Attention and associative learning in humans: An integrative review

Mike E. Le Pelley[1*], Chris J. Mitchell[2*], Tom Beesley[1],

David N. George[3], Andy J. Wills[2]

[1] UNSW Australia, Sydney, Australia

[2] Plymouth University, Plymouth, UK

[3] University of Hull, Hull, UK

[*] These authors contributed equally to this work

Running head: Human attentional learning

Please address correspondence to:

Dr Mike Le Pelley

School of Psychology, UNSW Australia

Sydney NSW 2052

Australia

Tel: +61 2 9385 1294

Email: m.lepelley@unsw.edu.au

Acknowledgment

Abstract

This article presents a comprehensive survey of research concerning interactions between associative learning and attention in humans. Four main findings are described. First, attention is biased towards stimuli that predict their consequences reliably (*learned predictiveness*). This finding is consistent with the approach taken by Mackintosh (1975) in his attentional model of associative learning in non-human animals. Second, the strength of this attentional bias is modulated by the value of the outcome (*learned value*). That is, predictors of high-value outcomes receive especially high levels of attention. Third, the related but opposing idea that *uncertainty* may result in increased attention to stimuli (Pearce & Hall, 1980), receives less support. This suggests that hybrid models of associative learning, incorporating the mechanisms of both the Mackintosh and Pearce-Hall theories, may not be required to explain data from human participants. Rather, a simpler model, in which attention to stimuli is determined by how strongly they are associated with significant outcomes, goes a long way to account for the data on human attentional learning. The last main finding, and an exciting area for future research and theorizing, is that *learned predictiveness* and *learned value* modulate both deliberate attentional focus, and more automatic attentional capture. The automatic influence of learning on attention does not appear to fit the traditional view of attention as being either *goal-directed* or *stimulus-driven*. Rather, it suggests a new kind of "derived" attention.

KEYWORDS: associative learning; attention; conditioning; reward learning

This article examines the relationship between two central phenomena in human cognitive psychology, attention and learning. Attention refers to the set of perceptual and cognitive processes that allows us to prioritise certain events for further analysis or action; as such, attentional processes play a role in almost all psychological phenomena. The focus of this article is the way in which attention to stimuli, and hence our prioritization of those stimuli, can change as a result of learning about them. This issue is of central importance, since across areas as diverse as social interaction, decision making, reasoning, and perception, people often experience stimuli on multiple occasions with important consequences. The evidence to be reviewed in this article shows that this experience may fundamentally change the way in which people perceive and interact with these stimuli, which will have an impact on understanding behaviour in these and all other areas of psychology.

The research conducted thus far in the domain of attentional learning ranges from the basic science of visual search, spatial cueing and attentional blink, through causal learning and the neuroscience of uncertainty to clinical applications such as the special hold that drug cues have over addicts. We review this literature, with a particular emphasis on the role of top-down and bottom-up attentional processes. We also offer a simple model to account for most of the current data and suggest some important avenues for future investigation.

Interestingly, research concerning interactions between attention and learning in humans has (to date) proceeded largely independently in two different psychological traditions. One stream has its roots in early descriptions of 'attentional' phenomena in animal conditioning. Here, past associative learning appears to change the way in which animals attend to their environment, showing that attention can be conditioned just as overt responses (lever-pressing, salivation etc) can be conditioned. This change in attention is then revealed by the way it influences later learning. Such findings have since been replicated and extended in studies of human associative learning. A second stream has arisen more recently in the

perception–cognition literature, in particular in the area of visual search. This work has shown that learning about stimuli influences subsequent search performance involving those stimuli. Our aim is to demonstrate that these two research streams (associative and perceptual-cognitive) are mutually complementary, and that considering both provides a fuller picture of the relationship between attention and learning than either in isolation.

The review is made up of four main sections. The first two sections examine models of attention and learning from the associative/animal conditioning tradition. One model proposes that predictive cues—those that signal reliable consequences—command more attention than non-predictive cues (Mackintosh, 1975). The second model proposes the opposite; that most attention is paid to cues whose consequences are uncertain (Pearce & Hall, 1980). We find that, in humans, there is a great deal of evidence in favour of the former, predictiveness principle and relatively little evidence to support the latter, uncertainty principle. The third empirical section examines the perception-cognition literature. Here, we find that people pay more attention to predictors of valuable outcomes than predictors of non-valuable outcomes. Overall, therefore, the current evidence suggests people tend to pay attention to cues that have reliably predicted valued outcomes in the past. Another major finding from this survey is that there is evidence for both top-down and bottom-up attentional processes with respect to both predictiveness and learned value. Hence, people deliberately pay attention to cues that predict high-value outcomes and, perhaps more interestingly, attention is also *captured* by these cues in a way that is outside of strategic control. The final section offers some theoretical analysis of the empirical evidence reviewed, and proposes a simple model of the roles of both predictiveness and learned value in human associative learning. We first provide a little historical context for the review and define its scope.

Background and Scope of the Current Article

The suggestion that the influence of attention might usefully be studied in the context

of learning was alluded to by Pavlov (1927), who noted that presenting a novel stimulus to an animal often elicited an orienting response (which he termed an "investigatory reflex", p12). The implication was that this orienting response reflected the animal's attention to the stimulus, with a larger orienting response reflecting greater attention to the stimulus. An alternative way of measuring attention was introduced by Lashley (1929), who noted that:

> "when any complex of stimuli arouses nervous activity, that activity is immediately organised and certain elements or components become dominant for reaction while others become ineffective… Such an organisation is in part described by… principles of attention. In any trial of the training series, only those components of the stimulating situation which are dominant in the organisation are associated. Other stimuli which excite the receptors are not associated because the animal is not set to react to them." (p242)

Here, Lashley argues that the extent to which a stimulus is learned about on a given trial will depend on the attention allocated to that stimulus. Hence, the speed with which a cue–outcome association is learned can be taken as a measure of the attention paid to the cue. Lashley's view was that the influence of attention was all or none: a stimulus could either be attended, in which case it would be learned about, or unattended, in which case nothing would be learned about it. This later gave way to a less extreme view (see Sutherland & Mackintosh, 1971) on which the allocation of attentional resources occurs on a continuum, such that differences in attention give rise to graded differences in the rate of learning about stimuli, otherwise referred to as the *associability* or *conditionability* of those stimuli.

These suggestions that attention might be studied by measuring orienting to cues (Pavlov) or their associability (Lashley) were originally made in the context of research in nonhuman animals. However, directly analogous approaches have since been used to study the interaction between attention and learning in humans. In subsequent sections we describe many studies of human learning that adopt one or both of these measures of attention.

The Scope of the Current Survey

As the earlier summary suggests, our interest here is in the way that learning affects attention; how does previous experience with a cue impact on the amount of attention paid to that cue later? Of course, another way in which attention and learning interact is that the intrinsic *physical salience* of a cue will affect the attention paid to it and this will, in turn, affect learning about the cue. Physical salience refers to the distinctiveness of a cue based on its physical properties; e.g., a cue may be salient based on its intensity (loudness, brightness, suddenness), or its dissimilarity from other items (e.g., a single red item among blue items). Although physical salience is not a major focus of the current article, it is worth mentioning briefly because experiments in this domain raise issues that will be important later.

Perhaps unsurprisingly, very intense cues (those that can be expected to receive most attention) tend to be learned about more rapidly than less intense cues. Since this finding is so uncontroversial—very intense cues tend to be prioritised whatever the task—we will examine it only briefly here. For example, Beck (1963) exposed participants to pairings of a tone with an air-puff directed at the eye. Participants were faster to develop a conditioned blink response to the tone if it was loud than if it was soft. That is, a cue's associability appears to be influenced by its intrinsic physical salience. This illustrates an idea that will be central to the arguments that follow – that associability, physical salience and psychological salience are intimately linked. One way to explain the observation that physically salient cues are learned about quickly, is to say that those stimuli are more psychologically salient: they have a greater impact on the perceptual system; they are more easily noticed or preferentially selected. It is well-known that stimulus selection can be driven by the perceptual properties of a stimulus – a stimulus is more likely to capture attention if it has an abrupt onset (Yantis & Jonides, 1984), or bright colour (Theeuwes, 1992). Such attentional biases driven by physical salience might be expected to affect the speed of learning about cues.

However, the finding that intense cues produce high responding on test does not, in itself, show that associability is increased by stimulus intensity. It could instead be that Beck's (1963) participants all *learned* at the same rate, but that responding to the loud tone on test was greatest because it was the most intense stimulus. That is, the effect observed may have been on performance, not on learning. Stronger evidence for an influence of stimulus intensity on learning (rather than performance) comes from the *overshadowing* effect in conditioning, first reported by Pavlov (1927) in animals and since demonstrated many times in humans. For example, in a study by Denton and Kruschke (2006), participants saw cues made up of dot patterns and were required (through trial and error) to learn the correct responses to those patterns (R1, R2 or R3). The patterns could be either of low physical salience (few dots in the display) or high salience (many dots). For one group of participants, a high-salience pattern in one colour (say red) was paired with a low-salience pattern in a second colour (say green). The correct response to this pair of patterns (known as a cue compound) was always response R1. On test, the low-salience green pattern was presented alone, and accuracy in producing response R1 was a poor 38% (chance responding was 33%). In contrast, for a second group of participants, the low-salience green pattern was trained with a *low*-salience red pattern. When this group was presented with the green cue on test, they were accurate in producing the correct R1 response (82% correct). Thus the accuracy of the response to exactly the same green stimulus depended on the salience of the stimulus that accompanied it during training. The implication is that the presence of the high-salience red pattern (but not the low-salience red pattern) overshadowed learning about the green pattern.

