King's Research Portal

DOI:

[10.1080/17470218.2016.1188407](https://doi.org/10.1080/17470218.2016.1188407)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Luque, D., Vadillo Nistal, M. A., Le Pelley, M. E., & Beesley, T. (2016). Prediction and uncertainty in associative learning: Examining controlled and automatic components of learned attentional biases. *Quarterly journal of experimental psychology*. <https://doi.org/10.1080/17470218.2016.1188407>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Prediction and uncertainty in associative learning: Examining controlled and automatic components of learned attentional biases

David Luque, Miguel A. Vadillo, Mike E. Le Pelley & Tom Beesley

To cite this article: David Luque, Miguel A. Vadillo, Mike E. Le Pelley & Tom Beesley (2016): Prediction and uncertainty in associative learning: Examining controlled and automatic components of learned attentional biases, *The Quarterly Journal of Experimental Psychology*, DOI: [10.1080/17470218.2016.1188407](https://doi.org/10.1080/17470218.2016.1188407)

To link to this article: <http://dx.doi.org/10.1080/17470218.2016.1188407>

Accepted author version posted online: 13 May 2016.

Submit your article to this journal [↗](#)

View related articles [↗](#)

View Crossmark data [↗](#)

Publisher: Taylor & Francis & The Experimental Psychology Society

Journal: *The Quarterly Journal of Experimental Psychology*

DOI: 10.1080/17470218.2016.1188407

Running head: Prediction and uncertainty

Prediction and uncertainty in associative learning: Examining controlled and automatic components of learned attentional biases

David Luque¹, Miguel A. Vadillo², Mike E. Le Pelley¹, & Tom Beesley¹

¹School of Psychology, UNSW Australia, Sydney, Australia

² Department of Primary Care and Public Health Science, King's College London, London, UK

Author note: Correspondence should be addressed to David Luque, School of Psychology, UNSW Australia, Sydney NSW 2052, Australia. Phone: +61 2 9385 3041. E-mail:

d.luque@unsw.edu.au. The authors would like to thank Janelle Fernandes for her assistance in running the experiments. This work was supported by Australian Research Council (ARC) Discovery Projects (grants DP140103268 and DP160103063). Mike Le Pelley was supported by an ARC Future Fellowship (FT100100260).

Abstract

It has been suggested that attention is guided by two factors that operate during associative learning: a predictiveness principle, by which attention is allocated to the best predictors of outcomes, and an uncertainty principle, by which attention is allocated to learn about the less-known features of the environment. Recent studies have shown that predictiveness-driven attention can operate rapidly and in an automatic way to exploit known relationships. The corresponding characteristics of uncertainty-driven attention, on the other hand, remain unexplored. In two experiments we examined whether both predictiveness and uncertainty modulate attentional processing in an adaptation of the dot probe task. This task provides a measure of automatic orientation to cues during associative learning. The stimulus-onset-asynchrony of the probe display was manipulated in order to explore temporal characteristics of predictiveness- and uncertainty-driven attentional effects. Results showed that the predictive status of cues determined selective attention, with faster attentional capture to predictive than to nonpredictive cues. In contrast, the level of uncertainty slowed down responses to the probe regardless of the predictive status of the cues. Both predictiveness- and uncertainty-driven attentional effects were very rapid (at 250 ms from cue onset) and automatically activated.

Keywords: attention, associative learning, dot probe, predictiveness, uncertainty.

Visual attention determines which stimuli are preferentially processed. It allows for the focusing of limited cognitive resources on important aspects of the environment, to the detriment of processing less important information. Thus, in order to understand our cognitive system, it is crucial to investigate which factors modulate this selection process. These factors are usually divided into two mutually exclusive functional categories according to whether attentional modulation is caused by physical characteristics of the stimuli (i.e., stimulus-driven modulation) or is caused by cognitive factors such as goal-directed intentions or motivations (i.e., goal-directed modulation) (e.g., Corbetta & Shulman, 2002; Yantis, 2000). For example, when driving we might use goal-directed, top-down attention to prioritize processing of events on the road ahead, and to ignore conversation from the backseat. But a physically salient event (e.g., a sudden bang from behind the car) will capture our attention in an bottom-up, stimulus-driven fashion regardless of our goals (e.g., Theeuwes, 1992; Folk, Remington, & Johnston, 1992).

However, it has recently been argued that there is a third category of influences on attentional selection that is neither fully goal-directed nor stimulus-driven, and which comes into play when people have had previous experience with stimuli. Specifically, it has been suggested that our attention is influenced by what we have learned about how stimuli relate to other events in the environment (Anderson, 2013; Awh, Belopolsky, & Theeuwes, 2012; Chelazzi, Perlato, Santandrea, & Della Libera, 2013; Le Pelley, Pearson, Griffiths, & Beesley, 2015; Le Pelley, Beesley, & Griffiths, in press; Le Pelley, Mitchell, Beesley, George, & Wills, in press).

While the suggestion of a relationship between learning and attention has received a great deal of recent interest, it is not a new idea. William James (1890/1983) wrote about *derived attention*: a form of attention to a stimulus that “owes its interest to association with some other immediately interesting thing” (p. 393). More importantly for current purposes, models of associative learning proposed over the last 40 years provide formal accounts of how learning and attention might interact (see Mitchell & Le Pelley, 2010). These models were developed largely on the basis of the results of studies of animal conditioning, but have since been applied to explain behaviour in studies of human learning. Such ‘attentional models’ of associative learning propose that the attention that is paid to a stimulus is influenced by the certainty or uncertainty of the predictions that it makes about other events.

In the associative learning literature, uncertainty is understood as the variance in the nature or magnitude of the outcome which follows a cue or an operant behaviour (Rushworth & Behrens, 2008). For instance, imagine an urn containing red balls and yellow balls. If the urn contains 50% red balls and 50% yellow balls, then the uncertainty about the colour of a randomly drawn ball will be greater than if the urn contained 80% red and 20% yellow. At the limit, if the urn contained only red balls, the uncertainty would be zero.

As noted above, attentional models of associative learning describe how uncertainty might influence which cue stimuli receive attentional priority for future learning. For example, according to the influential model of associative learning proposed by Mackintosh (1975), it is those cues that are the most reliable predictors of significant events in the environment (i.e., those cues that most consistently and accurately predict the events that follow them) that will receive the greatest amount of attentional processing. We term

this a *predictiveness-driven principle* for guiding attention. Phrased in terms of the uncertainty associated with cues, this model states that those cues that have, in the past, had more uncertain consequences will tend to receive less attention compared to more reliable cues that are presented alongside them. Thus, the Mackintosh model captures the selective nature of attention, in that it sees the learning system as seeking out the most reliable sources of information. Such attentional biases favouring certain over uncertain sources of information have been shown widely in the human and animal learning literature (for reviews, see Le Pelley, 2004, 2010; Le Pelley, Mitchell et al., in press; Pearce & Mackintosh, 2010).

However, when interacting with the world, our aim is not only to obtain well-predicted and reliable outcomes or rewards, but also to gather further information about the environment and hence reduce uncertainty. That is, adapted intelligent animals (and machine learning systems) devote time and effort towards exploring their environment (e.g., Oudeyer, Kaplan, & Hafner, 2007). The attentional system may have an important role to play as part of this uncertainty-reduction process (see Gottlieb, Oudeyer, Lopes, & Baranes, 2013). Consistent with this idea, computational modelling suggests that complementing a predictiveness-driven attentional process with an exploratory attentional mechanism might be an optimal information-processing strategy (e.g., Dayan, Kakade, & Montague, 2000). This *uncertainty-driven* attentional process is exemplified by the Pearce and Hall (1980) model of associative learning. According to this model, cues for which prediction errors have recently occurred (i.e., cues whose consequences are highly variable and hence uncertain), will attract more attention than certain cues (i.e., cues for which prediction errors have recently been minimal). Uncertainty-driven attention was proposed originally as a mechanism by which resource optimization is achieved. This is because uncertainty-driven

attention assigns most cognitive resources to processing those stimuli whose predictive status is currently poorly understood, which may be adaptive, since it may allow the true status of those cues to be clarified (Pearce & Hall, 1980). In addition, it has been pointed out that, even considering an unlimited processor, exploratory attention would be necessary in order to avoid local minima in the solution to learning tasks (Dayan et al., 2000). Suppose a predator learns that a particular type of abundant yellow beetle makes a tasty meal. If this predator used an exploitative attentional strategy, it would subsequently hunt exclusively for yellow beetles, and ignore other potential sources of food (say, somewhat rarer red spiders) that may also be palatable (Bond & Kamil, 1998). But the availability of different prey can change over time. Suppose that, for some reason, the once-abundant yellow beetles become less common than red spiders. Our 'pure exploitation' predator would now be at a disadvantage, since it would remain hunting for a scarce resource when an easier alternative is available. In contrast, a predator that occasionally changes its foraging strategy and explores other potential prey items (about which it is currently uncertain) would rapidly learn to increase its reliance on red spiders, and may thus be more likely to thrive.

