed on Z–/ZY– trials was quite different from the one observed on A + /AX+ trials. In the first run, only the nonblockers showed a pattern of transitional cost ($M = 0.49$, $SD = 0.58$) and recovery ($M = -0.37$, $SD = 0.59$). The blockers showed smaller transitional cost ($M = 0.10$, $SD = 0.27$) and recovery ($M = 0.15$, $SD = 0.50$) than the nonblockers, ts$(30) > 2.6$, ps $< .05$. The finding that blockers and nonblockers differed in their response to the two contingencies suggests that the RT differences observed for the A + /AX+ trials were not merely due to the novelty of Cue X but were instead tied to learning itself.

Though the Z–/ZY– trials control for any novelty-related changes in reaction times, one could instead argue that there might be other reasons that the pattern observed on these trials differs from that observed on blocking-related trials (A + /AX+). For example, one could argue that nonblockers focused more on the contingency of Z–/ZY– than on A + /AX+ and that blockers focused more on the A + /AX+ contingency than on Z–/ZY–. This would explain why blockers exhibited RT effects for the blocking-related trials (but not for the negative control trials) and why nonblockers showed RT effects for the negative controls (but not for the blocking-related trials). That is, there may have been a trade-off between the RT effects for A + /AX+ trials and the RT effects for Z–/ZY– trials. If this were the case, participants' responses to one contingency should predict their responses to the other contingency. To examine this possibility, we conducted an analysis of covariance (ANCOVA), testing the interaction between the RT effects (transitional cost vs. recovery) during A + /AX+ and subgroups (blocker vs. nonblocker) after first removing any variance accounted for by the RT effects during Z–/ZY– trials (both transitional cost and recovery). If the group differences were due to an attentional trade-off between A + /AX+ trials and Z–/ZY– trials as discussed above, the RT effects for A + / AX+ should be accounted for by the RT effects for Z–/ZY–, and this ANCOVA should not reproduce the interaction between the groups and blocking-related RT effects reported earlier. In contrast, results indicated that the interaction between RT effects on A + /AX+ trials and the subgroups was significant even after controlling for the RT effects on Z–/ZY– trials, $F(1, 28) = 6.46$, $p < .02$. This finding suggests that the difference between the two groups in their RT patterns on A + /AX+ trials during early Phase 2 cannot be accounted for simply by their behaviour on the negative control trials. Instead, it seems likely that the transitional cost and recovery effects reflect blocking-related processing.

Finally, to ensure that our post hoc classification of participants into blockers and nonblockers was not artificially producing the effects reported above, we conducted correlation analyses over the entire sample of participants. Specifically, we examined whether the RT changes seen near the Phase

1–Phase 2 transition were associated with the magnitude of the blocking effect exhibited by individual participants. Results revealed that for A+/AX+ trials, larger transitional costs were associated with larger blocking effects ($r = .45$, $p < .05$). Similarly, the greater the transitional recovery, the greater the blocking participants showed ($r = .45$, $p < .05$). In order to more thoroughly investigate the novelty explanation described above, we also analysed the correlation between blocking and the RT effects for Z–/ZY– trials. The result indicated that neither transitional costs nor recovery exhibited during these trials were correlated with the magnitude of blocking ($r < .1$, $p > .8$), again suggesting that novelty is unlikely to be driving the effects of interest.

Discussion

Experiment 1 demonstrated that the blocking effect was associated with a specific pattern of response times over the course of the learning sequence. Specifically, these results suggest that a cognitively demanding process may be invoked during the early portions of Phase 2. These effects were only observed in those participants that ultimately exhibited a blocking effect; participants that did not exhibit blocking did not show these RT effects at all. In addition, blockers and nonblockers differed in their responses to blocking-related trials, but not on control trials, suggesting that the RT effects could not be attributed to novelty.

These results begin to address the issues raised above. For example, these results are inconsistent with the predictions of RW, which suggests that Phase 2 should elicit little in the way of learning. In contrast, the results do inform both attentional and inferential accounts. For example, because of the associative nature of attentional accounts (Kruschke, 2001; Mackintosh, 1975a), it has frequently been assumed that error-driven shifts in attention are automatic (Wills et al., 2007). Our results are not consistent with this view, but may suggest that attentional shifts in blocking are somewhat demanding. With respect to inferential accounts, the current results may begin to clarify the time course of the deliberative processing previously uncovered by De Houwer and Beckers (2003). The pattern of RT effects found here suggests that whatever cognitively demanding processing is invoked in Phase 2 operates relatively early and only briefly.