This result follows naturally from an attentional interpretation (though other accounts of overshadowing have been offered, e.g., see Broomfield, McHugh & Reed, 2010; Miller & Matzel, 1988; Rescorla & Wagner, 1972). For the group trained with a compound of the low-salience green pattern and high-salience red pattern, the red pattern will command attention at

the expense of the green pattern during training, by virtue of its greater physical salience. Reduced attention to the green pattern will hinder learning about its relation with the response. A low-salience red pattern, on the other hand, will draw less attention away from the green pattern, and so learning about the green pattern will be more likely.

It is clear, then, that stimulus intensity and physical salience can influence learning. High intensity cues produce greater conditioned responding and prevent learning about other cues. These observations are not, in themselves, surprising. They do however highlight an issue that will be of central importance in the review presented here – the notion of competition. In overshadowing, a physically salient cue competes with a less salient cue in learning. Competition may be an important determinant, not only of the extent to which learning takes place, but also the way in which learning in turn alters psychological salience, by modulating attention.

Our concern in this article is the impact that associative learning, specifically, has on attention. There are, therefore, ways in which attention can change through experience that are not covered in this review. One large literature that will not be covered concerns the changes in attention that are brought about through mere exposure to cues, rather than the pairing of the cues with some outcome. For example, familiar stimuli are commonly learned about less quickly than novel stimuli—a phenomenon known as latent inhibition—consistent with a reduction in attention to those stimuli. Mere exposure can also change the way that stimuli are perceived (known as perceptual learning), which will have implications for later learning. Notably, while there do exist accounts of latent inhibition and perceptual learning that implicate associative learning and attention (e.g., Kruschke, 2001b; Lubow, 1989), there are also many alternative accounts that do not (e.g., Bouton, 1993; Miller & Matzel, 1988). Given the ambiguity of this research with regard to the issues addressed in the current article, we will not cover it here: for recent reviews of the literature on latent inhibition and

perceptual learning, we suggest Holmes and Harris (2010) and Mitchell and Hall (2014) respectively. Lastly, ideas related to those presented here also appear in other psychological disciplines that share some similarity with associative learning, but these other disciplines will not receive detailed treatment here. In particular, an important role for attention (like that described in the section on predictiveness next) has been identified in work on human categorization (see e.g., George & Kruschke, 2012 for review).

Having described what the current review is not about, we now consider the scope of the first two major empirical sections. These two sections examine the ways in which attention changes as a consequence of the predictiveness of cues. A predictive cue is one that is a consistent and reliable indicator of the events that follow it, whether these events refer to presence of an outcome (e.g., electric shock) or its absence (no shock). A nonpredictive cue is one that provides no information regarding the events that follow it (e.g., a stimulus that is sometimes followed by shock, and sometimes by no shock). Two (contrasting) views of the relationship between predictiveness and attention have been dominant in the associative learning literature. We refer to these as the *predictiveness principle* (attention increases towards cues when their outcomes are predicted) and the *uncertainty principle* (attention increases towards cues when their outcomes are surprising).

These principles are exemplified by theories proposed by Mackintosh (1975) and Pearce and Hall (1980) respectively. In searching for literature relating to these principles, we searched the ISI Web of Science database on 10/10/2015 for all articles in the fields of psychology, neuroscience, psychiatry, behavioral sciences and multidisciplinary sciences citing either of these key papers, and containing at least one of these keywords (where * denotes a wild card, such that associab* will pick up 'associability' and 'associable'): attention*, predict*, associab*, relevan*, irrelevan*, certain*, uncertain*. The resulting 828 hits were then reviewed for content and for references to further literature (28 extra articles

found), yielding a total of 143 articles relevant to the current discussion (marked as results of Search 1 in the reference list) which we cover in the sections to follow. Reasons for exclusion were: (1) The article did not present novel empirical data (articles in this category were reviews or modelling papers) – 137 articles; (2) The article described data from nonhuman animals only, whereas our focus is on humans – 366 articles; (3) The article presented data from humans but these data did not bear on the relationship between learning and attention (most articles in this category made only a passing reference to one of the key papers as an example of a theory of learning) – 171 articles; (4) The article concerned effects of mere exposure (latent inhibition or perceptual learning), which lie outside the scope of the current article as noted earlier – 39 articles. The first main section of this review looks at the evidence for the predictiveness principle in the human learning literature.

Predictiveness and Attention

One idea that has been dominant in the literature is that predictive stimuli—those with meaningful consequences—might 'stand out'. William James (1890/1983) introduced the concept of *derived attention*; a form of attention to a stimulus that "owes its interest to association with some other immediately interesting thing" (p393). Consistent with this idea, the predictiveness principle refers to the idea that cues become more psychologically salient as a result of their predictiveness with respect to important outcomes; more attention will be allocated to predictive cues than to nonpredictive cues.

The archetypal model of associative learning that incorporates the predictiveness principle, and its effect on associability, is that proposed by Mackintosh (1975), though other variants have been proposed before and since (e.g., Kruschke, 1992, 2001b, 2006; Lovejoy, 1968; Schmajuk & Moore, 1985; Sutherland & Mackintosh, 1971; Trabasso & Bower, 1968; Zeaman & House, 1963). Hence, when we talk of the Mackintosh model below we are

actually referring to a broad class of theories that embody the predictiveness principle.

Behavioural Evidence for the Predictiveness Principle

In this section we review studies that have provided evidence consistent with the predictiveness principle embodied by the Mackintosh model. These studies have examined three closely related phenomena of associative learning: 'simple' *learned irrelevance*, the *intradimensional–extradimensional shift effect*, and the *learned predictiveness effect*. We will see that there is good evidence for the idea that predictive cues attract more attention than non-predictive cues. Perhaps the strongest evidence is provided by the learned predictiveness effect (the third of the lines of evidence presented below). Here, previously predictive cues are learned about more quickly than previously non-predictive cues, and they also command more attention as measured by tracking eye-gaze. Following this discussion, we look more closely at learned predictiveness and ask whether the effect is the consequence of deliberate top-down attentional processes, or more automatic, bottom-up attentional capture.

Simple learned irrelevance. The simplest approach borrows directly from the learned irrelevance procedure originally developed by animal researchers (e.g., Mackintosh & Little, 1969). In a human version of this procedure (Myers, Oliver, Warren & Gluck, 2000), participants had to predict, on each trial, whether an outcome would occur (a rabbit appearing under a magician's hat). During training phase 1, for participants in a pre-exposed group, a red or green balloon was also presented on each trial, but the colour of this balloon was unrelated to whether the outcome occurred, which was randomly determined on each trial. That is, the colour of the balloon was non-predictive of the outcome. Then, in training phase 2, the colour of the balloon became predictive (e.g., whenever the balloon was red the rabbit appeared; whenever it was green no rabbit appeared). Participants who had been pre-exposed to the non-predictive balloon were slower to learn this colour–outcome relationship in phase

2 than control participants who had not been exposed to the coloured balloon during phase 1. This finding has been well replicated (Myers, Warren, et al., 2000; Orosz, Feldon, Gal, Simon & Cattapan-Ludewig, 2007; Orosz, Feldon, Gal, Simon & Cattapan-Ludewig, 2008; Orosz et al., 2011).

This finding follows naturally from the predictiveness principle. On this view, experience that the coloured balloon is non-predictive (i.e., irrelevant) with regard to the outcome during phase 1 results in a reduction in attention to this feature. If participants are paying less attention to the balloon, they will naturally be slower to learn its significance when it becomes relevant during phase 2; i.e., its associability will be reduced. In essence, this account proposes that pre-exposed participants learned to ignore the irrelevant balloon during phase 1. We refer to this effect as 'simple' learned irrelevance to contrast it with more complex examples discussed in the following sections.

However, the support for the predictiveness principle offered by this simple learned irrelevance effect is equivocal, because it is open to other interpretations. Firstly, the pre-exposed and non-pre-exposed (control) group differ not only in their experience of the predictiveness of the coloured balloon, but also in terms of its novelty. So the difference in learning of the colour–outcome relationship could stem from the fact that, for the control group, the coloured balloon is novel at the start of phase 2, while for the pre-exposed group it is not (cf. Bonardi & Hall, 1996). Moreover, the objective statistical relationship between the colours and the outcome differs between the two groups. For the control group, every presentation of the red balloon is paired with the rabbit. In contrast, the pre-exposed group have previously experienced trials on which the red balloon was not paired with the rabbit (during phase 1). Hence the overall conditional probability of the outcome given the red balloon is lower in the pre-exposed group during phase 2, and this may be why people in this group perceive the colour–outcome relationship as weaker (see Schmidt-Hansen & Le Pelley,

2012). Notably, neither of these accounts make any appeal to changes in attention as a result of learning about predictiveness. Stronger evidence for the predictiveness principle comes from studies that control more carefully for such alternative explanations.

The intradimensional–extradimensional shift effect. The studies described above investigate how previous experience of the relationship between a cue and an outcome influences subsequent learning about that same relationship. Another approach has studied how prior experience influences learning of a *new* cue–outcome relationship. An example is provided by studies of the *intradimensional–extradimensional shift* (ID-ED shift) effect.[1] For example, in the study by Eimas (1966), children learned to discriminate between cues that varied on two dimensions: colour (red or green) and shape (circle or star). Each trial presented a pair of cues that shared no features (i.e., the pair was either red circle and green star, or red star and green circle); children tried to pick the 'correct' cue, behind which was a winning chip. During the first discrimination task (D1), one of the dimensions predicted the solution of the discrimination. For group IDS, colour was the predictive dimension; e.g., the chip may always be behind the red cue, and never behind the green cue. Shape was nonpredictive, so across trials the chip was equally likely to be behind the circle or the star. In contrast, for group EDS shape was the predictive dimension (e.g., the chip was always behind the star and never behind the circle) and colour was nonpredictive.