Animal experimentation has shown results consistent with the operation of uncertainty-driven attention (Haselgrove, Esber, Pearce, & Jones, 2010; Kaye & Pearce, 1984a; Swan & Pearce, 1988; Wilson, Boumphrey, & Pearce, 1992), and the neural basis of this mechanism has been delimited to critical involvement of the prefrontal and amygdala regions (Fiorillo, Tobler, & Schultz, 2003; Roesch, Esber, Li, Daw, & Schoenbaum, 2012).

There are also experiments (though not many) which support a role for uncertainty-driven attention in human causal learning tasks (Beesley, Nguyen, Pearson, & Le Pelley, 2015; Griffiths, Johnson, & Mitchell, 2011; Hogarth, Dickinson, Austin, Brown, & Duka, 2008; for a review, see Le Pelley, Mitchell et al, in press).

From a theoretical point of view, predictiveness-driven and uncertainty-driven attentional processes are not necessarily exclusive. Indeed, modern attentional models of associative learning have strived to incorporate both principles, either in so-called ‘hybrid’ or dual-process models (e.g., Le Pelley, 2004; Pearce & Mackintosh, 2010), or in a flexible single-process account (Esber & Haselgrove, 2011). Supporting this view, Beesley et al. (2015) showed that overt attention to stimuli (measured via eye-tracking) could be determined by both predictiveness and uncertainty within the same task. These experiments used a *learned predictiveness* design (Le Pelley & McLaren, 2003; Lochmann & Wills, 2003) in which, on each trial, participants were presented with a compound of two cues and had to predict which outcome would occur following that compound; corrective feedback was provided. Only one cue from each pair (the *predictive* cue) was informative of the outcome, while the other cue was nonpredictive, since it was paired with each of the two outcomes equally often. When eye-gaze dwell times to predictive and nonpredictive cues were analysed, results complied with the predictiveness-driven principle; people spent longer looking at predictive cues than nonpredictive cues (see also Le Pelley, Beesley, & Griffiths, 2011). Beesley et al. also manipulated the uncertainty of each compound. In the *certain* condition, each compound had a deterministic relationship with its paired outcome (that is, the same outcome always followed a particular compound), such that participants experienced minimal prediction errors once the contingencies were learned. In contrast, for the *uncertain* condition, compound–outcome relationships were probabilistic (that is, each compound was typically, but not always, followed by one of the outcomes) and so occasional prediction errors were inevitable, even after asymptotic learning. Dwell-time analysis revealed that participants spent a greater proportion of the trial time attending to the cues in uncertain compounds as compared to cues in certain compounds, regardless of

the predictive status (predictive/nonpredictive) of those cues; i.e., a main effect of uncertainty, but no interaction between uncertainty and predictiveness. Thus, Beesley et al.'s results suggest that predictiveness-driven attention determines the *selection* of the most predictive stimulus from the environment, while uncertainty-driven attention reflects a mechanism which prioritises exploration of cues when recent prediction errors have been experienced.

The current article concerns the nature of the attentional processes underlying predictiveness-driven and uncertainty-driven attention. Recent research suggests that learning can exert an effect on attentional capture by stimuli that is rapid and automatic, in the sense that it occurs regardless of task demands and participants' ongoing goals (see Anderson, 2013; Awh et al., 2012; Chelazzi et al., 2013; Le Pelley et al. 2015; Le Pelley, Mitchell et al., in press). These previous studies considered the influence of learning about the value of rewards paired with stimuli (so-called *value-driven attentional capture*). In the experiments reported here, we investigated whether this pattern of rapid and automatic capture also applies to changes in attention that are driven by learning about the *variability* of outcomes paired with stimuli, in terms of predictiveness and uncertainty.

Toward this end, we explored the influence of predictiveness and uncertainty on attention using a variant of the spatial cueing task (Posner, Nissen, & Ogden, 1978). Specifically, we used an adaptation of the *dot probe* task (MacLeod, Mathews, & Tata, 1986), which we have previously used to study the operation of predictiveness-driven attentional processes during associative learning (Le Pelley, Vadillo, & Luque, 2013). In the current experiments, participants were initially trained on an associative learning (AL) task: on each trial, two cues were presented on the screen and participants made a

categorization response (does this pair of cues belong to the 'up' category or the 'down' category?), with immediate corrective feedback provided. Later, a dot probe task was superimposed on this AL task. Now when the two cues were presented on each trial, participants first had to respond as rapidly as possible to a probe stimulus that appeared over one cue (with equal likelihood of the probe appearing on either of the two cues), before subsequently making a categorization response as part of the AL task, just as before. Critically, this procedure allowed us to manipulate the predictiveness and uncertainty of cues and compounds in the AL task, and observe the resulting influence on attention to cues through response times on the dot probe task. That is, if the contingencies in the AL task were such as to cause participants to selectively attend to one cue over the other, then responses to the probe stimulus should have been faster if it appeared in the location of the attended cue compared to its counterpart (cf. MacLeod et al, 1986; Posner et al., 1978). Varying the timing of the probe onset allowed us to examine the time-course of attentional orienting and disengagement of attention to cues as a function of their predictiveness and uncertainty.

Using such a task, Le Pelley et al. (2013) showed that responses to the probe were faster when it appeared in the location of a predictive cue than a nonpredictive cue, if the probe appeared rapidly after cue-onset (with a stimulus-onset asynchrony [SOA] of 250-350 ms). Crucially, since the probe was equally likely to occur in the position of the predictive and nonpredictive cue, there was no advantage to be gained in directing greater attention to one type of cue than the other, prior to the probe presentation. Indeed, participants were explicitly informed that in order to respond to the probe as quickly as possible, their best strategy was to ignore the initially presented cues. Since attentional bias towards predictive cues was not required by the dot probe task or indeed an adaptive strategy with

regard to that task, the implication is that the observed bias reflected the operation of a process independent of participants' goal for this task (i.e., to localize the probe as quickly as possible). Consistent with this idea, Le Pelley et al. (2013) demonstrated that providing more time for participants to consciously process the stimuli—by increasing the SOA on dot probe trials to 1000 ms—significantly *weakened* the influence of predictiveness on dot probe responding. This supports the idea that the bias towards predictive cues observed in the short SOA condition was not a result of goal-directed, controlled processing but instead an *automatic*, rapid, and short-lived attentional process within the region of 250 milliseconds after cue-onset (for convergent evidence, see Feldmann-Wüstefeld, Uengoer, & Schubö, 2015).¹

The current experiments examined whether manipulations of uncertainty also influence attentional processing in the dot probe task, and hence whether this influence reflects a rapid and automatic process. It might be expected that in contrast to the predictiveness factor, uncertainty promotes the engagement of a more controlled, goal-directed process: Pearce and Hall (1980) described an uncertainty-driven increase in attention in their model as reflecting a “*controlled processing strategy*” (Pearce & Hall, 1980, p. 549, italics in original). Yet to date, this suggestion of a processing distinction between attentional processes remains untested in the human (or animal) associative learning literature.

The current experiments explored attention to cues during associative learning with designs that follow Beesley et al. (2015) in simultaneously manipulating the level of uncertainty for compounds of cues and the level of predictiveness of cues within each compound. Beesley et al.'s experiments did not permit an examination of whether

uncertainty-driven effects were produced by automatic and/or controlled changes in covert attention, since the time-course of participants' gaze was not assessed. By using the dot-probe task, the current experiments examine the time-course of attention by varying the SOA between presentation of the cues and appearance of the probe as a within-subject variable. Thus, the current experiments go beyond the experiments of Beesley et al. in providing an examination of the automatic/controlled nature of uncertainty-driven attention.

If uncertainty-driven attention reflects a controlled, goal-directed process (for instance, as a consequence of a volitional effort to find new information so as to minimize errors in the AL task), we should not necessarily expect to observe an influence of uncertainty (defined with regard to the AL task) on rapid responses to the dot probe, since participants have ample time to freely explore the cues following their dot probe response. In contrast, if uncertainty-driven attention is rapid and automatic (as seems to be the case for value-driven and predictiveness-driven attention), the manipulation of uncertainty would be expected to affect the responses to the dot probe. However, the specific pattern to be anticipated under this latter hypothesis is unclear. For instance, it is possible that uncertain cues may initiate a general increase in arousal or vigilance and hence decreased response times to probes. On the other hand, it may also be possible that uncertainty produces a rapid and automatic exploratory examination of the cues, in search of new information that will assist in reducing prediction error. Such diffuse deployment of attention may hinder the localization of the probe and hence result in increased response times in the dot probe task under conditions of uncertainty.

Experiment 1

Method

Participants and apparatus. For both experiments, we guided our decision on sample sizes using estimates from Le Pelley et al.'s (2013) Experiment 2, which is the closest existing procedure to the current experiments. Since we expected to replicate the predictiveness result found by Le Pelley et al. in their short SOA condition, we took the effect size from this condition for the current power analysis. The effect size from this experiment was $\eta_p^2 = .07$; considering that effect size, the needed sample for a power of 0.8 is 21 participants. Thus, we aimed to achieve 20-25 valid cases per SOA condition after the application of our rejection criteria (see above).