However, the conclusions suggested by Experiment 1 must be tempered by several methodological issues. For example, not all of our participants exhibited blocking effects, and we cannot be completely certain what the critical difference between the blockers and the nonblockers identified in Experiment 1 was. In addition, since the differences in blocking effects were not the result of any experimental manipulation, we cannot conclude that the changes in RTs were causally related to the blocking effect. Thus, Experiment 1 provides only relatively weak evidence that the slower response times in early Phase 2 reflected a mechanism responsible for the subsequent blocking effect. Experiment 2 seeks to further explore the relationship between the blocking effect and the cognitive processes operating early in Phase 2 by employing a cognitive load manipulation. In addition, we wished to test the relative importance of different portions of Phase 2.

EXPERIMENT 2

In Experiment 2, we imposed cognitive load during either the first half or the second half of Phase 2 in order to examine the effects on blocking. If, as Experiment 1 suggested, the increased cognitive processing early in Phase 2 reflects the operation of blocking-related processes, increasing cognitive load during the first half of Phase 2 should interfere with blocking effects more than increasing load late in Phase 2.

Method

Participants

Forty-nine undergraduate students at Stony Brook University participated in the experiment for partial course credit.

Materials

Participants were told that their task was to learn about the relationship between several new medications and allergic reactions. Pictures of 28 abstract shapes were used to represent the medications. They were randomly assigned to each role in the learning design for each participant. The allergic reaction was represented by a vertical bar (similar to Beckers et al., 2005) that was "filled" with green (representing no allergic reaction), half-filled in red (representing a moderate allergic reaction), or almost completely filled with red (representing a strong allergic reaction).

Design

The design was identical to that of Experiment 1 with the following modifications. An additivity pretraining phase (Beckers et al., 2005) was added before Phase 1. During the pretraining, two cues that were each followed by a moderate allergic reaction on their own were followed by a strong allergic reaction when they were paired together. Since we were mainly interested in Phase 2 of blocking, the additivity pretraining was included to increase the likelihood that all participants would exhibit blocking. Each trial type in the pretraining phase was repeated four times, and each trial type in Phase 1 and Phase 2 was repeated eight times. In Phase 2, each trial type was repeated four times in each half of the phase.

To manipulate cognitive load during the learning task, participants were required to perform a counting task, in which they had to count by threes, while simultaneously performing the learning task. There were three conditions, which differed with respect to when during the blocking sequence participants had to perform the concurrent load task. In the early-load condition, participants had to simultaneously perform the counting task during the first half of Phase 2 (the first 12 trials). In the late-load condition, participants performed the counting task during the second half of Phase 2 (the last 12 trials). In the control condition, participants performed the learning task without having to count at all. Participants went through the three conditions in one of three orders (i.e., control–early–late, early–late–control, and late–control–early), which counterbalanced the order in which each condition was administered.

Procedure

The procedure was identical to that of Experiment 1 with the following exceptions. First, no tones were played in Experiment 2. Second, in the early-load and late-load conditions, instructions for the counting task were intermixed with the trials in the learning sequence. Before the first loaded trial (i.e., 1st and 13th trials of Phase 2 in early-load and late-load conditions, respectively), an instruction screen was presented telling participants to begin counting by threes from a given, variable number. Participants were required to continue counting until a second instruction screen appeared (i.e., after the 12th and 24th trials of Phase 2 in early-load and late-load conditions, respectively). At this point, participants stopped counting and were asked to judge whether the last number they counted was greater than a probe number presented on the screen. There were no counting instructions in the control condition.

At the end of the learning sequence, participants judged how likely each medication was to cause allergic reaction on a scale that was only labelled "an allergic reaction is likely" on one end and "an allergic reaction is not likely" on the other end. Participants moved a cursor along the scale, and the judgements were later scaled so as to range from −1 to 1, respectively.

Results

Figure 4 shows the judgements for Cue X and the average judgements for the positive control cues (C and D). To analyse the influence of cognitive load on learning about the two types of cues, a 2 (cue: Cue X vs. control cues) × 3 (load: early-load vs. late-load vs. control condition) × 3 (condition order) three-way mixed ANOVA was conducted, with repeated measures on the first two factors. Lower bound correction was applied when necessitated by violations of sphericity in this and all subsequent analyses. The main effect of condition order and load was not significant (ps > .1). Main