Following training on discrimination D1, children were shifted to a new discrimination, D2, in which the cues contained new values on each dimension (colours were now blue and yellow, shapes were triangle and diamond). For both groups, colour was predictive and shape nonpredictive for dimension D2. So for group IDS, the predictive dimension for discrimination D2 was the same as that which was predictive for discrimination D1: these children experienced an intradimensional shift. For group EDS, the predictive dimension for D2 (colour) had been *nonpredictive* for D1: these children experienced an extradimensional

shift. Eimas (1966) found significantly faster learning of discrimination D2 by group IDS than group EDS. In other words, group IDS learned that colour was predictive in the first phase, which gave them an advantage in the second phase when colour was again predictive.

The ID-ED shift effect emerges naturally from the predictiveness principle; the idea that more attention is allocated to predictive than nonpredictive stimuli. For group IDS, training on discrimination D1 will result in greater attention to (and hence associability of) the predictive dimension of colour than to the nonpredictive shape. This will aid learning of discrimination D2, since D2 also relies on learning about colour, while shape provides no useful information. In contrast, for group EDS, training on D1 will result in lower attention to colour than to shape, which will impair learning about colour on discrimination D2. So overall, this account correctly anticipates faster learning of D2 in group IDS than group EDS.

The ID-ED shift procedure overcomes the problems with the simple learned irrelevance procedure in the previous section. Group IDS and EDS encounter the different stimuli exactly the same number of times, so stimulus novelty is equated in the two groups. Moreover, since all of the features are novel in discrimination D2 (none of blue, yellow, triangle, or diamond was experienced in D1), there can be no direct carryover of stimulus–outcome learning from D1 to D2. In other words, the objective stimulus–outcome relationships during D2 are identical for the two groups. The groups differ only in their prior experience of the predictiveness of the different dimensions (D1 training), and so this must drive the difference in acquisition of D2 – as anticipated by the predictiveness principle.

The ID-ED shift effect is consistent with the idea of greater attention to predictive cues than nonpredictive cues. A further question arises as to whether experience of predictiveness increases attention to cues (learned relevance) or nonpredictiveness decreases attention (learned irrelevance), or both. These possibilities can be tested by comparing previously predictive or nonpredictive dimensions with novel dimensions which were not experienced

during D1. Studies taking this approach have revealed mixed results: some have found stronger evidence for learned relevance than learned irrelevance (Lawrence, Sahakian, Rogers, Hodges & Robbins, 1999; Owen et al., 1993), others have found stronger evidence for learned irrelevance than learned relevance (Cools, Rogers, Barker & Robbins, 2010; Gauntlett-Gilbert, Roberts & Brown, 1999; Maes et al., 2004), and one study found similar effects of both (Goldstone & Steyvers, 2001). Moreover, there are limitations to these datasets and to this approach more generally. Firstly, only Goldstone and Steyvers and Maes et al. provide statistics that allow the size of the two effects to be compared; for the other studies, the patterns are numerical but not necessarily statistically significant. Secondly, Owens et al. and Lawrence et al. did not equate prior exposure to predictive and nonpredictive dimensions. Finally, the comparison of predictive or nonpredictive dimensions against a novel (i.e., unfamiliar) dimension in discrimination D2 is problematic in principle; it raises the possibility that a difference in familiarity (rather than predictiveness) may influence the rate of learning about different dimensions during D2. As such, the extent to which the overall ID-ED shift effect reflects learned relevance or irrelevance remains unclear, and it may not be possible to resolve this issue for certain. As a consequence, for the remainder of this article we use the term *learned predictiveness* without making any specific commitment to whether this reflects changes with regard to predictive stimuli, nonpredictive stimuli, or both.

Finally in this section we note an issue regarding the level at which attention operates. The ID-ED shift effect shows that learning can influence the attention that is paid to a stimulus *dimension* (e.g. colour, shape) (see also Goldstone, 1994; Nosofsky, Gluck, Palmeri, McKinley & Glauthier, 1994). That is, experience of the predictiveness of certain feature-values on the dimension during discrimination D1 (e.g., red, green) has an effect on attention to new features from that same dimension in discrimination D2 (blue, yellow). One possible

interpretation of this is that the effect of learning on attention operates at the level of dimensions; learning that red and green are predictive results in greater attention to the dimension of colour (for examples of models implementing this approach, see Kruschke, 1992; Nosofsky, 1986; Sutherland & Mackintosh, 1971). The alternative view, embodied by Mackintosh's (1975) model, has attention operating at the level of individual features; learning that red and green are predictive results in greater attention to red and green (cf. Kalish & Kruschke, 2000). In order to account for the ID-ED shift effect (where the features change between training and transfer discriminations) Mackintosh suggested that learned attention could generalize between features based on their similarity, and this similarity will typically be greater for features belonging to the same dimension (red and yellow) than features belonging to different dimensions (red and triangle).

The ID-ED shift effect is amenable to either of these accounts – attention at the level of dimensions, or features. Other phenomena are less ambiguous, however; in many cases below (including the next example) all cues come from the same dimension, and participants must learn which particular features on that dimension are predictive and which are nonpredictive. In these examples it is clear that variations in learned attention apply to individual features, and not to dimensions.

The learned predictiveness effect. The ID-ED shift effect prevents direct carryover of stimulus–outcome learning from D1 to D2 by changing the cues between the two phases (e.g., colours are red and green in D1, but blue and yellow in D2). An alternative option is to change the outcomes between the two phases, and this is the approach taken in studies of the so-called *learned predictiveness effect*, developed by Le Pelley and McLaren (2003) based on a procedure originally used by Lochmann and Wills (2003).

On each trial of Le Pelley & McLaren's (2003) study, participants were told the foods eaten by a fictitious patient, and had to predict what sort of allergic reaction the patient would

suffer as a result. Immediate feedback allowed participants to learn the relationships between foods (cues) and different types of allergic reactions (outcomes). In training Phase 1, participants learned about patient 'Mr X', who suffered from either sweating or dizziness. During this phase, certain cues (e.g., apple) consistently predicted which reaction would occur. Others (e.g., banana) were paired equally often with sweating or dizziness and so were nonpredictive. Training Phase 2 introduced a new patient ('Mr Y') who ate the same foods as Mr X but suffered different types of reactions (nausea and itching). Specifically, each trial featured a pair of foods, one of which had previously been predictive for Mr X, and the other of which had been nonpredictive. Thus, using the examples above, participants might experience trials on which Mr Y ate apples and bananas, and suffered nausea as a result. Crucially, during Phase 2, all cues were now equally predictive of the outcomes with which they were paired. So both apples and bananas predicted perfectly that nausea would occur. Despite this, a final test revealed that participants had learned more about the previously predictive cues than the previously nonpredictive cues; e.g., they learned the relationship between apples and nausea more strongly than the relationship between bananas and nausea.

This finding suggests that the predictiveness of cues established in Phase 1 produced a persistent bias in the associability of those cues that carried over to, and influenced learning in, Phase 2. The advantage for predictive cues over nonpredictive cues is consistent with the idea that experience of predictiveness influences attention in the manner suggested by the predictiveness principle. And like the ID-ED shift effect, the learned predictiveness effect is not confounded with stimulus familiarity (since all cues are presented the same number of times) or differences in the objective cue–outcome relationships (since all cues are equally predictive of the outcomes with which they are paired in Phase 2).

This learned predictiveness effect demonstrated by changing outcomes between two phases of training has been replicated many times.[2] Extensions of the basic procedure suggest

that: (1) Associability can be influenced by absolute predictiveness (how well do apples predict sweating?) rather than a comparison of relative predictiveness (do apples predict sweating better than bananas do?) (Beesley & Le Pelley, 2010; Kattner, 2015; Le Pelley, Turnbull, et al., 2010; Livesey et al., 2011). (2) The influence of predictiveness can be contextually mediated. Suppose that apples are predictive and bananas are nonpredictive in context C1 (e.g., a particular grocery), but apples are nonpredictive and bananas predictive in context C2. Under these circumstances, people will learn new cue–outcome relationships faster for apples than bananas if these are encountered in context C1, but will learn faster about bananas than apples in context C2 (Lucke et al., 2013; Uengoer & Lachnit, 2012; Uengoer, Lachnit, Lotz, Koenig & Pearce, 2013: see George & Kruschke, 2012, and Kruschke, 1996, for similar findings from an ID-ED shift procedure). These findings suggest that, under some circumstances, attention can operate at the level of cue compounds, e.g., the configuration of apples and context C1 is predictive, whereas the configuration of bananas and context C1 is nonpredictive, so the former configuration receives more attention than the latter (see George & Pearce, 1999, p372).

Beyond Associability: Evidence from Eye Gaze

The studies cited above follow Lashley's (1929) lead in using the associability of a cue as an index of the attention. This focus on associability reflects the origins of studies of learned predictiveness in the animal conditioning literature; conditioning researchers are practised in measuring learning, and so they naturally use rate of learning as a proxy for attention. However, the focus on associability separates this literature from studies of attention as it is typically understood in the perceptual–cognitive tradition. So we might ask, does learning about predictiveness influence *only* the associability of cues, or does it also influence other aspects of stimulus processing that are more typically taken as evidence of changes in attention? Several recent studies have addressed this question, and have shown that

predictiveness modulates stimulus processing in a way that is characteristic of attention as it is understood in the wider perceptual–cognitive literature. Here, we shall first briefly consider studies that have used eye-tracking. In the next section, where we consider the top-down/bottom-up nature of learned attentional biases, we will cover other methods of assessing attention based on the attentional blink and spatial cueing.