Based on our experience with this type of task, we estimated that around 3/4 of all participants would pass our rejection criterion. In order to achieve 20-25 valid cases per condition, seventy UNSW Australia students participated for course credit in the Experiment 1. Participants were randomly allocated to either the 250ms or 1000ms SOA conditions. They were tested in individual enclosed cubicles, using standard PCs with 58.4 cm monitors (1920 × 1080 resolution, 120 Hz), at a viewing distance of approximately 60 cm. Stimulus presentation was controlled by the Cogent 2000 toolbox (Cogent 2000 team, Wellcome Trust, London, UK) running under MATLAB (Mathworks Inc.). Participants made all responses with their right hand, using the arrow keys of custom keyboards which provide average response latencies of around 1 millisecond (DirectIN keyboard, Empirisoft, New York).

Stimuli. Cues were eight coloured polygons, which differed in colour and the thickness of their 'spikes' (see Figure 1B). Colours (RGB) and relative luminance (on a 255-level scale) for

colours were as follows: red (R255, G0, B0; L54), yellow (R230, G230, B51; L217), green (R0, G204, B51; L150), turquoise (R51, G255, B255; L212), blue (R0, G128, B255; L110), magenta (R255, G51, B255; L109), brown (R153, G102, B0; L105), and salmon (R255, G128, B128; L155). The polygon shapes were framed by white squares with sides subtending 4.7° visual angle. For each participant these stimuli were randomly assigned to play the roles of the various cues shown in Table 1. The cue stimuli were presented on the horizontal midline of the screen, on either side of a small, central fixation cross. The distance from the centre of the cross to the centre of each square subtended 4.7° . The probe was a white square which subtended 0.67° . This appeared superimposed centrally on one of the stimuli. The screen background was black.

Design. The experiment contained three phases (see Table 1): the *Pretraining* phase established the cue–outcome contingencies in the absence of the dot probe task; *Phase 1* continued training these cue–outcome contingencies in the presence of the dot probe task; and *Phase 2* involved a manipulation of the uncertainty of the relationships between compounds of cues and the outcomes with which they were paired.

For the associative learning (AL) task in Pretraining and Phase 1, participants were required to make either an *up* or *down* categorization response on each trial; Table 1 shows the correct response to each of the eight cue-compounds that were presented. Each compound contained one predictive cue (labelled p1, p2, p3 and p4 in Table 1) and one nonpredictive cue (labelled n1, n2, n3 and n4). As can be seen in Table 1, predictive cues consistently indicated the correct categorization response during these phases; for instance, whenever p1 appeared, the correct response was always response R1 (*up* or *down*, counterbalanced across participants). Thus, once the contingencies were learned, the

participants could make these responses accurately using only the information provided by the predictive cues and therefore perfectly predict the correct response on each trial. In contrast, nonpredictive cues provided no information regarding the correct response (e.g., for half of the appearances of n1 the correct response was R1, for the other half it was R2, see Table 1), and therefore the outcome could not be anticipated using these nonpredictive cues alone.

The left/right positions in which cues appeared was counterbalanced within each block. Hence each block of the Pretraining phase contained two presentations of each of the eight compounds shown in Table 1, one with the predictive stimulus on the left and one with the predictive stimulus on the right. This resulted in 16 trials per block and the Pretraining phase involved four such blocks.

The dot probe task was superimposed on the AL task in Phase 1. We needed to ensure that the probe was equally likely to appear over either cue of each compound (e.g., that the probe was equally likely to appear over cue p1 as it was to appear over cue n1), such that there was no reason for participants to *strategically* orient attention to one or other cue prior to the appearance of the probe. Consequently, each block of Phase 1 comprised 32 trials: every combination of the 8 compounds, the left/right positioning of cues, and the left/right location of the probe. Trials within a block were presented in a random order. There were eight such blocks in Phase 1.

Four compounds were presented in Phase 2, again with the left/right position of the cues and left/right location of the probe counterbalanced. In addition to these factors we also manipulated uncertainty. During Phase 2, each of the four compounds shown in Table 1 was presented 24 times. For the uncertain compounds, on 16 of these 24

presentations (67%) the correct response was consistent with the response that had been correct for that compound during Phase 1, while on the remaining 8 trials (33%) the correct response was inconsistent with prior training. Within the consistent and inconsistent trial types, cue and probe location were counterbalanced, as before. For the certain compounds, the correct response was always consistent with the correct response in Phase 1. Trial order was randomized.

The SOA of the probe (250 or 1000 ms) was manipulated between-subjects, and the same SOA applied throughout Phase 1 and Phase 2.

Procedure. Initial instructions described the AL task: participants were told that on each trial a pair of stimuli would appear and that they were required to make a response using either the up or down arrow key. They were informed that their task was to learn the correct response for each stimulus pair. Participants then completed the Pretraining phase. Each trial began with the presentation of a central fixation cross, followed after 500 ms by two cue stimuli. One second after the presentation of the cues, the text 'UP or DOWN?' appeared centred on the screen in the space between the two cues (30-point Arial font for the words 'UP' and 'DOWN, 20-point for the word 'or'). This way, the cues were still visible even with the text on the screen. Participants made a categorization response using the up or down arrow keys and this response was allowed only when the words 'UP or DOWN?' were on the screen. If the response was correct, no explicit feedback was provided and the next trial began after an inter-trial interval of 1s. If the response was incorrect, then the message 'incorrect' appeared for 3s, followed by the inter-trial interval. Participants could take as long as they liked to make the categorisation response in all phases.

Instructions prior to Phase 1 stated that participants would now have to perform an additional task on each trial, which was to respond as rapidly as possible to the location of a small white square, using the left and right arrow keys. Figure 1A shows a schematic of a typical trial in Phases 1 and 2. After an initial fixation interval of 500 ms, the two cues appeared. Then, after an SOA of 250 ms (short SOA condition) or 1000 ms (long SOA condition), the probe appeared, superimposed on the centre of one of the cues. This probe remained on-screen until participants made the correct response (left arrow key for a probe presented on the left; right arrow key for a probe on the right). Immediately after making the correct response to the location of the probe, the probe disappeared and participants made their categorization response as before. The cues were visible for the duration of the trial (see Figure 1A). Participants were told that in order to respond to the white square (the probe) as quickly as possible, the best strategy was to ignore the coloured polygons until after they had made their left/right response to the square. Specifically, this instruction was as follows: “Try to respond to the square as fast as you can. To do so, it is best if you ignore the two figures until you have responded to the location of the square” (the instructions also showed the text underlined). Participants completed the experiment in a single 45-minute session.

Results and discussion

Data pre-processing. Averaged across all participants, accuracy on the AL task increased during training and reached a high level by the end of Phase 1 (see Figure 1D). Since our results in Phase 2 depended on appropriate acquisition of the associative contingencies in Phase 1, participants who performed with less than 60% accuracy during the final three

blocks of Phase 1 were removed from the analysis. As a consequence of this criterion, the data from sixteen participants were not analysed further (final sample, $N = 54$, with 27 in each SOA condition).

For the analysis of dot probe RTs we adopted a pre-processing pipeline very similar to that used by Le Pelley et al. (2013, Experiments 2 and 3). Since the dot-probe task provided no explicit feedback on response accuracy, trials in which an incorrect response was made at any point during the presentation of the probe were not analysed (1% of all trials); nor were trials in which responses were very fast (under 150 ms) or very slow (over 1000 ms) (5% of all trials). Finally, as a guard against within-participant outliers, trials with RTs to the probe lying more than 2.5 standard deviations from each participant's mean were also excluded from analysis (1% of the remaining trials).

Statistical analyses. All tests were performed at the $\alpha = 0.05$ significance level. For repeated-measures ANOVA, Greenhouse-Geisser alpha correction was applied when necessary.

Dot probe task: Response time. Figure 2 shows RTs to the probe averaged across blocks for Phases 1 and 2. Regarding the Phase 1 results, a 2 (SOA: short vs. long) $\times$ 2 (Predictiveness: probe on predictive cue vs. probe on nonpredictive cue) ANOVA yielded a main effect of predictiveness, $F(1,52) = 6.41$, $p = .014$, $\eta_p^2 = .11$, and a significant predictiveness $\times$ SOA interaction, $F(1,52) = 14.18$, $p < .001$, $\eta_p^2 = .21$. The main effect of SOA was not significant, $F(1,52) = 0.21$, $p = .652$, $\eta_p^2 < .01$. To explore the interaction further, paired t -tests examined the influence of the predictiveness factor at each level of the SOA factor. These revealed that RTs were faster when the probe was positioned over predictive stimuli compared to nonpredictive stimuli at short SOA, $t(26) = 4.03$, $p < .001$, $d = 0.78$, but there was no

significant difference in performance at long SOA, $t(26) = 0.99$, $p = .333$, $d = -0.19$. Thus, predictiveness-driven attentional capture was observed when attention was measured 250 ms after cue onset, and this effect had dissipated by 1000 ms, replicating the findings of Le Pelley et al. (2013).