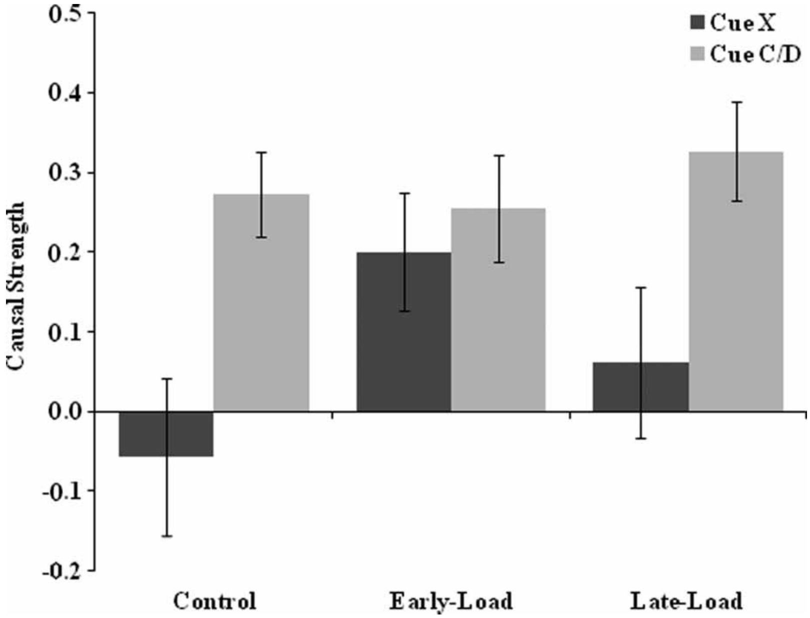

Figure 4. *Judgements (on a scale from −1 to 1) for Cue X and control cues (C and D) in early-load, late-load and control conditions in Experiment 2. Judgements for Cue X differed across conditions whereas judgements for the control cues did not. Significant blocking effects were observed in both the late-load and control conditions (and did not differ from each other). In contrast, no blocking effect was observed in the early-load condition.*

effect of cue was significant, $F(1, 92) = 17.30$, $p < .001$. The two-way interaction of load and condition order was not significant, $F(4, 92) = 1.04$, $p > .3$, but the interaction between cue and condition order was significant, $F(2, 46) = 3.54$, $p < .05$. However, neither the judgements for Cue X nor the judgements for the controls varied significantly across the three runs, ps $> .14$. Further analyses suggest that blocking effect (i.e., subtracting Cue X judgement from control cue judgements) was greater in the second run than in the first run, $t(48) = 2.07$, $p < .05$, but the magnitude of the blocking effect did not differ between the first and third runs or between the second and third runs (ps $> .1$). Most critical to the current investigation, we also observed a significant interaction between cue and load, $F(1.0, 92) = 4.17$, $p < .05$. Post hoc analyses indicate that judgements for Cue X differed across conditions, $F(2, 92) = 3.60$, $p < .05$, but judgements for control cues did not ($p > .3$). The three-way interaction was not significant $F(2.0, 92) = 2.23$, $p > .1$.

To investigate the influence of cognitive load on blocking, we compared participants' judgements for Cue X and their average judgements for the positive control cues (subtracting the former judgements from the latter). The traditional blocking effect would thus be observed if participants judged Cue X to be less likely to cause allergic reaction than the positive control cues. Analyses suggested that participants showed a significant blocking effect both in the control condition ($M = 0.33$, $SD = 0.59$), $t(48) = 3.92$, $p < .001$, and in the late-load condition ($M = 0.27$, $SD = 0.65$), $t(48) = 2.84$, $p < .01$. However, the blocking effect in the early-load condition was not significant ($M = 0.05$, $SD = 0.39$), $t(84) = 0.98$, $p = .33$. A one-way repeated measure ANOVA showed that blocking effects varied across conditions, $F(2, 96) = 3.88$, $p < .05$. Post hoc comparisons indicated that the blocking effect in the early-load condition was significantly smaller than that in the late-load condition, $t(48) = 2.08$, $p < .05$, and in the control condition, $t(48) = 3.35$, $p < .005$. The blocking effects in the control and late-load conditions were not different from each other, $t(48) < 1$. These results suggest that our cognitive load manipulation selectively influenced learning of Cue X but not the learning for control cues. Participants' exhibited standard blocking effects in the late-load and control conditions; despite repeated pairings, Cue X did not come to be associated with the outcome. In contrast, such blocking effects were not observed in the early-load condition.

Discussion

Experiment 2 showed that the blocking effect was modulated by cognitive load. However, blocking was affected only when cognitive load was increased during the early portions of Phase 2. In contrast, blocking remained entirely intact when cognitive load was increased only in the later portions of Phase 2. In Experiment 1, we observed different response patterns in blockers and nonblockers during the transition between Phase 1 and Phase 2. Extending these findings, Experiment 2 demonstrates that the cognitively demanding process engaged early in Phase 2 is causally related to blocking. Interfering with the cognitively demanding process operating during this critical period disrupts blocking. Load introduced only later in Phase 2 had no effect whatsoever. Moreover, load did not affect learning of all cues equally. Instead, cognitive load only altered how participants learned the relationship between Cue X and the outcome; learning about control cues was unaffected. Taken together, these results suggest that the RT effects observed in Experiment 1 are likely to reflect the transient operation of a process that is critically related to the exhibition of blocking. Furthermore, we note that the disruption in the early-load condition could not be compensated for later in the learning sequence.