One of the most obvious features of visual attention is that it tends to coincide with where our eyes are looking; indeed, it is well-established that an eye movement to a given location is always preceded by a spatial shift of attention to that location (Deubel & Schneider, 1996). Attention indexed by eye movements is known as *overt attention*. There now exist several studies of the relationship between overt attention and predictiveness. For example, Le Pelley et al. (2011) tracked participants' eye movements as they performed the learned predictiveness task of Le Pelley and McLaren (2003). During training Phase 1, participants spent longer looking at cues that predicted the outcome on each trial than cues that were nonpredictive. Moreover, this bias in eye gaze towards cues that had been predictive in Phase 1 persisted during training Phase 2, in which all cues were now equally predictive of the outcomes with which they were paired. So not only was learning during Phase 2 influenced by participants' prior experience of predictiveness, their overt attention was too. The finding of greater eye gaze to predictive than non-predictive cues has now been replicated many times (Beesley, Nguyen, Pearson & Le Pelley, 2015; Haselgrove et al., 2015; Hoffman & Rehder, 2010; Le Pelley et al., 2011; Le Pelley, Mitchell & Johnson, 2013; Mitchell et al., 2012; Rehder & Hoffman, 2005a; 2005b; for related evidence from electroencephalography, see also Luque, Morís, Rushby & Le Pelley, 2015; Welham & Wills, 2011). This evidence, together with the earlier learned predictiveness findings, firmly establishes that predictive cues are higher than non-predictive cues in both associability and overt attention. But what is the nature of this attentional bias? In the next section we address

perhaps the most important question of this kind; are participants deliberately attending to previously predictive cues, or is attention being captured in an involuntary fashion?

Controlled and Automatic Attentional Processes

The preceding sections have described findings demonstrating that predictive stimuli are learned about more rapidly, and receive more overt attention, than nonpredictive stimuli. These findings are consistent with the predictiveness principle, and strongly implicate attentional processes in the bias towards learning about previously predictive stimuli. But how do these attentional processes operate? Some aspects of the data seem straightforward to explain. Take, for example, the learned predictiveness task of Le Pelley and McLaren (2003). Recall that, in Phase 1, one of the foods presented on each trial was always paired with the same type of allergic reaction (e.g., apple was always paired with nausea), whereas the other food was equally paired with both types of reaction (e.g., banana was sometimes paired with nausea and sometimes with sweating). So to respond correctly, participants need to identify the predictive cue on each trial, while the nonpredictive cue should be ignored because it does not allow correct predictions. Of course, therefore, participants will spend longer looking at the predictive cues in order to make accurate predictions. The question then is, why does the attentional bias persist into Phase 2, when all cues are equally predictive?

One possibility is that the predictiveness effect is a top-down, controlled attentional bias. Participants may reason that "Apples were predictive in Phase 1 – perhaps they are also important in Phase 2". That is, the bias towards previously-predictive cues in Phase 2 is a conscious, top-down decision to look for cues that participants have learned are important in the current context. The alternative is that learning about predictiveness influences attention in a more automatic, bottom-up way – changing the extent to which cues 'grab' attention in a way that is outside of people's strategic control. We shall consider these two options in turn.

Controlled attention. Mitchell et al. (2012) investigated whether top-down attentional processes can influence the learned predictiveness effect. Between Phases 1 and 2, participants in the *Change* group were instructed that the stimuli that had determined the outcome in Phase 1 (the predictive cues, e.g. apples) were *unlikely* to determine the outcome in Phase 2. If the learned predictiveness effect was sensitive to top-down attentional processes, these instructions should have a profound effect on learning in Phase 2: participants might be expected to switch attention from the previously predictive cues to previously nonpredictive cues (such as bananas). In the *Continuity* group, participants were told that the predictive cues in Phase 1 were also likely to determine the outcome in Phase 2. In this group, the standard learned predictiveness effect was expected.

The results suggested that the learned predictiveness effect could indeed be modulated by top-down control. The reversal instructions given to the Change group produced a learned predictiveness effect in the opposite direction to that observed in the Continuity group. During Phase 2, participants in the Change group learned more about cues that had been *non*predictive in Phase 1, and paid greater overt attention to these cues (measured using eye-tracking), compared to those that had been predictive in Phase 1. This finding shows that participants in the Change group were aware of the predictive status of cues in Phase 1 and were able to use that knowledge to strategically allocate their attention to previously non-predictive cues in Phase 2. Furthermore, Mitchell et al. (2012) saw no evidence for any automatic bias towards previously predictive cues; there was no evidence for any overall bias across both groups towards previously predictive cues. It should be noted, however, that subsequent studies using the same approach have found that, although instructions do have a profound influence on the effect, there can also be a residual predictiveness effect (a bias towards previously predictive cues when averaged across both groups) that survives reversal instructions (Don & Livesey, 2015; Shone et al., 2015). This residual bias could reflect a

bottom-up contribution to the learned predictiveness effect that is resistant to the influence of instructions on strategic allocation of attention. However, the validity of this conclusion rests on the assumption that the instructions given to the Change and Continuity groups were equally effective in producing top-down changes in attention, but this may not be the case. In particular, a residual predictiveness effect would be observed if some participants in group Change, for whatever reason, did not make use of the reversal instructions and instead assumed continuity between the two phases.

Taking these data together, then, there is strong evidence that knowledge of predictiveness can influence attention in a controlled, top-down fashion. There is also evidence (albeit equivocal) consistent with the existence of another, more automatic process that is resistant to instructional manipulations. The next section describes further studies providing clearer evidence for an automatic influence of predictiveness on attention (further evidence for automatic processes will be presented in the section on *learned value* below).

Automatic attention. Studies investigating the possibility that predictiveness training can influence automatic attentional capture have used procedures based on the attentional blink effect (Glautier & Shih, 2015; Livesey, Harris & Harris, 2009), and spatial cueing (Haselgrove et al., 2015; Le Pelley, Vadillo & Luque, 2013). We review these two lines of evidence in turn.

Attentional blink. In the attentional blink paradigm, participants must identify two targets in a rapid sequence of visual stimuli appearing in the same location on a screen. The effect is called the attentional blink because identification of the second of the two targets is often impaired if it appears 200-500ms after presentation of the first target. This effect is widely assumed to reflect attentional limitations during encoding that restrict perceptual awareness of targets (Raymond, Shapiro & Arnell, 1992). Consistent with this interpretation, the attentional blink is reduced for targets that are higher in perceptual salience (e.g., letters

among numbers, Chun & Potter, 1995), emotional salience (e.g., negatively valenced words, Anderson & Phelps, 2001), or semantic salience (e.g., one's own name, Shapiro, Caldwell & Sorensen, 1997), and which are hence more likely to capture attention.

Livesey et al. (2009: see Glautier & Shih, 2015, for a replication) demonstrated that the susceptibility of stimuli to the attentional blink varies as a function of predictiveness. On each trial participants saw a rapidly presented series of letters, and then had to make either a left or right button-press response as quickly as possible (Figure 1). The appearance of certain target letters (predictive cues) in the series predicted which response (left or right) participants would be required to make. Other target letters (nonpredictive cues) did not predict the correct response. In a separate, subsequent test phase, participants showed an advantage in detecting previously-predictive target letters in the rapidly-presented letter streams, relative to previously-nonpredictive targets. Detection of the predictive cues was, therefore, less impaired by the attentional blink than was detection of the nonpredictive cues.

This finding is certainly consistent with the idea that associative learning effectively increases the salience of (and extent of attentional capture by) predictive stimuli, thereby increasing their detectability. Importantly, the participant's task was to detect all targets presented during the test phase. Thus, there was no obvious advantage to the participant in biasing attention towards one stimulus and away from another; the implication is that the bias observed was automatic and not under the participant's control.

Spatial cueing. People are faster to respond to events that occur in locations to which they are attending. This idea is the basis of the spatial cueing task (Posner, 1980; Posner, Nissen & Ogden, 1978), developed to measure shifts of visual attention. Recent experiments have used a variant of this task to investigate the possibility of an automatic attentional bias towards the location in which predictive stimuli appear (Le Pelley, Vadillo, et al., 2013). On each trial of an initial categorization phase, two stimuli appeared: a green square (which

could take one of two different shades of green) and a set of oblique lines (which could take one of two different orientations). Participants categorized this stimulus pair into one of two categories, with immediate corrective feedback provided on each trial. For half of the participants, the shade of green predicted category membership and the orientation of the lines was nonpredictive. For the remaining participants, the reverse was true.

A subsequent test phase used a spatial cueing procedure to assess whether the categorization task produced an attentional bias towards predictive cues and away from nonpredictive cues. On each trial a pair of stimuli (a green square and set of oblique lines), appeared on opposite sides of the screen for 150 ms. These stimuli then disappeared, and after 200 ms a triangle (the probe) could appear in the location formerly occupied by one of them. Thus the stimulus-onset asynchrony (SOA) between the cueing stimuli and the probe was 350 ms. This probe was equally likely to appear in either location. Participants' task was to respond to the appearance of the probe as rapidly as possible by pressing a single key.

Responses to the probe were faster when it appeared in the location previously occupied by the predictive cue than when it appeared in the location of the nonpredictive cue (a pattern recently replicated by Haselgrove et al., 2015). This pattern again indicates an attentional advantage for the predictive cue, consistent with the predictiveness principle. It suggests that the predictive stimulus captured participants' spatial attention and hence sped responses to events occurring in that location; in this case, the onset of the probe. This attentional capture occurred even though (1) it was not required by the task, (2) it was not adaptive with regard to that task, and (3) the short SOA meant that there was little time for participants to consciously process and respond to the stimuli on each test trial.

Le Pelley, Vadillo, et al. (2013) also demonstrated that providing more time for participants to consciously process the stimuli—by increasing the SOA on test trials to 1000 ms—significantly *weakened* the influence of predictiveness on probe responding. This

finding suggests that the pattern observed at short SOA is not a result of conscious, controlled processing but instead reflects a rapid and automatic effect of predictiveness on attentional capture. In contrast, a long SOA provides sufficient time for participants to use controlled processes to correct for the automatic attentional orienting caused by presentation of the stimuli, returning attention to the centre of the display (cf. Klauer, Roßnagel & Musch, 1997). These findings represent strong evidence for a non-volitional component in the learned attentional bias towards previously predictive cues.