Probe RTs in Phase 2 were analyzed by repeated-measures ANOVA with factors of predictiveness, SOA, and the uncertainty associated with each compound (certain compounds vs. uncertain compounds). This revealed a main effect of predictiveness, $F(1,52) = 6.96$, $p = .011$, $\eta_p^2 = .12$, with faster RTs when the probe appeared in the location of a predictive cue than a nonpredictive cue. There was also a main effect of uncertainty, $F(1,52) = 43.89$, $p < .001$, $\eta_p^2 = .46$, with faster RTs to probes on trials featuring certain compounds than uncertain compounds. The main effect of SOA was not significant, $F(1,52) = 1.75$, $p = .191$, $\eta_p^2 = .03$, and there were no significant interaction effects, all $F_s(1,52) \leq 1.92$, $p_s \geq .171$, $\eta_p^2_s < 0.4$.

Dot probe task: Accuracy. We analysed accuracy in the dot probe task in order to assess whether the uncertainty and/or predictiveness effects detected in RTs to the probe could be accounted for by a speed-accuracy trade-off. In general, mean accuracy of dot probe responses was high ($M = .990$, $SEM = .01$). First, we assessed whether the SOA $\times$ predictiveness interaction for RTs in Phase 1 was produced by a speed-accuracy trade-off. For this, we conducted a 2 (SOA) $\times$ 2 (Predictiveness) ANOVA on the dot probe accuracy data. This yielded a significant main effect of SOA, $F(1,52) = 5.08$, $p = .028$, $\eta_p^2 = .09$, but no significant effect of predictiveness, $F(1,52) = 2.38$, $p = .129$, $\eta_p^2 = .04$, nor an interaction, $F(1,52) = 0.46$, $p = .499$, $\eta_p^2 < .01$. The main effect of SOA was produced by less accurate responses in the short SOA condition ($M = .987$, $SEM = .002$) than the long SOA condition (M

= .992, $SEM = .002$). This analysis suggests that the $SOA \times$ predictiveness interaction obtained in the RT dependent variable during Phase 1 is not due to a speed-accuracy trade-off.

The same strategy was applied to the Phase 2 data. Since we obtained predictiveness and uncertainty main effects in the RT results, we assessed whether there was any difference in accuracy regarding these two independent variables. A 2 (Predictiveness) $\times$ 2 (Uncertainty) ANOVA revealed no significant main effects nor interactions, all $F_s < 1$. Thus, none of the effects found in the RTs to the probe can be attributed to speed-accuracy trade-offs.

To summarize, RTs were faster when the probe appeared in a location pre-cued by a predictive, rather than a nonpredictive, stimulus. During Phase 1 this effect was evident only in the short SOA condition, and not in the long SOA condition, replicating the effects obtained by Le Pelley et al. (2013). These data are therefore consistent with the idea that predictive stimuli elicit rapid attentional capture, but that this effect also dissipates rapidly, to the extent that it plays no further observable role by 1000 ms. [While the corresponding predictiveness $\times$ SOA interaction did not reach significance in Phase 2, we note that for the certain compounds (which are comparable to those experienced in Phase 1), paired t-tests revealed a significant effect of predictiveness at short SOAs, $t(26) = 2.47$, $p = .020$, $d = 0.47$, but not at long SOAs, $t(26) = 0.02$, $p = .980$, $d < 0.01$].

Attentional models of associative learning (e.g., Esber & Haselgrove, 2011; Le Pelley, 2004; Pearce & Hall, 1980; Pearce & Mackintosh, 2010) predict that the manipulation of uncertainty should affect the attention paid to uncertain compounds experienced during Phase 2. Experiment 1 examined the effect of any such influence of uncertainty on responses in the dot probe task. We observed slower responses to the probe when it was

pre-cued by an uncertain compound than a certain compound, and this effect did not interact with the length of the SOA. These data suggest that the process underpinning this effect of uncertainty was activated rapidly and automatically by the onset of the cues and persisted for at least 1000 ms [Indeed, a significant simple effect of uncertainty was evident in both SOA conditions: short SOA, $t(26) = 6.01$, $p < .001$, $d = 1.16$; long SOA, $t(26) = 3.59$, $p = .001$, $d = 0.69$].

The increase in response times to the probe in the uncertain condition suggests a modulation of attentional processing by these cues, whereby the uncertainty of the associated outcome decreases the attentional resources devoted to detecting the probe. The fact that this effect was observed in the short SOA condition indicates that the effect was rapidly elicited by the onset of the cues. Moreover, this effect was observed in a task which did not require participants to encode the identity of the cues. Cue identity was important only for responding on the AL task, and this response was made after participants had made responded to the dot probe. Indeed, participants were explicitly instructed to not pay attention to the two cues during the dot probe task. The finding of an effect of uncertainty on dot probe responding under these conditions is compatible with the hypothesis that uncertainty-driven attention, like predictiveness-driven attention, reflects a rapid process that operates independently of participants' ongoing task goals. We return to the theoretical discussion of this effect in the General Discussion, but first address possible alternative accounts of this key finding in Experiment 2.

Experiment 2

An alternative account of the effect of uncertainty on dot probe responding observed in Experiment 1 focuses on the conflict that exists between responding in the dot probe task and the AL task. Recall that on each trial participants responded to the position of the probe, and then made a choice response as to which action was correct (*up* or *down*). As a result of the inevitable prediction errors that would be made to the uncertain compounds in Phase 2, participants may have come to associate these compounds with both prediction responses (*up* and *down*). It is possible that these associations then led to greater response competition on the dot probe task: a cue associated with both *up* and *down* may interfere with probe responding more than a cue associated with either the *up* or *down* response. A further alternative account along the same lines appeals to an inhibitory control process. During early phases, participants may develop habitual response tendencies to all compounds (e.g., respond *up* to p1 & n1, respond *down* to p4 & n4, etc.). During Phase 2, the errors that are occasionally produced by these responses for uncertain compounds may then lead to automatic engagement of an inhibitory process that suppresses this habitual tendency. Since response inhibition needs sufficient time in order to have an effect on controlling responses (Osman, Kornblum, & Meyer, 1986), it is likely to be advantageous to start the inhibition process at the moment of cue onset, and so it may well be initiated before the resolution of the dot probe response. Assuming that this inhibitory process generalizes beyond the specific response that is the intention of the inhibition (that is, it results—to some extent—in suppression of all responses, not just the habitual *up* or *down* response), it will lead to a general slowing of response times to the probe on uncertain trials, as observed.

Experiment 2 was designed to minimize any influence of response interference or inhibition produced by uncertainty variations in the AL task on dot probe performance.

Observing an influence of uncertainty on dot probe responding under these conditions would undermine the alternative, non-attentional accounts advanced in the previous paragraph, and thus increase support for the account based on attentional resource re-allocation raised in the Discussion of Experiment 1.

In Experiment 2 we reduced the response interference between the AL and the dot probe tasks by substantially changing the experimental procedure. First, dot probe and AL tasks were programmed to occur on separate and alternating trials, and participants were prompted about which type of task was to be completed before each trial began. Separating the two tasks should reduce the extent to which the AL response tendencies were activated by the cues during the dot probe task (and vice versa), and hence competitive or inhibitory processes elicited by these tendencies should be minimized. In support of this rationale, it is well-established that behavioral effects of interfering habitual responses are diminished when participants have the opportunity to anticipate such a conflict (e.g., Logan & Zbrodoff, 1982). Thus, external signals can be used to bias the relevant set of stimulus-response representations, leading to a reduction in response interference. Clear evidence of a conflict resolution process has also been shown in EEG experiments. These experiments show that neural markers of conflict resolution (e.g., N200) are activated at signals-of-conflict onset, with the magnitude of these markers drastically reduced during the actual conflicting trial (e.g., Correa, Rao, & Nobre, 2009). To further reduce interference between the two tasks, each task had a different set of responses which required the use of different hands.

Splitting the two tasks into separate trials has a further advantage, in terms of minimizing the potential for overlap between task goals. We argued earlier that, in Experiment 1, the best strategy was to ignore cue identity until after the dot probe response

had been made, and only then to identify the cues in order to decide on the correct categorization response for the AL task. Indeed, participants were explicitly informed that they should ignore the cues until after they had responded to the dot probe. However, it is possible that (some) participants may have nevertheless prioritized the AL task and begun preparing their categorization response before the dot probe appeared. Under these conditions, task goals from the AL task—relating to the identity of the cues—may have ‘bled into’ the dot probe task. On this account, the pattern of response times for the dot probe task in Experiment 1 may have resulted from a goal-directed attentional strategy being applied to the AL task. Separating the dot probe and AL tasks in Experiment 2 (and explicitly informing forewarning participants of the task to be performed on each trial) allowed us to rule out this possibility. There was now no strategic reason at all for participants to encode cue identity on dot probe trials, since there was no upcoming categorization response to prepare on these trials.

In sum, the procedure in Experiment 2 was designed to minimize the potential impact of any competition between the responses required by the different tasks, or to generalization of motor response inhibition. In addition, Experiment 2 provided a stronger test of the hypothesis that predictiveness- and uncertainty-driven attentional capture are independent of participants’ ongoing task goals.