Traditional models of associative learning (e.g., Rescorla & Wagner, 1972) are not able to account for these results. For example, RW suggests that all trials in Phase 2 should be processed in an identical manner, and therefore different portions of the phase should not have been differentially affected by the load manipulation. As for the attentional account, as discussed earlier, because these models do not specify the nature of the attentional shifts, it is unclear whether these models predict any relationship between cognitive load and blocking. The current findings also help to extend the results of De Houwer and Beckers (2003) who first reported that cognitive load reduced blocking. These authors argued that this finding supported the inferential account, which suggests that blocking results from a cognitively demanding logic. Our results go beyond this initial demonstration to suggest a unique need for cognitive resources early in Phase 2.

SIMULATIONS

We believe that the pattern of results observed in the current study is one that is not predicted by current theories of learning. However, we do believe that it may be possible to amend existing theories in order to accommodate the current findings. For example, it seems plausible that the inferential account could be modified so as to conform to the dynamics suggested by the current studies. Such possibilities remain largely speculative because the inferential account, a relatively new proposal, has yet to be described in great enough detail to examine detailed predictions.

Attentional accounts, on the other hand, have been formalized mathematically and thus lend themselves to more rigorous evaluation. As described above, associative processes have traditionally been characterized as automatic and independent of cognitive resources (Evans, 2003; Evans & Over, 1996; Kahneman, 2003; Sloman, 1996; Stanovich & West, 2000). Thus, the error-related attentional shifts suggested to produce blocking have also been assumed to be automatic (Wills et al., 2007). For this reason, it is not clear how to reconcile the associative processes described by attentional theories with the influence of a demanding concurrent task observed in the current study. To examine the potential mechanisms by which cognitive load could influence blocking, we simulated the behaviour of the attentional learning model. Specifically, we simulated the condition of Experiment 2 in order to evaluate three different mechanisms by which cognitive load might influence blocking effects.

General method

We utilized a version of Mackintosh's attentional learning model (Kruschke, 2001; Mackintosh, 1975a) because this algorithm includes an attentional mechanism that might be disrupted under cognitive load. This model, like that of Rescorla and Wagner (1972), consists of a simple network in which input nodes represent cues (e.g., conditioned stimuli), and output nodes represent outcomes (e.g., unconditioned stimuli). Each cue–outcome relationship has an associated weight, V, which represents the associative strength of the intervening relationship. In the model, learning consists of adjusting these weights according to Equation 1.

$$\Delta V_i = \alpha_i \beta \left(\lambda - \frac{\sum_{j=\text{present cues}} \alpha_j V_j}{\sum_{j=\text{present cues}} \alpha_j} \right) \qquad 1$$

In this equation, λ is an indicator of whether the outcome is present ($\lambda = 1$ in our simulations) or absent ($\lambda = 0$). The β parameter is an outcome-specific learning rate parameter ($\beta = .5$ in our simulations). The α parameter is a cue-specific learning rate parameter but can also be thought of as the attentional weight of each cue (α was initialized to .3 in our simulations). The parenthetical quantity in Equation 1 is referred to as prediction error and is simply the difference between the observed outcome (λ) and the predicted value of the outcome. This prediction is made by computing a weighted average of the strengths of the cues present on that trial. The strength of each present cue is weighted by its attentional weight (α). Thus, the more a learner attends to a cue, the more that cue's strength influences the prediction. The resulting quantity, ΔV, is then added to the strength of each cue present on that occasion. Cues that are absent never have their strengths adjusted.

The model also adjusts the attentional weights of each cue present on that trial. The first step in adjusting the attentional weights is to determine the amount of error attributable to each present cue. To compute the error attributable to a given cue, A, the model computes the difference between the associative strength of that cue and the outcome (i.e., $|\lambda - V_A|$). This quantity is then compared to the error attributable to the remaining present cues (i.e., $|\lambda - V_X|$, where X is the set of present cues excluding Cue A; when Cue A is the only cue present, its attentional weight is not adjusted). The attentional weight assigned to Cue A, α_A, is then adjusted according to Equation 2.

$$\begin{aligned} \Delta\alpha_A &= \gamma \cdot |\lambda - V_A| \; \mathit{if} \; |\lambda - V_A|1 < \left|\lambda - \sum V_X\right| \\ \Delta\alpha_A &= -\gamma \cdot |\lambda - V_A| \; \mathit{if} \; |\lambda - V_A| > \left|\lambda - \sum V_X\right| \end{aligned} \qquad 2$$

The parameter γ determines the rate at which the value of α is adjusted in response to cue-specific prediction error. Psychologically speaking, γ controls the speed of the learning-related attentional shifts.

To simulate Experiment 2, we presented the attentional model with a standard blocking paradigm that included both the blocking-related trials (consisting of Cues A and X) and the positive control trials (consisting of Cues C and D). To simulate Phase 1, we first presented the models with eight A+ trials (as in the current study). We then presented the models with both AX+ and CD+ trials, eight times each (also as in the current study).