The studies reviewed in this section suggest that learning about predictiveness can influence the extent to which cues will subsequently command attention. This is true even when attending to those cues is not required by the task that participants are performing (as in the cueing task of Le Pelley, Vadillo, et al., 2013), or when there is no advantage to be gained by biasing attention towards predictive cues (as in the attentional blink task of Livesey et al., 2009). The implication is that the influence of predictiveness on attention may, at least in part, reflect a change in the likelihood that cues will automatically capture attention, regardless of a person's goals or intentions. We shall see later that this evidence for an automatic effect of predictiveness on attention is complemented by a larger body of evidence demonstrating that learning about *value* can also influence automatic attentional capture.

Summary: The Predictiveness Principle Affects Top-Down and Bottom-Up Attention

In this section on the predictiveness principle, we have seen that predictiveness exerts an influence on many different measures of stimulus selection, not merely learning rate – these include eye gaze (overt attention), spatial cueing, and susceptibility to the attentional blink. So models in which the effect of predictiveness is limited to an influence on associability (e.g., Mackintosh, 1975) do not go far enough. Instead it seems that learning about predictiveness exerts a more general effect on the psychological salience of stimuli that influences many aspects of their processing; not only how they are learned about, but how

they perceived and demand responses (in terms of orienting) in their own right.

More generally, the data reviewed here support the idea that attention can be learned just as other behavioural responses are learned. These studies indicate that the attentional bias towards previously predictive stimuli is multiply determined. In part it seems likely to be the consequence of deliberate and controlled (top-down) attention toward cues deemed important as a consequence of their history. However, there is also evidence that predictiveness produces an attentional bias that is automatic and outside of participants' control.

In the next section we consider a quite different model of conditioned attention, also developed in the context of animal learning, proposed by Pearce & Hall (1980). Here we ask whether people, at least under some circumstances, pay particular attention to cues that fail to predict important outcomes – about which they are uncertain.

Uncertainty and Attention

The Pearce–Hall Model

In the previous section, we described studies in which better predictors of outcomes demand more attention. However, Dickinson (1980) has argued that it actually makes little sense to continue to devote the majority of learning resources to those cues whose consequences are already well-known. Rather, the most efficient approach is for attentional resources to be focussed on cues whose predictive status is currently unknown, thus promoting rapid learning about the true significance of those cues. This notion was formalized in an alternative form of attentional learning theory suggested by Pearce and Hall (1980). At the heart of this theory is the suggestion that attention will be greatest (and hence learning will be most rapid) for stimuli whose consequences are poorly predicted and hence surprising. Following Griffiths, Johnson and Mitchell (2011), we term this the *uncertainty principle*, since it anticipates greater attention to cues whose predictive status is uncertain. On the surface at least, this

suggestion is at odds with the predictiveness principle outlined earlier, which holds that attention will be greatest for stimuli that predict their consequences most accurately.

In rats, there is evidence in support of both the uncertainty principle and the predictiveness principle (for reviews, see Le Pelley, 2004; Pearce & Mackintosh, 2010). This pattern of findings has motivated the development of *hybrid* theories of attentional learning, which attempt (in various ways) to reconcile and combine both the predictiveness and uncertainty principles (George & Pearce, 2012; Le Pelley, 2004, 2010; Pearce, George & Redhead, 1998; Pearce & Mackintosh, 2010). We have already seen ample support in human learning for the predictiveness principle. If hybrid models are to have any relevance to human learning and attention, there must also be strong evidence for the Pearce-Hall model's uncertainty principle in humans.

Evidence for the Uncertainty Principle in Humans

There are three lines of research that have attempted to test the uncertainly principle in humans. Firstly, over the past 15 years, neuroscientists have shown a great deal of interest in the notion of uncertainty. We review attempts to locate the brain areas in which uncertainty is processed. Our conclusion is that current neuroscientific evidence for the uncertainty principle—specifically that attention goes to *cues* about which people are uncertain—is not at all strong. Secondly, a model-fitting approach has been taken. Some studies have found that models in which attention to cues is variable, and determined by uncertainty, fit empirical learning data better than models in which attention is fixed. However, findings are mixed and not all possible models (e.g., predictiveness-based models) have been tested. The evidence for the uncertainty principle from model-fitting studies is again, therefore, not compelling. Finally, some experimental tests of the qualitative predictions of the Pearce-Hall model have shown evidence for the uncertainty principle. Unfortunately, these findings have not proven to be robust to replication. Hence, overall, our conclusion is that the uncertainty principle

does not currently receive strong empirical support from studies of human learning. However, this lack of support may reflect the possibility that existing experiments are not always optimally designed to detect an influence of uncertainty on attention in humans. Towards the end of this section we consider ways in which future studies might be parameterized so as to be most sensitive to detecting an influence of uncertainty on 'attentional exploration'.

Encoding of unsigned prediction error in the brain. Imagine pairing a cue, A (which might be a light, or a tone, etc) with an outcome (which might be food, money, shock etc). *Prediction error* describes the discrepancy between the magnitude of the outcome that actually occurs (usually represented as λ) and the magnitude of the outcome that is predicted by cue A (usually represented as V_A). Hence the prediction error is given by $(\lambda - V_A)$. If the outcome is under-predicted (so $\lambda > V_A$), prediction error will be positive; if the outcome is over-predicted ($\lambda < V_A$), prediction error will be negative. More generally, the absolute value of the prediction error, $|\lambda - V_A|$, describes the accuracy with which cue A predicts the outcome. If the outcome on a given trial is well-predicted by cue A, the absolute prediction error will be small; if the outcome is poorly predicted by A and hence surprising (either over- or under-predicted) the absolute prediction error will be large. Hence this *unsigned* prediction error (uPE) provides an index of the predictive uncertainty associated with cue A. There are (at least) three critical aspects that uPEs must fulfil if the uncertainty principle is to be supported: (1) uPEs produced by outcomes must be encoded somewhere in the brain; (2) these uPEs must then be 'attached' to cues; and (3) this must result in a change in attention to these cues. We consider these aspects below.

Several fMRI studies with humans have found brain regions whose activity correlates across trials with uPE, with greater activity generated by presentation of surprising outcomes than predicted outcomes. That said, the exact region implicated varies between studies: either amygdala (Boll, Gamer, Gluth, Finsterbusch & Buchel, 2013; but see Prevost, McCabe,

Jessup, Bossaerts & O'Doherty, 2011, for a null finding), hippocampus (Ploghaus et al., 2000; Vanni-Mercier, Mauguiere, Isnard & Dreher, 2009), substantia nigra (D'Ardenne, Lohrenz, Bartley & Montague, 2013; Dreher, Kohn & Berman, 2006), cingulate cortex (Bach, Hulme, Penny & Dolan, 2011; Behrens, Woolrich, Walton & Rushworth, 2007; Hauser et al., 2014), prefrontal cortex (Turner et al., 2004), temporoparietal junction (Eickhoff, Pomjanski, Jakobs, Zilles & Langner, 2011), superior frontal gyrus (Ploghaus et al., 2000), or cerebellum (Ploghaus et al., 2000). Other studies using electroencephalography (EEG) have shown that that a distinct component of electrophysiological activity elicited by feedback, known as the feedback-related negativity (FRN), may encode uPE; the FRN is thought to have its source in the anterior cingulate cortex (Huang & Yu, 2014; Mas-Herrero & Marco-Pallares, 2014). The reason for the difference in highlighted brain regions across these studies is not clear, though it most likely relates to the wide variety of different procedures. Some of these studies used 'standard' associative learning tasks in which people were presented with cues and predicted outcomes (e.g., Boll et al., 2013; Turner et al., 2004), while others used gambling tasks (e.g., Behrens et al., 2007; Prevost et al., 2011), and others used financial forecasting tasks (e.g., D'Ardenne et al., 2013). In some cases the degree of uncertainty varied over time (Boll et al., 2013; Hauser et al., 2014), while in others uncertainty was sustained (e.g., Behrens et al., 2007; Vanni-Mercier et al., 2009). Some studies used rewarding outcomes (typically money, e.g. Behrens et al., 2007; D'Ardenne et al., 2013; Hauser et al., 2014), others used aversive outcomes (pain, e.g., Boll et al., 2013; Ploghaus et al., 2000), and others used neutral outcomes (Turner et al., 2004). Future work should aim to clarify which aspects are critical for determining where uPEs are encoded.

To reiterate, the studies cited above reported brain regions whose activity correlated with the uPE generated by presentation of an *outcome*. Such studies show that uPE is encoded in the brain, which is a prerequisite of the uncertainty principle. But the uncertainty

principle relates to the effect of uPE on processing of the *cue*, not the outcome: that is, cues previously associated with uncertainty (i.e., with large uPEs) should command more processing resources than cues previously associated with certainty (small uPEs). While this issue has received less attention, some fMRI studies have shown that amygdala activity elicited by presentation of a cue correlates with the size of uPEs previously experienced in relation to that cue (Aron et al., 2004; Boll et al., 2013; Li, Schiller, Schoenbaum, Phelps & Daw, 2011; see also Preuschoff, Hart & Einhauser, 2011).

While these latter imaging data show a link between cues and uPEs, they still fall short of confirming the uncertainty principle. This is because the imaging data tell us little about the *psychological* effect of uncertainty on cue processing. These data merely reveal that cues associated with uncertainty elicit greater activity in certain brain areas; perhaps this activity has nothing to do with attentional processing. For example, uncertain cues will have previously been associated with more errors in outcome prediction than certain cues; perhaps the uPE-related activity elicited by such cues relates to frustration, rather than a difference in attentional processing.