Method

Participants and apparatus. The AL task of Experiment 2 was easier for participants to learn as compared with Experiment 1 (see below). Therefore, we expected that not that many participants would fail to meet the data selection criteria. Hence, in order to equal the

number of 'valid' cases in each experiment (aiming 20-25 per condition), we reduced the number of participants in Experiment 2. Fifty-five UNSW Australia students participated for course credit. The apparatus was identical to Experiment 1.

Stimuli. Stimuli were the same as those used in Experiment 1. Since Experiment 2's design needed fewer stimuli (see Table 2), only a subset of the stimuli were used (see Figure 1B; the subset of stimuli used in Experiment 2 are framed).

Design. The design was similar to that of Experiment 1 and included three phases: Pretraining, Phase 1 and Phase 2 (see Table 2). However, Experiment 2 involved half the number of cues and compounds as in Experiment 1; this made the AL stimulus-response contingencies much easier for participants to learn. Specifically, in Pretraining and Phase 1, participants encountered four different compounds and in Phase 2 there were two compounds. One of these compounds (p1 & n1 in Table 2) maintained a certain relationship with the AL response (i.e., R1 was the correct response on all Phase 2 trials of the AL task). The other compound, p2 & n2, transitioned to an uncertain relationship with the previously-established AL response, R2. For example, for participants in the counterbalancing condition in which *down* was the correct response for p2 & n2 trials throughout Phase 1 (i.e., R2 was *down*), the correct response for this same compound was *down* on just two-thirds of the Phase 2 trials, and *up* (i.e., R1) was the correct response on the remaining one-third of trials. Cue locations and probe location were counterbalanced for each compound as in Experiment 1. The pretraining phase included six blocks, with each block comprising eight trials. Each block in Phase 1 comprised 16 dot probe and 16 AL trials, due to the counterbalancing of the probe position. There were seven blocks in Phase 1. Phase 2 comprised 96 dot probe and 96 AL trials. As in Experiment 1, the additional number of trials

as compared with Phase 1 was needed in order to implement the probabilistic relationship of the uncertain condition. The SOA of the probe (250 ms or 1000 ms) was manipulated between-subjects.

Procedure. The procedure was similar to that of Experiment 1, with the following exceptions. During Phase 1 and Phase 2, the AL and dot probe tasks were now programmed as independent, alternating trials. That is, the first trial of each phase involved only the AL task; the next trial involved only the dot probe task; the next trial returned to the AL task, and so on. In both tasks, corrective feedback was provided after every incorrect response.

In order to distinguish the two tasks as clearly as possible, and with the aim of avoiding any response interference, the responses required in the different tasks were more distinct than in Experiment 1. For the dot probe task the responses were again the left and right arrow keys. However, the responses for the AL task were now the keys 'A' and 'Z'. Participants were instructed to respond to the two tasks using different hands (left hand for the AL task and right hand for the dot probe task).

The fixation cross at the start of each trial was replaced by a sign which indicated if the upcoming trial was going to be an AL trial or a dot probe trial. This sign appeared in the centre of the screen for 500 ms before the cues were presented. For the AL trials the sign was the letter 'A' above the letter 'Z'; for the dot probe trials, the sign was two white arrows pointing left and right (see Figure 3A). These signs remained on the screen throughout each trial. The screen with the text 'UP or DOWN?' was not presented in AL trials, since the 'A-Z' sign worked as a reminder for the response options available in these trials.

Results and discussion

Data pre-processing. We used the same selection criterion as in Experiment 1. As a consequence of this criterion, six participants were removed from the final sample (final $N = 49$, with 25 participants in the short SOA condition). As in Experiment 1, RTs from trials with incorrect responses to the probe were not analysed. These responses amounted to 0.8% of all dot probe trials. As in Experiment 1, probe responses that were very fast (under 150 ms) or very slow (over 1000 ms) were also deleted (8% of all trials). Finally, dot probe trials with RTs lying more than 2.5 standard deviations from each participant's mean were also excluded from analysis (0.5% of the remaining trials).

Dot probe task: Response times. Figure 4 shows RTs to the probe averaged across blocks for Phases 1 and 2. Regarding the Phase 1 results, a 2 (SOA: short vs. long) $\times$ 2 (Predictiveness: predictive vs. nonpredictive) ANOVA yielded a main effect of predictiveness, $F(1, 47) = 14.86, p < .001, \eta_p^2 = .24$, and a significant predictiveness $\times$ SOA interaction, $F(1, 47) = 10.14, p < .001, \eta_p^2 = .18$. The main effect of SOA was not significant, $F(1, 47) = 2.14, p = .15, \eta_p^2 = .04$. Paired t -tests at each level of SOA revealed significantly faster RTs when the probe was cued by a predictive stimulus than a nonpredictive stimulus in the short SOA condition, $t(24) = 4.08, p < .001, d = 0.82$, but not in the long SOA condition, $t(23) = 0.70, p = .494, d = 0.14$. This replicates the pattern observed in Experiment 1. It is notable that the effect size for this predictiveness effect in the short SOA condition ($d = 0.82$) was similar to that observed in Experiment 1 ($d = 0.78$), despite the considerable procedural change of separating the two task components into distinct trials.

Regarding the Phase 2 data, an ANOVA with factors of SOA, predictiveness, and uncertainty revealed a main effect of SOA, $F(1,47) = 8.10, p = .007, \eta_p^2 = .15$, with faster RTs

in the long SOA condition. There was also a main effect of uncertainty, $F(1,47) = 7.40$, $p = .009$, $\eta_p^2 = .14$, and a significant uncertainty x SOA interaction, $F(1,47) = 18.22$, $p < .001$, $\eta_p^2 = .28$. To examine this interaction further, comparisons were made between responses to certain and uncertain compounds within each SOA condition. This revealed that RTs were significantly faster to probes appearing over certain compounds than uncertain compounds in the short SOA condition, $F(1,24) = 17.76$, $p < .001$, $\eta_p^2 = .42$, but that there was no significant difference in the long SOA condition, $F(1,23) = 2.03$, $p = .168$, $\eta_p^2 = .08$. No other main effects or interactions in the omnibus ANOVA were significant ($F_s < 1.1$, $p_s > .3$, $\eta_p^2_s < .03$).

Dot probe task: Accuracy. As in Experiment 1, we assessed whether effects detected on the RTs to the probe were due to speed-accuracy trade-offs. In general, the mean accuracy of dot probe responses was very high ($M = .992$, $SEM = .002$). Regarding Phase 1, a 2 (SOA) $\times$ 2 (Predictiveness) ANOVA did not yield any significant effects [SOA: $F(1, 47) = 0.75$, $p = .391$, $\eta_p^2 = .02$; predictiveness: $F(1, 47) = 3.21$, $p = .08$, $\eta_p^2 = .06$; SOA x predictiveness interaction: $F(1, 47) = 2.11$, $p = .153$, $\eta_p^2 = .04$]. Note that the marginal effect of predictiveness reflects a pattern opposite to that expected by a speed-accuracy trade-off (predictive condition: $M = .993$, $SEM = .002$; nonpredictive condition: $M = .989$, $SEM = .003$).

Regarding Phase 2, since we obtained an interaction between SOA and uncertainty factors in the RT results, we assessed any difference in accuracy regarding these two factors. A 2 (SOA) $\times$ 2 (Uncertainty) ANOVA revealed no significant main effects nor interactions, all $F_s(1, 47) \leq 2.77$, $p_s \geq .1$, $\eta_p^2_s \leq 0.06$.

To summarize, the dot probe data from Phase 1 replicate those of Experiment 1 in showing a performance advantage when the probe appeared in the location of a predictive

cue, rather than a nonpredictive cue, but only at the short SOA. This finding therefore offers converging evidence of an influence of within-compound differences in predictiveness on rapid attentional capture. Moreover, Experiment 2 replicated the influence of uncertainty on dot probe RTs in Phase 2, with slower responses when the probe was cued by an uncertain compound than a certain compound. Importantly, these effects were observed even though the associative learning and dot probe tasks were separated in Experiment 2; and indeed, effect sizes were larger than those observed in Experiment 1 despite this change. This suggests that these effects are not a product of response competition or response inhibition resulting from interference between the two tasks. Moreover, the findings of Experiment 2 results reinforce the hypothesis that effects of predictiveness- and uncertainty-driven attention occur independently of participants' ongoing task goals.

The results of Experiment 2 differed from those of Experiment 1 in two notable ways. Firstly, while the influence of uncertainty on Phase 2 dot probe responses did not differ significantly as a function of SOA in Experiment 1, the effect was significantly greater at short SOA than long SOA in Experiment 2. This discrepancy is discussed further in the General Discussion. Secondly, Experiment 2 did not find a significant effect of predictiveness on dot probe responses in Phase 2. This may be a consequence of reduced sensitivity in Experiment 2. Separating the tasks resulted in a general improvement in RT performance on the dot probe task in Experiment 2 (mean RT = 426 ms, as compared to 556 ms in Experiment 1). Consequently, response time may have been nearer to floor in Experiment 2, potentially reducing the sensitivity of the experiment to detect the more subtle effect of predictiveness.