As discussed earlier, what is less clear is how to simulate the cognitive load manipulations in Experiment 2. Because of this uncertainty, we evaluated three different variants of the above attentional model. First, we tested a variant in which load was assumed to decrease the speed with which learners could shift attention away from redundant cues (and toward predictive cues). Specifically, two values of γ were employed: one for the loaded portions of the learning sequence (γ_{Lord} = 0.1) and one for the unloaded portions of the learning sequence (γ_{NoLord} = 0.9). We refer to

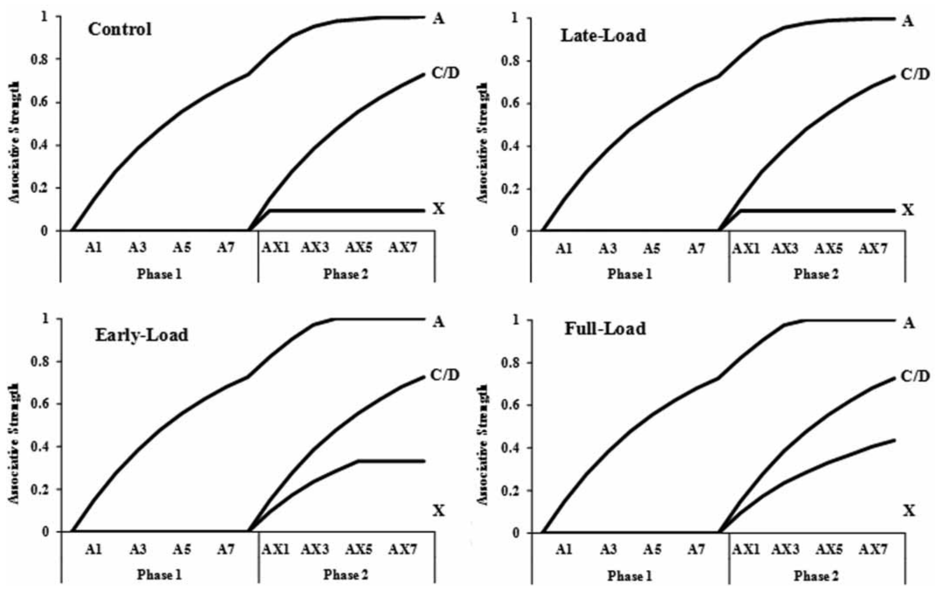

Figure 5. *Results from model simulations of Experiment 2. These acquisition curves illustrate the behaviour of the top-down attention model in four experimental conditions (see text for details). Under typical conditions, the model quickly shifts attention away from Cue X, preventing any learning about this cue. When the model is unable to shift attention early in the second half of Phase 2, the inability to shift attention had no consequence because the blocking-related shifts in attention have already completed by this time.*

this variant as the top-down attention model. Second, we tested a variant in which load was assumed to decrease the rate at which learners could acquire the cue–outcome associations themselves. Specifically, two values of β were employed: one for the loaded portions of the learning sequence (β_{Load} = 0.3) and one for the unloaded portions of the learning sequence (β_{NoLoad} = 0.5). We refer to this variant as the learning rate model. Third, we assumed that load diminished both the speed of attentional shift and the rate at which learners could acquire the cue–outcome associations themselves. That is, two parameter changes were made during the loaded portions of the sequence (γ_{Load} = 0.1, β_{Load} = 0.3), and two were made during the unloaded portions of the learning sequence (γ_{NoLoad} = 0.9, β_{NoLoad} = 0.5). We refer to this variant as the combination model.

We simulated each of the three conditions employed in Experiment 2. For the early-load and late-load conditions, the load-related parameter changes described above were implemented for the appropriate portions of the sequence. In the early-load simulation, the parameter changes were implemented during the first eight trials of the simulation (the first four AX+ trials and the first four CD+ trials). In the late-load simulation, the parameter changes were implemented during the last eight trials of the simulation (the last four AX+ trials and the last four CD+ trials). To simulate the control condition, the entire simulation was run with the standard (no load) parameter values. For the sake of completeness, we also simulated a full-load condition in which load was presented during the entire simulation.

Results

The results of these simulations are illustrated in Figure 5. We first discuss the results from the top-down attention model. As can be seen in Figure 5, the top-down attention model exhibits greater blocking effects in the control condition than in the early-load condition. In contrast, this model exhibits similar levels of blocking in the control condition than in the late-load condition.