In summary, while studies of encoding of uPEs in the brain are revealing in terms of functional brain anatomy, they do not convincingly address the uncertainty principle. The next section reviews an approach that has the potential to provide somewhat stronger support.

Uncertainty and cue processing I: Model-fitting studies. The uncertainty principle states that cues previously associated with greater uncertainty will receive greater attention. Following a theme of this article, studies testing this claim have measured the effects of uncertainty on both the rate of learning about cues (i.e, their associability) and the extent of orienting to them.

One strand of research has used a model-fitting approach. For example, participants in Boll et al.'s (2013) study experienced one of three different visual cues on each trial. During

an initial acquisition phase, one of the cues (CS100) was paired with delivery of electric shock every time it appeared, another (CS-) was never paired with shock, and the third (CS50) was paired with shock on a random half of its presentations. Hence CS50 was associated with greater predictive uncertainty than CS100 or CS-. In a subsequent phase, these relationships were reversed so that now CS100 was never paired with shock, and CS- was always paired with shock (CS50 continued to be paired with shock on half of trials). The onset of this reversal phase (which occurred without warning) was thus a period of increased uncertainty, when outcomes were particularly surprising. On every trial, participants rated their expectancy of shock immediately after presentation of the cue. Changes in these shock-expectancy ratings were used to track participants' learning of the cue–outcome relationships. Boll et al. tested how well two different models of learning fit the shock-expectancy data. In the first model, based on the Pearce–Hall theory, the associability of each cue was determined by its previously experienced predictive uncertainty (i.e., by the size of the uPEs on previous trials with this cue). This model implements the uncertainty principle. In the second model, the associability of each cue was fixed. Critically, Boll et al. found that the Pearce–Hall type model provided a significantly better fit to their learning data than the fixed learning-rate model. Similar results have been reported elsewhere (Bai, Katahira & Ohira, 2014; Behrens et al., 2007; Cavanagh, 2015; Li et al., 2011). However, other studies have found evidence favouring a fixed rate model over a Pearce–Hall type model (Payzan-LeNestour, Dunne, Bossaerts & O'Doherty, 2013), or have found a negative correlation between uPEs and associability (D'Ardenne et al., 2013). Finally, Prevost et al. (2011) found evidence favouring a fixed rate model for learning which of two responses to make under conditions of uncertainty, but a Pearce–Hall model for learning whether a cue indicated the current trial would produce an appetitive, aversive or neutral outcome.

We can see this as a quantitative approach. In these studies the fixed-rate models

generally provide a reasonable account of the learning data, but the variable-rate, Pearce–Hall type models explain more of the variance. This is unsurprising. The fixed-rate model is a special case of a variable-rate model in which the parameter controlling changes in rate is set to zero, so the variable-rate model will always fit at least as well as the fixed-rate model. So in the end, whether a fixed- or variable-rate model is preferred comes down to how heavily the greater flexibility of the variable model is penalized; there is debate over the most appropriate way to implement this penalty (Kenny, 2015; Pitt & Myung, 2002). It is also possible that the greater variance explained by the more flexible model may reflect *overfitting* – fitting noise in addition to systematic patterns (Pitt & Myung, 2002; Wills & Pothos, 2012). Perhaps more importantly, the model-fitting studies cited above consider only two options: a fixed-rate model versus a variable-rate model incorporating the uncertainty principle. They do not consider alternative versions of a variable-rate model; in particular, a version incorporating the predictiveness principle – perhaps such a model would provide the best fit of all. In other words, these studies fail to consider the *relative adequacy* (Wills & Pothos, 2012) of predictiveness- and uncertainty-based models. Hence the specific support offered by these model-fitting studies to the uncertainty principle remains rather weak.

Uncertainty and cue processing II: Qualitative studies. A more satisfying approach would come from a focussed demonstration of a *qualitative* pattern of behaviour that is anticipated by the uncertainty principle, and which cannot be reconciled with other views (such as the predictiveness principle, or a fixed-rate approach). Below we consider the handful of human studies that have taken this approach, with somewhat mixed results.

Griffiths et al. (2011) found some evidence that surprising outcomes increase attention to the cues that preceded them. This study used an allergist task in which participants learned about the strength of allergic reactions (weak, moderate or strong) produced when Mr X ate different foods. In an initial phase of training, the target cue was paired with an outcome of

moderate strength. In the second phase, the strength of the outcome paired with the target increased. According to Pearce-Hall, the initial training, in which the target cue is established as predictive of the moderate outcome, should decrease attention to that target and hence slow learning in the second phase (as compared to a control condition in which the target cue did not appear in the initial training phase). This effect, known as *negative transfer* (Hall & Pearce, 1979), was observed in Experiment 1. However, negative transfer is not itself a unique prediction of the Pearce-Hall model. To further test the Pearce-Hall account, half the participants in a second experiment were given some surprising trials inserted between the two phases (the target was presented, but no outcome occurred on these trials). The surprise should increase the uncertainty regarding the predictiveness of the target cue. According to the uncertainty principle this should produce an increase in attention, and hence increase the speed with which the new relationship between the target and the strong outcome is learned in Phase 2. This is exactly what was observed. Griffiths et al.'s demonstration of release from negative transfer lends unique support to the Pearce-Hall model.

It should be noted, however, that unpublished attempts to replicate the negative transfer effect of Griffiths et al. (2011), by the original researchers, have been mixed. There is no doubt that training with a moderate outcome slows later learning with a larger outcome. In contrast, the critical effect of the surprising trials on later learning is typically weak, and not always replicated. Moreover, other studies using related procedures have found no evidence for a restoration of attention to a cue paired with a surprising outcome (i.e., no release from negative transfer: Packer, Siddle & Tipp, 1989; Siddle, 1985; Siddle, Booth & Packer, 1987; see also Lipp, Siddle & Dall 1993; Schutzwohl, 1998).

A somewhat similar mixed pattern emerges when we examine another source of evidence for the uncertainty principle. In Hogarth, Dickinson, Austin, Brown, and Duka's (2008) design, a compound of cues A and X was always followed by an outcome (AX+: in

one experiment the outcome was a 50 dB tone; in another it was a more aversive, 97 dB noise), C and X were never followed by the outcome (CX–), and B and X were followed by the outcome 50% of the time (BX+/–). Thus A and C were consistent predictors of outcome and no outcome respectively, while B was a poor predictor. Differences in attention to these cues were assessed by comparing eye gaze to the unique cue on each trial (A, B, or C) with gaze to the common cue (X). In two experiments, Hogarth et al. found that participants spent longer looking at the unique cue than the common cue on all three types of compound trial. However, this bias in gaze was greater on BX trials than on AX and CX trials. The implication is that overt attention was greater to less predictive cues than to more predictive cues (see Trick, Hogarth & Duka, 2011, for a similar, but not identical, result).

This finding led Hogarth, Dickinson, Austin, et al. (2008) to suggest that 'attention for learning' operates according to the uncertainty principle (though an effect on learning rate was not reported). Again, however, there is inconsistency in the findings. Austin and Duka (2010) replicated Hogarth et al.'s design, with only minor differences in procedure, using either a 97 dB noise (as used by Hogarth et al.) or a 102 dB noise as the outcome. Regardless of which noise was used, this study showed the greatest attentional bias to cue A (predictive of outcome), followed by B (nonpredictive), and then cue C (predictive of no outcome). This result is quite different from the B > A = C pattern observed by Hogarth et al. Austin and Duka did report eye gaze data consistent with the Pearce–Hall model, but only when positively-valenced outcomes (monetary gains) were used. However, in a further study using positively-valenced outcomes (cigarettes; participants were smokers), Austin and Duka (2012) again found the A > B > C pattern that conflicts with Pearce–Hall. Until further replication establishes a consistent set of findings, it is unclear how these data should be treated.

The uncertainty principle: conclusions. Overall, the evidence from human studies in support of the uncertainty principle is not as strong as that in support of the predictiveness

principle. While there are findings that are consistent with the uncertainty principle, they do not always provide unequivocal support, and must be balanced against null findings using similar procedures. It is, therefore, perhaps premature to apply to humans the kinds of hybrid models of attentional learning developed in the context of animal learning. However, the relative lack of evidence for the Pearce–Hall process may simply reflect a trend in the types of procedures that have been used to study attentional learning in humans. Beesley et al. (2015) have noted that the predictiveness principle can be seen as a pattern of *attentional exploitation*, in which cognitive resources become focused on the most predictive sources of information available so as to best exploit these sources. In contrast, the uncertainty principle embodies *attentional exploration*, in which cognitive resources are deployed to explore currently uncertain but potentially useful sources of information for predicting future events. In studies that do not involve substantial monetary rewards (the vast majority of studies of predictiveness in humans), it seems fair to assume that participants' primary motivation is to complete the experiment as quickly as possible and with minimum effort. It is plausible that these are circumstances favouring exploitation of whatever useful information they have available, rather than an effortful (and possibly fruitless) search for new information. Perhaps if participants were more strongly encouraged and motivated to explore all potential sources of information, then evidence for a Pearce–Hall type process might emerge: people might allocate attention in a way that helps them to understand the causal status of the cues whose outcomes are currently uncertain.

What might encourage participants to explore and therefore reveal evidence for the uncertainty principle? One factor that may be important is the complexity of the experimental design. Exploration of all sources of information would seem more likely in simple designs that involve only a small number of cues and relationships, where the memory demands are relatively low and it might be possible to establish the causal status of all of the cues. In

effect, under such simple conditions, participants have the capacity to explore for further information (and may have more motivation to do so; perhaps when confronted with a very simple design, participants reason that there must be more to the experiment and hence undertake a deeper search). In this regard it is notable that studies finding evidence of attentional exploration have indeed tended to use much simpler designs than studies finding evidence of attentional exploitation. For example, participants in Griffiths et al.'s (2011) experiments (which found exploration) experienced at most three different cues in any phase of training, while Le Pelley and McLaren's (2003) study (which found exploitation) involved eight different cues which were experienced in overlapping cue compounds.