General Discussion

Recent research suggests that the 'derived attention' first described by James (1890/1983) can modulate the extent to which stimuli automatically capture attention in a way that is independent of the physical salience of those stimuli (Anderson, 2013; Awh, et al., 2012; Chelazzi et al., 2013; Le Pelley et al. 2015). These studies have typically considered the impact of learning about the size of the reward that is associated with stimuli. The current study instead investigated the influence of learning about the variability of outcomes that are paired with stimuli; that is, learning about the *uncertainty* associated with stimuli. Previous research and theorizing suggests that such learning engages two related but distinct attentional mechanisms, which we refer to as predictiveness-driven (cf. Mackintosh, 1975) and uncertainty-driven (cf. Pearce and Hall, 1980) processes. The current study aimed to shed light on the nature of these mechanisms; in particular, whether they reflect rapid and relatively automatic, or slower and more controlled, attentional processes.

In two experiments, attention to cues was measured by using a dot probe task that was conducted jointly with an associative learning (AL) task. A learned predictiveness design was used for the AL task in both experiments, in which some cues were perfect predictors of the correct categorization response, while other cues were nonpredictive of the correct response. Our findings replicated those of Le Pelley et al. (2013) in demonstrating that probes appearing in the location of predictive cues elicited faster responses than probes appearing over non-predictive cues. In a second phase we manipulated the uncertainty surrounding some compound cues (with respect to the AL task) and observed a significant effect on response times to probes: when probes appeared over cues in uncertain compounds, responses to probes were significantly slower than when probes appeared over

cues in certain compounds. This novel effect of uncertainty did not interact with predictiveness: responses to probes over cues in uncertain compounds were slow for both the predictive and the nonpredictive cues.

This pattern of results suggests that the perception of an uncertain compound initiates a mechanism which interferes with some of the processes needed for responding to the dot probe task (either detecting the probe, or executing the response). Experiment 2 assessed whether the effect of uncertainty was mediated by an influence on response execution, by reducing the potential for interference between the AL and dot probe tasks. In this experiment, the AL and dot probe tasks occurred on separate and alternating trials, and participants were informed which type of task was to be completed before the trial began. To further reduce the potential interference between the two tasks, each task had a different set of responses which required the use of different hands. This manipulation was clearly successful in reducing response interference between the two tasks in that dot probe responses were generally faster in Experiment 2. Under these conditions, no effect of uncertainty was observed in the long (1000 ms) SOA condition. This suggests that the influence of uncertainty observed in the long SOA condition of Experiment 1 may have been a consequence of interference between the tasks. For instance, uncertain compounds could activate relatively slow controlled inhibitory mechanisms with the aim of diminishing the number of errors in the AL task. When the two tasks were separated in Experiment 2 this inhibitory mechanism would not have been engaged.

Importantly, an uncertainty effect was observed in the short (250 ms) SOA condition of Experiment 2, and the magnitude of this effect was similar to that observed in Experiment 1. Given the separation of the two task components and the use of distinct

motor responses (left and right hands), we argue that this uncertainty effect is unlikely to have been produced by response interference. The observation of the effect in the 250ms SOA condition suggests instead that this attentional modulation was rapidly and automatically initiated by the presence of the uncertain cues. We describe this effect as automatic in nature since orienting attention to cues was not instructed in the dot probe task: there was no need for participants to identify the cue stimuli in order to respond rapidly to the probe, and furthermore participants were explicitly informed that they would gain no advantage by doing so. In other words, the attentional effects were observed despite the ongoing task goals of the dot-probe task (to attend centrally). The suggestion of an automatic influence is compatible with Experiment 2's finding of an effect of uncertainty under a short prime–target SOA of 250 ms but not a long SOA of 1000 ms. Previous priming experiments have shown that the use of controlled response strategies is dramatically reduced under prime–target SOAs of 300 ms or shorter (e.g., Favreau & Segalowitz, 1983; Koivisto, 1997; Neely, 1977; Ortells, Fox, Noguera, & Abad, 2003; Pylkkänen & Marantz, 2003; for a review, see Neely, 1991; for similar results in an associative learning procedure, see Morís, Cobos, Luque, & López, 2014). On this account, the pattern of dot probe responding observed at short SOA would reflect the rapid and automatic effect of uncertainty-driven attention. The longer SOA would allow time for participants to use controlled processes to move attention back to the centre of the screen, in line with the task demands of the dot probe task.

We have described the modulation of attention by uncertainty as leading to an increase in attentional allocation, but yet we observed *slower* response times to probes over uncertain cues (at short SOA) in both experiments. This raises the possibility that one could interpret the data from the current Experiment 1 and 2 as indicating that those cues

associated with uncertainty automatically *repel* attention. By this account, the slower RTs to probes on uncertain trials are due to inattention to the cues in general. However, this account is inconsistent with the findings of Beesley et al. (2015), who showed (using eye-tracking as a measure of attentional processing) that participants spent longer looking at cues in uncertain compounds as compared with cues in certain compounds; this is the opposite of the pattern that would be expected if cues associated with uncertainty repelled attention.

So, Beesley et al. (2015) observed enhanced *orienting* to cues in uncertain compounds, while the current experiments demonstrate a deficit in probe *detection* for cues in uncertain compounds. How can we reconcile this difference, and what does this tell us about the nature of the uncertainty-driven attentional mechanism? We propose that uncertain cues automatically engage an exploratory process of information-gathering: when these cues appear, attention may be drawn to novel features of the stimuli that have not been processed in an attempt to resolve the prediction errors that have been associated with these cues on previous trials (producing the difference in orienting reported by Beesley et al.). Assuming that this exploratory process must be engaged before any response actions are initiated, this account would also explain the increase in dot probe RTs for cues in uncertain compounds observed in the current experiments. In other words, on this account responses to the probe were slowed down because participants' cognitive resources were consumed with a process of exploration for new information from the cues.

Expanding on this idea, Figure 1 shows that the shapes used as cues in these experiments differ in two important respects. Firstly, and most obviously, the stimuli differed in colour. It thus seems likely that participants used colour to distinguish between

predictive and nonpredictive cues during Phase 1. However, some shapes also differed in the thickness of the 'spikes' which projected from the central circle, although these differences were clearly less salient than the differences in colour (see Wang, Yu, & Zhou, 2013). If the introduction of uncertainty in Phase 2 did indeed promote exploration of hitherto-unexplored differences in stimulus features, it seems possible that these spikes might constitute such features. Consistent with this idea, previous work in categorization has shown that changes in attention to the values of a low-salience attribute occur only after the values of another, more salient attribute have already been associated with categories (Kersten, Goldstone, & Schaffert, 1998). Future experimental work will test this account in several ways, for example by manipulating the complexity of the stimuli so as to promote or hinder this exploratory process, or by using forced choice recognition tests against similar foils to test for enhanced memory for different features of uncertain cues.

In a sense, the effect of uncertainty on attentional processing observed in the current experiments might be thought of as a manifestation of the orienting response (OR) to stimuli in animals (e.g., Kaye & Pearce, 1984a, 1984b; Pearce & Kaye, 1985). The OR was first defined by Pavlov (1927) as an investigatory reflexive response to new stimuli. More recently it has been shown that an OR is elicited not only to new stimuli, but also to cues that have recently been associated with prediction error (e.g., Kaye & Pearce, 1984a). Kaye and Pearce (1984b) also showed that the OR operated at the level of the cue-compound: if a compound of two cues included a novel cue and a reliable predictive cue, the OR to the compound was weak and rapidly decreased with further training. Thus, the influence of uncertainty found in the current experiments could be explained as the consequence of an automatic OR elicited by uncertain compounds. This account can also be reconciled with the notion that uncertain stimuli receive enhanced processing in an information-gathering

process: the initiation of an OR would slow down any subsequent responses since attention is focused on the processing of the lesser-known characteristics of the cues (i.e., the peripheral spikes of our cues).

This account of the effect of uncertainty has implications for attentional theories of associative learning. In essence, we are proposing that the experience of an associative prediction error has two different consequences. Firstly, the associative strength of the cue(s) involved will change as a result of an error-correcting learning mechanism, and this may lead to a reduction in the perceived salience of those cues in line with the reduction in their predictiveness. Secondly, and likely in parallel, experience of a prediction error will increase the likelihood that the agent will show an automatic OR to stimuli (or features of the stimuli) that are less well-explored, in line with the increase in the uncertainty regarding the predictive status of those stimuli. This pattern of changes in attention caused by an error signal is not easily reconciled with current attentional theories of associative learning. For instance, dual-process models (e.g., Le Pelley, 2004; Pearce & Mackintosh, 2010) incorporate two salience parameters, one determined by learned predictiveness and the other by uncertainty. Although the rules by which these parameters change are different, both parameters have the same effects on behavior: they modulate the likelihood that selective attention will be allocated to a cue, and—as a consequence—modify the rate of learning about the cue. Thus, both forms of attention would facilitate rapid orienting to cues (which should produce rapid responses in a dot probe task), and would also facilitate new learning about these cues in new contexts. Our results seem to point towards a more radical differentiation between uncertainty- and predictiveness-driven attention. The current data show that uncertainty-driven attention *does not facilitate rapid responses*. On the contrary, it slows down responses because attention appears to be focused on an information-