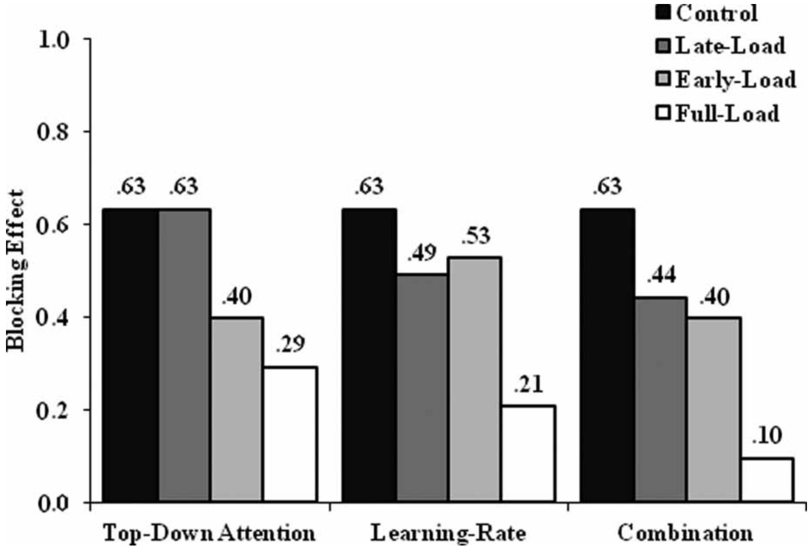

Figure 6. *Simulated blocking effects of the three models in all four conditions. The top-down attention model is successful in reproducing the pattern observed in Experiment 2. The learning-rate and combination models largely failed to mirror the observed pattern across conditions.*

This pattern mirrors the behavioural findings reported in Experiment 2. In addition, the top-down attention model exhibits larger blocking effects in the control condition than in the full-load condition. This finding mirrors the results reported by De Houwer and Beckers (2003). Perhaps surprisingly, the top-down attention model exhibited relatively similar levels of blocking in the early-load and full-load conditions.

The behaviour of the top-down attention model stems from the fact that the model quickly shifts attention away from Cue X and toward Cue A when first encountering the AX+ trials in Phase 2. When unloaded, these shifts occur quickly, preventing the model from associating Cue X with the outcome. Because these shifts occur early during Phase 2, cognitive load imposed during this period results in the model "erroneously" associating Cue X and the outcome. Once the load is removed, the model does ultimately shift attention away from Cue X. However, by this point the model does not unlearn the association between Cue X and the outcome it acquired during the early stages of Phase 2. In the late-load condition, the model shifts attention away from Cue X early in Phase 2. Once the load is imposed, the learned attentional biases are already in place, so no further attentional switching is needed, and thus no obvious consequences of the load are observed.

We next turn to the learning rate model. As can be seen in Figure 6, all experimental conditions elicited similar levels of blocking. The learning rate model did exhibit smaller blocking effects in the full-load condition than in the control condition, again consistent with the findings of De Houwer and Beckers (2003). However, if anything, the learning rate model exhibited slightly larger blocking effects in the early-load condition than in the late-load condition. Given that the exact opposite pattern was observed in Experiment 2, this particular variant does not seem particularly promising in light of the current data. This failure appears to stem from the fact that blocking is defined as the difference in two associative strengths (i.e., V_X and V_C). Changes to the global learning rate reduce the absolute magnitude of the resulting associative strengths, but have less of an influence on the relative magnitudes.

Finally, we turn to the combination model. As can be seen in Figure 6, the combination model exhibits greater blocking effects in the control condition than in the early-load condition. In addition, the combination model exhibits larger blocking effects in the control condition than in the full-load condition as reported by De Houwer and Beckers (2003). However, the combination model exhibits somewhat lower levels of blocking in the late-load condition than in the control. Thus, this model predicts that the early-load and late-load manipulations should have similar consequences. This prediction is intuitive, but contradicts the pattern observed in Experiment 2. Also, unlike the top-down attention model, the combination model exhibited smaller blocking effects in the full-load than in the early-load condition.

Discussion

The results of the current simulations are intended as a proof of concept to illustrate that traditional attentional models can be modified so as to render their behaviour consistent with the current study. Specifically, when we assumed that cognitive load influenced the ability to shift attention, model behaviour mirrored the behaviour of participants in Experiment 2. When we instead assumed that load simply slows acquisition, or slowed acquisition and influenced attentional shifts, the models produced behavioural patterns that were less consistent with the behaviour observed in the current study. These results support the idea that the nature of error-driven attentional shifts may not be entirely automatic; they are at least partially dependent on the availability of cognitive resources. The simulation results also suggest that early portions of Phase 2 in blocking are critical in that it is difficult to undo what is learned during this period. Indeed, interference with attentional shifting during early portions of Phase 2 appears to disrupt blocking nearly as much as interference administered during all of Phase 2.

We also note that the stimulations reaffirm the value of examining the time course of blocking-related processes as we have done in the current study. All three variants of the attentional model were capable of accounting for the reduced blocking effects previously observed in the full-load condition (e.g., De Houwer & Beckers, 2003). It was only when examining the effects of load at specific points during the learning sequence that we were able to distinguish between the three attentional models (this point is addressed further in the General Discussion).