A second property that might influence the exploration/exploitation balance is the overall amount of uncertainty present in the experiment design. In the majority of studies that have found evidence of exploitation, the outcome that occurred on every trial was perfectly predicted by the cues that were presented. So once participants had learned the various cue–outcome relationships, they could make a correct prediction on every trial, and hence these designs have zero uncertainty (e.g., Bonardi et al., 2005; Eimas, 1966; Kruschke, 1996; Le Pelley & McLaren, 2003; Livesey & McLaren, 2007; Lochmann & Wills, 2003; Whitney & White, 1993). Under these circumstances, when perfect performance is possible, there seems little value in exploring the cues for new sources of information – instead it makes sense to exploit the predictive relationships that are already known. In contrast, in the study by Hogarth et al. (2008), the outcome on BX trials was unpredictable: the noise occurred on a random half of BX trials. Similarly, in Griffiths et al.'s (2011) negative transfer study, the size of the outcome paired with the target cue changed in an unpredictable way over the course of training. Under these circumstances of uncertainty, it makes sense to explore the cues for additional pieces of information that may allow more accurate predictions to be made in future. For example, perhaps participants in Hogarth et al.'s study thought that there

was some subtle difference in the B cue that was presented on each BX trial, that predicted whether or not the outcome will occur, and hence attended to this cue in order to try and identify this putative predictive feature. Of course in reality there was no such predictive information on BX trials, but participants were not aware of this and so may have continued searching under the belief that more accurate performance was possible.

Evidence for the role of overall uncertainty in determining the exploration/exploitation balance comes from a recent study by Beesley et al. (2015). On each trial, two cues were presented as a compound; one of these cues predicted the outcome that would occur, while the other did not. Eye-tracking showed that participants paid more overt attention to the predictive cue in each compound, revealing evidence for the predictiveness principle. However, the absolute accuracy with which the predictive cues allowed the outcome to be predicted was systematically varied (between-subjects in one experiment and within-subjects in another). In the 'certain' condition, the predictive cues allowed the outcome to be predicted with 100% accuracy; in the 'uncertain' condition, the predictive cues allowed the outcome to be predicted with only 67% accuracy. Notably, overt attention to all cues (both predictive and nonpredictive) was greater in the uncertain condition than the certain condition; that is, participants spent a greater proportion of each trial looking at the cues in the uncertain condition. This finding of greater attention to cues under conditions of greater uncertainty accords with the uncertainty principle embodied by the Pearce–Hall model, and with the suggestion that uncertainty drives attentional exploration of cues for prediction-relevant information.

Taken together, Beesley et al.'s (2015) findings suggest that uncertainty might influence attention based on the overall uncertainty of the prediction made by *all* presented cues (i.e., given this set of cues, how uncertain am I of the outcome that will occur?), while predictiveness operates at the level of the individual cues (which of the N cues that were

presented was individually the best predictor of the outcome?). This suggestion is consistent with the original characterizations of uncertainty-driven and predictiveness-driven processes in the models of Pearce & Hall (1980) and Mackintosh (1975) respectively, and offers one way in which the predictiveness and uncertainty principles—which at first glance seem to be mutually contradictory—might be reconciled (see Le Pelley, 2004, 2010, for a hybrid model along these lines). However, while Beesley et al.'s findings are promising, research into the experimental parameters and procedures that might reveal evidence for the uncertainty principle is at an early stage; as such, the possibilities raised above are currently largely speculation. Indeed, the notion of overall uncertainty cannot be the whole story. As far as we are aware, it is true that all studies supporting the uncertainty principle have used procedures in which the outcome cannot be predicted perfectly on each trial. However, there are also a handful of studies that have used procedures in which the outcome cannot be predicted perfectly, but which have yielded evidence consistent with the predictiveness principle (and counter to the uncertainty principle: Kattner, 2015; Le Pelley, Turnbull, et al., 2010; Livesey et al., 2011). Future research should aim to establish systematically what, if any, are the factors that determine whether a pattern of attentional exploitation (following the predictiveness principle) or exploration (following the uncertainty principle) will be observed.

Summary: Predictiveness versus uncertainty

The preceding sections describe a body of evidence showing that the attention paid to a stimulus is influenced by the accuracy with which it predicts subsequent events. We have seen that there is strong evidence supporting the operation of the predictiveness principle in humans – the idea that more attention will be paid to cues that are more accurate predictors of subsequent events. There is also some evidence (certainly less abundant and arguably less compelling) in support of the uncertainty principle. This is (superficially at least) the opposite

of the predictiveness principle, and describes a process whereby greater attention is allocated to cues paired with outcomes that are less accurately predicted (i.e., more surprising). Hybrid models incorporating both of these principles have been put forward in the context of animal learning theory, though it remains for future empirical work to establish more convincingly whether both of these mechanisms operate in humans.

Before moving on, we note briefly that there are certain phenomena demonstrating a relationship between learning and attention that are anticipated by both the predictiveness principle and the uncertainty principle. A prime example is *blocking*. In blocking, pairings of cue A with an outcome are followed by pairings of an AB compound with the same outcome. Less is learned about the contingency between cue B and the outcome than for a control cue D, trained in a CD compound where C was not previously paired with the outcome. Thus learning about B on AB→outcome trials is blocked by prior learning of the A→outcome relationship. Several studies with humans show that blocking treatment results in a reduction to attention to the blocked cue B relative to the control cue D. These include studies measuring attention in terms of associability (Beesley & Le Pelley, 2011; Griffiths & Le Pelley, 2009; Kruschke & Blair, 2000; Le Pelley, Beesley & Suret, 2007; Le Pelley, Oakeshott & McLaren, 2005; Liu & Luhmann, 2013; Mitchell, Harris, Westbrook & Griffiths, 2008), eye gaze (Beesley & Le Pelley, 2011; Eippert, Gamer & Buchel, 2012; Kruschke, Kappenman & Hetrick, 2005; Le Pelley, Beesley & Griffiths, 2014; Wills, Lavric, Croft & Hodgson, 2007), recognition memory (Griffiths & Mitchell, 2008), and 'neural markers' of selective attention recorded using electroencephalography (EEG: Wills et al., 2007; but see also Kopp & Wolff, 2000, for a null result). The pattern of reduced attention to B emerges from the predictiveness principle because B is a poorer predictor of the outcome on AB→outcome trials than is A, as a result of prior training with A alone. The same pattern emerges from the uncertainty principle because prior training with A renders the outcome on

AB→outcome trials unsurprising, compared to the outcome on CD→outcome trials which is surprising because neither cue has been pretrained. That is, the overall uncertainty associated with the AB compound is lower than that associated with the CD compound, and so—on the basis of the idea (raised earlier) that overall uncertainty may be the crucial determinant of attention (Pearce & Hall, 1980)—the uncertainty principle anticipates a reduction in attention to cues belonging to a low-uncertainty compound (A and B) relative to cues in a high-uncertainty compound (C and D).[3]

Other studies that are consistent with both the predictiveness and uncertainty principles have examined how attention to contexts varies as a function of the information learned in them (Abad, Ramos-Alvarez & Rosas, 2009; Callejas-Aguilera & Rosas, 2010; Gawronski, Rydell, Vervliet & De Houwer, 2010; Leon, Abad & Rosas, 2011; Nelson, Lamoureux & Leon, 2013; Rosas & Callejas-Aguilera, 2006). In essence, in these studies people initially learn that whenever Mr X eats a particular food (say steak) in restaurant R he suffers illness, but other foods eaten in this restaurant do not cause illness. Later, this relationship is extinguished: steak eaten in restaurant R no longer causes illness. Evidence suggests that this treatment results in an increase in attention to the context (restaurant R: see Nelson et al., 2013; Rosas & Callejas-Aguilera, 2006: see also Dunsmoor, Campese, Ceceli, LeDoux & Phelps, 2015; Lucke, Lachnit, Stuttgen & Uengoer, 2014). This finding follows from the predictiveness principle because, during the extinction phase, the restaurant is a better predictor of 'no illness' than is steak (since steak has a strong pre-trained association with illness). It also follows from the uncertainty principle because, on early extinction trials, the occurrence of 'no illness' after eating steak at restaurant R is surprising and hence (according to the Pearce–Hall model) will result in an increase in attention to all presented stimuli, which includes the restaurant.

To reiterate, these findings, like those reviewed in previous sections, provide further

evidence of a relationship between attention and learning about predictiveness. But because they do not provide selective support for any particular account of this relationship we will not consider them further here.

Learned Value

In the preceding discussion, we have considered the effect on attention of learning about differences in the *predictiveness* of cues; that is, the consistency or variability of the outcome that follows a particular cue. A predictive cue is one that is consistently followed by the same outcome, while a nonpredictive cue is one that is followed (unpredictably) by different outcomes. However, an emerging body of research suggests this is not the whole story; that attention to cues is also influenced by the motivational value of the outcome events they are paired with – what we term the *learned value* of the stimuli. That is, while attention tends to increase towards cues that reliably signal their outcomes, even greater attention will be paid to cues that reliably signal the outcomes that have the highest value.

To date, the majority of studies investigating the effect of learned value on attention have manipulated the monetary reward associated with cues (though later we shall see some alternative examples). In essence these studies all take a similar approach. Suppose cue X is consistently paired with a relatively large reward (say 10¢), and cue Y is consistently paired with a smaller reward (1¢). Under these circumstances, X and Y have equal *predictiveness*, since both are perfect predictors of the outcome with which they are paired. The critical question is whether learning about these cue–outcome relationships leads to greater attention being paid to the *high-reward cue* X than to the *low-reward cue* Y.