seeking process, diverting resources away from processing of and/or responding to other stimuli (such as the dot probe target) appearing in the same location. Further evidence for a distinction between the behavioural effects of predictiveness-driven and uncertainty-driven attention comes from the existing literature. While demonstrations of an influence of predictiveness-driven attention on the rate of learning about cues are abundant, there is very little evidence of a similar effect of uncertainty-driven attention on learning rate in studies of humans (Beesley et al., 2015; Kattner, 2015; Le Pelley, Turnbull, Reimers, & Knipe, 2010; for a recent review, see Le Pelley, Mitchell et al., in press).² Hence, it seems that predictiveness-driven and uncertainty-driven processes may reflect two distinct mechanisms, which impact upon behavior in different ways. Specifically, we suggest that predictiveness-driven attention is appropriately characterized by a change in the perceived salience of a cue, wherein more predictive cues have higher salience, and hence are learned about more rapidly. On the other hand, uncertainty-driven attention would increase after prediction errors, and engage the agent in a resource-limiting *attentional exploration* process, where previously ignored cues or features of the environment are attended in a search for further information. It remains for future research to establish whether, and if so under what circumstances, this exploration process influences the salience of the explored (and non-explored) stimulus features, and the rate of subsequent learning about them. It seems likely that these influences may be moderated by such factors as stimulus complexity and changes in context (for further discussion, see Le Pelley, Mitchell et al., in press).

To sum up, we have shown in the current experiments that attention to cues is rapidly and automatically modulated by predictiveness-driven and uncertainty-driven attention. This constitutes a considerable challenge to the usual characterization of uncertainty-driven attentional mechanism as the output of a controlled top-down

mechanism (Pearce & Hall, 1980). Exploration is usually considered as a volitional behaviour, a “...refined capacity, demanding careful regulation.” (Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006, p. 876). In a similar vein, exploratory behaviors are usually portrayed as voluntary actions which “...temporarily suspend routine stimulus-based control and switch the control of the motor apparatus from sensory to volitional input.” (Haggard, 2008, p. 938). In addition, imaging studies have found that exploratory actions are positively correlated with an increase in the activity of prefrontal cortex systems (e.g., Badre, Doll, Long, & Frank, 2012; Daw et al., 2006). Since these areas have previously been associated with cognitive control, these results reinforce the idea that exploration-related behaviours are the consequence of a rational, controlled decision-making mechanism (e.g., Daw, Niv, & Dayan, 2005). Our results for the first time challenge these ideas, suggesting that exploration can be initiated very rapidly and in a relatively automatic way.

References

- Anderson, B. A. (2013). A value-driven mechanism of attentional selection. *Journal of Vision*, 13, 1–16. doi: 10.1167/13.3.7
- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16, 437–443. doi: 10.1016/j.tics.2012.06.010
- Badre, D., Doll, B. B., Long, N. M., & Frank, M. J. (2012). Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron*, 73, 595-607. doi: j.neuron.2011.12.025
- Beesley, T., Nguyen, K. P., Pearson, D., & Le Pelley, M. E. (2015). Uncertainty and predictiveness determine attention to cues during human associative learning. *Quarterly Journal of Experimental Psychology*, 68, 2175-2199. doi:10.1080/17470218.2015.1009919
- Bond, A. B., & Kamil, A. C. (1998). Apostatic selection by blue jays produces balanced polymorphism in virtual prey. *Nature*, 395, 594-596. doi:10.1038/26961
- Chelazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual selective attention. *Vision Research*, 85, 58-62. doi: 10.1016/j.visres.2012.12.005
- Corbetta, M., & Shulman, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3, 201-215. doi: 10.1038/nrn755
- Correa, Á., Rao, A., & Nobre, A. C. (2009). Anticipating conflict facilitates controlled stimulus-

response selection. *Journal of Cognitive Neuroscience*, 21, 1461-1472. doi:

10.1162/jocn.2009.21136doi:jocn.2009.21136

Dayan, P., Kakade, S., & Montague, P. R. (2000). Learning and selective attention. *Nature Neuroscience*, 3, 1218-1223. doi:10.1038/81504

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8, 1704-1711. doi: 10.1038/nn1560

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441, 876-879. doi:10.1038/nature04766

Esber, G. R., & Haselgrove, M. (2011). Reconciling the influence of predictiveness and uncertainty on stimulus salience: A model of attention in associative learning. *Proceedings of the Royal Society B*, 278, 2553-2561. doi:10.1098/rspb.2011.0836

Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-278. doi: 10.1146/annurev.psych.59.103006.093629

Favreau, M., & Segalowitz, N. S. (1983). Automatic and controlled processes in the first- and second-language reading of fluent bilinguals. *Memory & Cognition*, 11, 565-574. doi: 10.3758/BF03198281

Feldmann-Wüstefeld, T., Uengoer, M., & Schubö, A. (2015). You see what you have learned. Evidence for an interrelation of associative learning and visual selective attention. *Psychophysiology*, 52, 1483-1497. doi: 10.1111/psyp.12514

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and

uncertainty by dopamine neurons. *Science*, 299, 1898-1902.

doi:10.1126/science.1077349

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human perception and performance*, 18, 1030-1044. doi: 10.1037/0096-1523.18.4.1030

Gottlieb, J., Oudeyer, P.-Y., Lopes, M., Baranes, A. (2013). Information seeking, curiosity and attention: Computational and neural mechanisms. *Trends in Cognitive Science*, 17, 585-596. doi: 10.1016/j.tics.2013.09.001

Griffiths, O., Johnson, A. M., & Mitchell, C. J. (2011). Negative transfer in human associative learning. *Psychological Science*, 22, 1198-1204. doi: 10.1177/0956797611419305

Haggard, P. (2008). Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience*, 9, 934-946. doi:10.1038/nrn2497

Haselgrove, M., Esber, G. R., Pearce, J. M., & Jones, P. M. (2010). Two kinds of attention in Pavlovian conditioning: evidence for a hybrid model of learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 36, 456-470. doi:10.1037/a0018528

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. doi: 10.1080/17470210701643439

James, W. (1890/1983). *The principles of psychology*. Cambridge, MA: Harvard University Press.

Kattner, F. (2015). Transfer of absolute and relative predictiveness in human contingency

learning. *Learning & Behavior*, 43, 32-43. doi: 10.3758/s13420-014-0159-5

Kaye, H., & Pearce, J. M. (1984a). The strength of the orienting response during Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 90-109. doi:10.1037/0097-7403.10.1.90

Kaye, H., & Pearce, J. M. (1984b). The strength of the orienting response during blocking. *Quarterly Journal of Experimental Psychology*, 36B, 131-144. doi:10.1080/14640748408402199

Kersten, A. W., Goldstone, R. L., & Schaffert, A. (1998). Two competing attentional mechanisms in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 1437-1458. doi: 10.1037/0278-7393.24.6.1437

Koivisto, M. (1997). Time course of semantic activation in the cerebral hemispheres. *Neuropsychologia*, 35, 497-504. doi: 10.1016/S0028-3932(96)00100-5

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. doi: 10.1080/02724990244000179

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. doi:10.1080/02724990344000141

Le Pelley, M. E. (2010). Attention and human associative learning. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour* (pp. 187-215). Oxford: Oxford University Press.

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. doi: 10.1037/a0021384

Le Pelley, M. E., Beesley, T., & Griffiths, O. (in press). Associative learning and derived attention in humans. In R. A. Murphy & R. C. Honey (Eds.), *The Wiley Blackwell handbook on the cognitive neuroscience of learning*. Chichester: Wiley.

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

Le Pelley, M. E., Pearson, D., Griffiths, O., & Beesley, T. (2015). When goals conflict with values: Counterproductive attentional and oculomotor capture by reward-related stimuli. *Journal of Experimental Psychology: General*, 144, 158-171. doi: 10.1037/xge0000037

Le Pelley, M. E., Turnbull, M. N., Reimers, S. J., & Knipe, R. L. (2010). Learned predictiveness effects following single-cue training in humans. *Learning & Behavior*, 38, 126-144. doi: 10.3758/LB.38.2.126

Le Pelley, M. E., Vadillo, M. A., & Luque, D. (2013). Learned predictiveness influences rapid attentional capture: Evidence from the dot probe task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39, 1888-1900. doi: 10.1037/a0033700

Lochmann, T., & Wills, A. J. (2003). Predictive history in an allergy prediction task. In F. Schmalhofer, R.M. Young, & G. Katz (Eds.). *Proceedings of EuroCogSci 03: The European Cognitive Science Conference* (pp. 217-222). Mahwah, NJ: Lawrence Erlbaum Associates.