Lastly, we briefly note that the results of our simulations and, to some extent, Experiment 2 are not entirely consistent with the findings reported by De Houwer and Beckers (2003). In their study, they found that blocking was reduced when a concurrent secondary task was imposed both during the learning sequence and as postlearning judgements were made. There are several possible explanations for this difference in results. An obvious difference is that De Houwer and Beckers (2003) employed a secondary task that required participants to discriminate between two tones played at random intervals whereas participants in our Experiment 2 counted backwards by threes. It is likely that counting backwards is more cognitively taxing than tone discrimination. Moreover, even though De Hower and Beckers did not find significant effect when load was not imposed as participants made judgements (their Experiment 1), there was still numeric differences between their easy and difficult load conditions, which is consistent with our findings. On the other hand, although imposing cognitive load during the judgement phase is potentially interesting, such a procedure also disrupts nonlearning processes (e.g., the processes by which acquired representations are transformed into causal judgements), which were not explored in the current study.

GENERAL DISCUSSION

Two experiments investigated the processes that unfold over the sequence of learning in a standard blocking paradigm. Results from both experiments indicate that the demand for cognitive processing increases during the early portions of Phase 2, when the novel Cue X is first introduced. In Experiment 1, the blocking effect was strongly associated with a slowing of secondary-task RTs during this period, suggesting an increased engagement in processing. The evidence provided by

Experiment 1 is relatively weak, however, because our identification of blockers and nonblockers was not based on any experimental manipulation. Thus, we cannot be certain what produced the different behaviour in the two groups. However, Experiment 2 provides converging evidence using a cognitive load manipulation. Load eliminated the blocking effect when administered during the early portion of Phase 2, but not when administered during the late portion of Phase 2.

The results of the current experiments contradict the predictions of traditional models of associative learning. For instance, according to RW (Rescorla & Wagner, 1972), learners in a blocking procedure should not be surprised when confronted with the AX+ trials at any point during Phase 2 because they have already learned to expect the outcome based on prior learning about Cue A. Therefore, there is no reason to predict any heterogeneity during Phase 2. Such an account cannot explain our observation that processing early in Phase 2 was different from processing during later portions. Moreover, associative processes are presumed to be automatic, fast, and not dependent on cognitive resources (Evans, 2003; Evans & Over, 1996; Kahneman, 2003; Sloman, 1996; Stanovich & West, 2000). Therefore, traditional associative models, including RW, cannot accommodate the effect of cognitive load manipulations (De Houwer & Beckers, 2003). Thus, the RT effects observed in Experiment 1 and the effect of load observed in Experiment 2 are simply inconsistent with these models.

Attentional theories of learning (Kruschke, 2001; Mackintosh, 1975a), on the other hand, suggest that learners in a blocking paradigm need to first learn that Cue X is a redundant predictor of the outcome and then shift their attention away. However, attentional models do not specify whether attentional shifts are demanding or consume cognitive resources. Despite recent empirical findings supporting attentional accounts (e.g., Kruschke et al., 2005; Le Pelley, Beesley, & Griffiths, 2011; Wills et al., 2007), the nature of these attentional processes has remained uncertain. For this reason, it is not entirely clear what pattern of results attentional accounts should predict in the current studies. Given that attentional theories of learning are associative in nature, it is natural to assume that their operation, including any attentional mechanisms, is automatic and thus free of cognitive demands. However, we presented simulations that suggest that attentional accounts can be modified to shift attention via cognitively demanding, top-down processes (e.g., executive control) in order to more easily explain the findings reported here. Indeed, these simulations suggest that it is crucial to assume that attentional shifts in learning are at least partially top-down and depend on cognitive resources; simulating other plausible consequences of cognitive load proved unsuccessful in replicating the current results.

The current results may also help constrain the recently proposed inferential account (Beckers et al., 2005; De Houwer & Beckers, 2003; Lovibond, 2003), which suggests that learners apply logical rules to infer the status of Cue X. Advocates of inferential accounts do not explicitly specify when such analytic processing occurs, but our results suggest that the early portions of Phase 2 are critical to blocking. In Experiment 2, the early-load condition completely eliminated the blocking effect, which suggests that load allowed learning about the typically blocked Cue X and that learners were unable to "undo" this learning later in Phase 2. In contrast, in the late-load condition, blocking-related processes may have been able to finish before load was imposed, leaving blocking intact.

Process versus rule-based models

As demonstrated in the current experiments, investigating the processes that unfold over the course of sequence of trials in a blocking paradigm can be useful for discriminating among plausible accounts of blocking. Given the prominence of blocking as an empirical phenomenon, theorists must all construct their models so as to accurately mimic the gross features of the observed behaviour. As a consequence, the only difference between competing models is frequently in the mechanisms proposed to achieve the agreed-upon behaviour. By measuring (in Experiment 1) and manipulating (in

Experiment 2) resources that may be implicated in the processes underlying blocking, we were able to present more specific criteria by which theories can be evaluated.