Logan, G. D., & Zbrodoff, N. J. (1982). Constraints on strategy construction in a speeded

- discrimination task. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 502-520. doi:10.1037/0096-1523.8.4.502
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298. doi: 10.1037/h0076778
- MacLeod, C., Mathews, A., & Tata, P. (1986). Attentional bias in emotional disorders. *Journal of Abnormal Psychology*, 95, 15-20. doi: 10.1037/0021-843X.95.1.15
- Mitchell, C. J., & Le Pelley, M. E. (Eds.). (2010). *Attention and associative learning: From brain to behaviour*. Oxford, UK, UK: Oxford University Press.
- Morís, J., Cobos, P. L., Luque, D., & López, F. J. (2014). Associative repetition priming as a measure of human contingency learning: evidence of forward and backward blocking. *Journal of Experimental Psychology: General*, 143, 77-93. doi: 10.1037/a0030919
- Moors, A., & De Houwer, J. (2006). Automaticity: a theoretical and conceptual analysis. *Psychological Bulletin*, 132, 297-326. doi: 10.1037/0033-2909.132.2.297
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, 106, 226-254. doi: 10.1037/0096-3445.106.3.226
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of correct findings and theory. In D. Besner & G. W. Humphreys (Eds.), *Basic processing in reading: Visual word recognition* (pp. 264-336). Hillsdale, NJ: Erlbaum.
- Ortells, J. J., Fox, E., Noguera, C., & Abad, M. J. F. (2003). Repetition priming from attended vs. ignored single words in a semantic categorization task. *Acta Psychologica*, 114, 185-210. doi: 10.1016/j.actpsy.2003.08.002

- Osman, A., Kornblum, S., & Meyer, D. E. (1986). The point of no return in choice reaction time: controlled and ballistic stages of response preparation. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 243-258. doi: 10.1037/0096-1523.12.3.243
- Oudeyer, P. Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *Evolutionary Computation, IEEE Transactions on*, 11, 265-286. doi: 10.1109/TEVC.2006.890271
- Pavlov, I. P. (1927). *Conditioned reflexes*. Oxford: Oxford University Press.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532-552. doi: 10.1037/0033-295X.87.6.532
- Pearce, J. M., & Kaye, H. (1985). Strength of the orienting response during inhibitory conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 11, 405-420. doi: 10.1037/0097-7403.11.3.405
- Pearce, J. M., & Mackintosh, N. J. (2010). Two theories of attention: A review and a possible integration. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour* (pp. 11-40). Oxford: Oxford University Press.
- Perlman, A., & Tzelgov, J. (2006). Interactions between encoding and retrieval in the domain of sequence-learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 118-130. doi: 10.1037/0278-7393.32.1.118
- Posner, M. I., Nissen, M. J., & Ogden, W. C. (1978). Attended and unattended processing modes: The role of set for spatial location. In H. L. Pick & I. J. Saltzman (Eds.), *Modes of perceiving and processing information* (pp. 137-157). Hillsdale, NJ: Erlbaum.

Pylkkänen, L., & Marantz, A. (2003). Tracking the time course of word recognition with MEG.

Trends in Cognitive Sciences, 7, 187-189. doi: 10.1016/S1364-6613(03)00092-5

Roesch, M. R., Esber, G. R., Li, J., Daw, N. D., & Schoenbaum, G. (2012). Surprise! Neural correlates of Pearce–Hall and Rescorla–Wagner coexist within the brain. *European Journal of Neuroscience*, 35, 1190-1200. doi: 10.1111/j.1460-9568.2011.07986.x

Rushworth, M. F., & Behrens, T. E. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience*, 11, 389-397. doi: 10.1038/nn2066

Shone, L. T., Harris, I. M., & Livesey, E. J. (2015). Automaticity and cognitive control in the learned predictiveness effect. *Journal of Experimental Psychology: Animal Learning and Cognition*, 41, 18-31. doi: 10.1037/xan0000047

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. doi:10.1037/0097-7403.14.3.292

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, 51, 599-606. doi: 10.3758/BF03211656

Wang, L., Yu, H., & Zhou, X. (2013). Interaction between value and perceptual salience in value-driven attentional capture. *Journal of vision*, 13, 5. doi: 10.1167/13.3.5

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. doi:10.1080/02724999208250600

Yantis, S. (2000). Goal-directed and stimulus-driven determinants of attentional control. In S. Monsell & J. Driver (Eds.), *Control of Cognitive Processes: Attention and Performance XVII* (pp. 73–103). Cambridge, MA: MIT Press.

Footnotes

¹ Here, we follow recent literature (e.g., Feldmann-Wüstefeld et al., 2015; Shone, Harris, & Livesey, 2015) in taking automatic as a synonym for ‘independent of ongoing task goals’ (see also Awh et al., 2012; Moors & De Houwer, 2006). We also use automatic to indicate the rapid nature of the effects (e.g., Evans, 2008). The results that we report in this article are compatible with these two characteristics of automaticity. However, we note that stricter definitions of automaticity do exist. For instance, other characterizations of automaticity require that the effect should be, not only independent of the actual goals, but also counterproductive (Perlman & Tzelgov, 2005). The procedures used in the current experiments do not allow us to determine whether our effects would still be considered automatic under such stricter criteria.

² Notably, studies with nonhuman animals provide stronger evidence for an influence of uncertainty on learning rate (e.g., Kaye & Pearce, 1984a; Haselgrove, Esber, Pearce, & Jones, 2010). Indeed, it has been proposed that attentional mechanisms relating to uncertainty could differ fundamentally across species (Haselgrove et al., 2010).

Figure legends

Figure 1. Panel A shows the trial structure for Phases 1 and 2 of Experiment 1. Participants were instructed to respond the position of the probe (left or right) as fast as possible. This probe was presented 250 or 1000 ms after cue onset, depending on the SOA condition. After the dot probe response, participants made a predictive response for the associative learning task (up or down). If this predictive response was incorrect, error feedback was provided (not shown). Panel B shows the eight stimuli used as cues. The four stimuli also used for Experiment 2 are framed. Mean response times (panel C) and mean proportions of 'probable outcome' responses (panel D) are shown for the associative learning task. Error bars represent standard error of the mean.

Figure 2. Mean response times to the probe in the Experiment 1. The upper panel shows results in the short SOA condition (SOA = 250 ms); the lower panel shows results for the long SOA condition (SOA = 1000 ms). Error bars represent standard error of the mean.

Figure 3. Panel A shows the trial structure for the associative learning task (top) and the dot probe task (bottom). Feedback was provided when participants made an error on the associative learning task (feedback screen not shown). Panel B shows the mean response times and panel C shows the proportion of 'probable outcome' responses for the associative learning task. Error bars represent standard error of the mean.

Figure 4. Mean response times to the probe in Experiment 2. The upper panel shows results in the short SOA condition (SOA = 250 ms); the lower panel shows results for the long SOA condition (SOA = 1000 ms). Error bars represent standard error of the mean.

Table 1. Experiment 1 design

Pretraining	Phase 1	Phase 2
4 blocks x 8 trials	8 blocks x 32 trials	96 trials
Only associative learning task	Dot probe task and associative learning task	
p1 & n1 → R1	p1 & n1 → R1	<u>Certain compounds</u>
p1 & n2 → R1	p1 & n2 → R1	p1 & n1 → R1
p2 & n1 → R2	p2 & n1 → R2	p2 & n2 → R2
p2 & n2 → R2	p2 & n2 → R2	
p3 & n3 → R1	p3 & n3 → R1	<u>Uncertain compounds</u>
p3 & n4 → R1	p3 & n4 → R1	p3 & n3 → R1 (67%) / R2 (33%)
p4 & n3 → R2	p4 & n3 → R2	p4 & n4 → R2 (67%) / R1 (33%)
p4 & n4 → R2	p4 & n4 → R2	

Note. p1–p4 denote cues that were predictive with regard to the associative learning task during Pretraining and Phase 1; n1–n4 denote cues that were nonpredictive. R1 and R2 denote the correct categorization response for each cue-compound (*up* and *down*, counterbalanced). During Phase 2, *Certain compounds* are those for which the same categorization response was correct through Phase 2; *Uncertain compounds* are those for which one categorization response was correct on two-thirds of appearances in Phase 2, while the other response was correct on one-third of appearances. The probe was equally likely to appear in the location of all cues during Phase 1 and Phase 2.

Table 2.Experiment 2 design

Pretraining	Phase 1	Phase 2
6 blocks x 8 trials	7 blocks x 16 trials	96 trials
Only associative learning task	Dot probe task and associative learning task	
<p>p1 & n1 → R1</p> <p>p1 & n2 → R1</p> <p>p2 & n1 → R2</p> <p>p2 & n2 → R2</p>	<p>p1 & n1 → R1</p> <p>p1 & n2 → R1</p> <p>p2 & n1 → R2</p> <p>p2 & n2 → R2</p>	<p><u>Certain compound</u></p> <p>p1 & n1 → R1</p>
		<p><u>Uncertain compound</u></p> <p>p2 & n2 → R2 (67%) / R1 (33%)</p>

Note. p1–p2 denote cues that were predictive during Pretraining and Phase 1; n1–n2 denote cues that were nonpredictive. R1 and R2 denote the correct categorization response for each cue-compound (*up* and *down*, counterbalanced). During Phase 2, *Certain compounds* are those for which the same categorization response was correct throughout Phase 2; *Uncertain compounds* are those for which one categorization response was correct on two-thirds of appearances in Phase 2, while the other response was correct on one-third of appearances. The probe was equally likely to appear in the location of all cues during Phase 1 and Phase 2.