The problem with alternative models making similar behavioural predictions is not unique to the phenomenon of blocking. Indeed, this problem is pervasive in the field of learning. As mentioned earlier, research in blocking has largely focused on factors that affect judgements elicited at the end of the learning sequence. Similarly, researchers in the field of human contingency learning frequently seek to compare learners' final contingency judgements to normative standards. Models developed with such a focus tend to specify the rules for how the final judgements should be computed based on the summary of observations in a contingency matrix (e.g., Cheng, 1997; Cheng & Novick, 1992; Jenkins & Ward, 1965). These models avoid specifying how pieces of information are (or should be) processed so as to achieve these judgements.

However, there is mounting evidence that information is not processed in the same way on every trial in a contingency learning task. For example, Liu and Luhmann (2012) had participants perform a simple tone-discrimination task as they simultaneously completed a contingency learning task. Participants' responses to the tones were slower on trials that presented information that was inconsistent with the majority of the previous trials. These results suggest that trials are not processed uniformly. Instead, information in a given trial appears to be analysed with respect to previous information, with inconsistent information demanding deeper processing.

Research on so-called order effects has similarly shown that covariation information is not processed equally (e.g., Dennis & Ahn, 2001; Lopez, Shanks, Almaraz, & Fernandez, 1998; Luhmann & Ahn, 2011). This work manipulates presentation order while holding the set of covariation information constant. This work frequently finds that presentation order influences final judgements, suggesting that the position within a learning sequence at which information is encountered affects how it is processed. Participants' judgements often reflect the information presented at either the beginning (Dennis & Ahn, 2001) or the end (Lopez et al., 1998) of the trial sequence. These findings have been taken as evidence that covariation information is not processed uniformly and that post-learning contingency judgements are best thought of as reflecting dynamic processes that unfold over the course of learning. Results in the current experiments similarly suggest that covariation information is processed differently depending on when during the course of learning it is encountered. Discovering how the relevant processes unfold over time should help to further discriminate among competing theories and bring fresh insights into the study of learning.

Original manuscript received 3 August 2011
Accepted revision received 24 July 2012
First published online 13 September 2012

REFERENCES

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

Beckers, T., Miller, R. R., De Houwer, J., & Urushihara, K. (2006). Reasoning rats: Forward blocking in Pavlovian animal conditioning is sensitive to constraints of causal inference. *Journal of Experimental Psychology: General*, *135*, 92–102.

Beesley, T., & Le Pelley, M. E. (2011). The influence of blocking on overt attention and associability in human learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *37*, 114–120.

Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgement. *Memory and Cognition*, *18*, 537–545.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.

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

De Houwer, J., & Beckers, T. (2002a). Higher-order retrospective revaluation in human causal learning. *The*

Quarterly Journal of Experimental Psychology, *55*, 137–151.

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

De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *The Quarterly Journal of Experimental Psychology*, *56*, 345–357.

Dennis, M., & Ahn, W. (2001). Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory & Cognition*, *29*, 152–164.

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

Evans, J. St. B. T. (2003). In two minds: Dual-process accounts of reasoning. *Trends in Cognitive Sciences*, *7*, 454–459.

Evans, J. St. B. T., & Over, D. E. (1996). *Rationality and reasoning*. Hove, UK: Psychology Press.

Jenkins, H. M., & Ward, W. C. (1965). Judgement of contingency between responses and outcomes. *Psychological Monographs*, *79*, 1–17.

Kahneman, D. (2003). A perspective on judgment and choice. *American Psychologist*, *58*, 697–720.

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

Kruschke, J. K. (2001). Toward 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 and Review*, *7*, 636–645.

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

Le Pelley, M. E., Beesley, T., & Griffiths, O. (2011). Overt attention and predictiveness in human contingency learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *37*, 220–229.

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

Liu, P., & Luhmann, C. C. (2012). *Evidence for on-line processing during causal learning*, Unpublished manuscript.

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

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

Luhmann, C. C., & Ahn, W. (2011). Order effects during learning: Expectations and interpretations. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *37*, 568–587.

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

Mackintosh, N. J. (1975b). Blocking of conditioned suppression: Role of the first compound trial. *Journal of Experimental Psychology: Animal Behavior Processes*, *1*, 335–345.

Mitchell, C. J., Lovibond, P. F., Minard, E., & Lavis, Y. (2006). Forward blocking in human learning sometimes reflects the failure to encode a cue–outcome relationship. *Quarterly Journal of Experimental Psychology*, *59*, 830–844.

Pineno, O., & Miller, R. R. (2007). Comparing associative, statistical, and inferential reasoning accounts of human contingency learning. *The Quarterly Journal of Experimental Psychology*, *60*, 310–329.

Posner, M., & Boies, S. (1971). Components of attention. *Psychological Review*, *78*, 391–408.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical Conditioning II: Current research and theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, *37*, 1–21.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3–22.

Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, *28*, 303–333.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, *23*, 645–726.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*, 222–236.

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

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.