Three key empirical papers have provided a great deal of the impetus behind the recent body of research investigating the influence of learned value on attention: Anderson, Laurent and Yantis (2011b), Della Libera and Chelazzi (2009), and Raymond and O'Brien (2009). In

searching for literature relating to this field, we searched ISI Web of Science on 19/10/2015 for all articles in the fields of psychology, neuroscience, psychiatry, behavioral sciences and multidisciplinary sciences citing one or more of these key papers. The resulting 195 hits were reviewed for content and for references to other relevant articles (4 found). This yielded 86 articles relevant to the current discussion (marked as results of Search 2 in the reference list), which we cover in the sections to follow. Reasons for exclusion were: (1) The article did not present novel empirical data (articles in this category were reviews or modelling papers) – 14 articles; (2) The article described data from nonhuman animals only, whereas our focus is on humans – 9 articles; (3) The article presented data from humans, but did not bear on the question of how learning changes attention to cues – 86 articles. Most articles in this category made only a passing reference to one of the key papers as providing an example of a factor (reward) that influences attention (47 articles). Others described studies that did not involve learning a relationship between a stimulus and a value-related outcome (reward or punishment: 24 articles), or did not measure attention (7 articles). Finally 8 articles described studies measuring learning of an arbitrary attentional response: e.g., when cue X appears in the centre of the screen, participants are rewarded for making a saccade to the left. These studies are irrelevant for current purposes, as we are interested in how reward influences the attention that is paid to the cue itself.

The evidence reviewed below comes in two kinds. In some studies, the cues that predict the valued rewards are task-relevant – participants must do something with them in order to complete a task. In other studies, the cues that predict valued rewards are task-irrelevant: participants are never required to attend to them—indeed, must ignore them—in order to complete the task. For example, when looking through a drawer for my keys, the keys are the task-relevant stimulus. Other items in the drawer (balls of string, batteries etc) that are not the object of my search—and which may even impede my search—are task-irrelevant stimuli.

There is strong evidence for a top-down, deliberate, value-based attentional bias when the cues in question are task-relevant. We review this evidence first. Following this, we describe evidence from studies using task-irrelevant cues, which shows that learned value can also result in bottom-up, automatic attentional capture.

Reward Influences Attention to Task-Relevant Stimuli

It is now well established that the learned value of stimuli influences performance on tasks in which they are task-relevant. For example, in a study of visual search, Stormer, Eppinger and Li (2014) had participants locate, and report the orientation of, a T shape (the target) among rotated L shapes (distractors) on each trial. They received a reward for each correct response, with the size of this reward depending on the colour in which the target T was presented. For example, a correct response to a red T may yield a relatively large reward (5¢), but a green T may yield smaller reward (1¢). Each of the distractors was rendered in a distinct, randomly chosen colour (but not one of the colours in which the target could appear; red or green in this example). Over the course of training, participants became faster to respond to targets rendered in the high-reward colour than the low-reward colour. This finding suggests that learning about the rewards associated with stimuli modulated attention to those stimuli; greater attention allocated to the high-reward target colour than the low-reward colour.

A crucial question, addressed earlier in the context of learned predictiveness, is whether this learned value effect is under participants' control (top-down) or is automatic (bottom-up). A top-down account of Stormer et al.'s findings would argue that, having learned that locating a red target was worth more than locating a green target, people strategically reallocated attentional resources to prioritize processing of red stimuli (cf. Duncan, 1989; Wolfe, Butcher, Lee & Hyle, 2003). A bottom-up account supposes that training changes the extent to which those colours automatically capture attention regardless of the participant's intentions. In this case, reward learning makes the high-reward colour more psychologically

salient, so that it 'leaps out' from the display and is located more quickly. While Stormer et al.'s findings demonstrate an influence of learned value on attention, they do not allow us to decide between these alternative accounts of that influence. This is because the stimuli that predicted reward (red and green colours) were task-relevant throughout the task – they defined the targets to which participants responded in order to earn that reward. Consequently participants' goals (locate valuable stimuli) were aligned with the reward structure of the task, such that top-down and bottom-up accounts anticipate the same pattern of behaviour.

Several other recent studies have found similar evidence consistent with an effect of learned value on attention to task-relevant, reward-related stimuli.[4] All of these examples could potentially be explained by a top-down influence of learned value on attention, a bottom-up influence, or a combination of the two. In some cases it would appear that top-down processes are contributing, since attention is influenced by verbal instructions regarding reward relationships even in the absence of further stimulus–reward training that would be required to alter bottom-up attention (Le Pelley, Mitchell, et al., 2013; Stankevich & Geng, 2015). In other cases it seems that bottom-up processes may be at work. For example, prior association with large reward facilitates visual search for stimuli even if these stimuli are colour singletons (e.g., a single red target circle among several blue distractors: Kiss et al., 2009; Kristjansson et al., 2010). It is often assumed that such singleton search primarily reflects the operation of rapid, bottom-up attentional processes due to the physical salience of the target, so it seems plausible that reward might be modulating the operation of this bottom-up process (see also O'Brien & Raymond, 2012; Stankevich & Geng, 2015, for further evidence consistent with a rapid influence of learned value).

So, we have evidence consistent with an influence of learned value on top-down attention. This is unsurprising, and should be uncontentious. If I have learned that responding to stimuli appearing in location X produces a large reward, while responding to stimuli in

location Y produces small reward, then clearly I ought to prioritise location X rather than location Y so as to maximize my payoff (see Chelazzi, Perlato, Santandrea & Della Libera, 2013; Maunsell, 2004). More exciting is the possibility that learned value might exert an automatic effect on attentional capture that is outside of participants' strategic control. While the studies reviewed above are (arguably) suggestive of such an effect, work using task-irrelevant cues has yielded more conclusive evidence.

Value-Modulated Attentional Capture by Task-Irrelevant Stimuli

As noted above, the reason the studies in the previous section have trouble discriminating between goal-directed and automatic accounts is because people's goals were aligned with the task's reward structure. Studies using task-irrelevant cues have solved this problem by effectively decoupling participants' goals from the reward structure, so that the relationship between cues and reward is incidental to the task participants are engaged in. Most of this work has been in visual search and related tasks requiring allocation of spatial attention, and hence we focus on this area, though we will also cover approaches based on nonspatial tasks.

Learned value in spatial tasks. There are three ways in which spatial tasks have been used to look at learned value for task-irrelevant cues. Firstly, cues can be task-relevant in a training phase where they are paired with reward, but then task-irrelevant on a later test. On test, attention to the pretrained cues is measured by the extent to which they distract attention from the task participants are engaged in. We refer to these as *training phase – test phase* procedures. Secondly, in *intertrial priming* procedures, cues can predict reward on one trial, and then serve as distractors on the next trial. Finally, cues can be scheduled to signal reward, but never be the target of the task participants are performing. Below we review evidence from each of these three approaches for an influence of learned value on automatic attentional capture.

Training phase – test phase procedures. Perhaps the most common procedure for studying the relationship between learned value and attention was developed by Anderson et al. (2011b). Each trial of an initial training phase presented six differently coloured circles (Figure 2a). Each display contained a target circle, which could be red or green. Participants responded as rapidly as possible to the orientation (vertical or horizontal) of a line inside the target circle. Fast, correct responses were rewarded, with the amount of reward related to the colour of the target (red or green) on that trial. One of the target colours (the high-value colour) was typically paired with high reward (5¢); the other, low-value target colour was typically paired with low reward (1¢).

After extensive training on this task, participants moved on to a test phase, in which no further rewards were delivered. During the test phase, the target on each trial was now defined by shape; say, a diamond among circles (Figure 2b). Hence colour was now task-irrelevant, since it did not define the target. Critically, people were slower to respond to the shape-singleton target if the display contained a distractor shape in the high-value colour (known as high-value distractor trials) than if neither of the reward-related colours was present in the display (no-distractor trials: note, the comparison here does not include low-value distractor trials, we return to this issue later). This finding suggests that participants' prior experience with the high-value target colour during the training phase made it more likely to distract attention away from the target in the test phase. Notably, since colour was task-irrelevant in the test phase, there was no reason to strategically orient attention to red or green stimuli. The implication is that the high-value distractor colour produced automatic and involuntary capture of attention. Several subsequent studies have replicated the finding of slower responses during the test phase on high-value distractor trials compared to no-distractor trials (Anderson, Faulkner, Rilee, Yantis & Marvel, 2013; Anderson, Laurent & Yantis, 2013, 2014; Anderson, Leal, Hall, Yassa & Yantis, 2014; Anderson & Yantis, 2012;

Laurent, Hall, Anderson & Yantis, 2015; Qi, Zeng, Ding & Li, 2013; Roper, Vecera & Vaidya, 2014; Sali, Anderson & Yantis, 2014; Wang, Yu & Zhou, 2013).

However, there is a problem of interpretation here. Attentional capture by the high-value colour (say, red) during the test phase may not be driven by its learned value, but because this colour defined the target during the preceding training phase. So the fact that red continues to capture attention during the test phase may simply reflect a difficulty in inhibiting an old task-set (locate red and green stimuli) and adopting a new one (locate shape singleton), regardless of reward. In other words, red may capture attention because it has been responded to as a target many times in the past, an effect known as *selection history* (Awh, Belopolsky & Theeuwes, 2012; Kyllingsbaek, Schneider & Bundesen, 2001).

Anderson et al. (2011b; see also Qi et al., 2013; Wang et al., 2013) attempted to rule out this non-reward-based account. They repeated their original study, but provided no rewards during the initial training phase. In this *unrewarded control* condition, no significant difference in response time was observed during the test phase between trials featuring a distractor in a colour that had been a target during training (red or green) versus no-distractor trials. This null result suggests that, in this procedure, selection history alone does not produce capture by former targets. The implication is that capture by the former high-reward target in the rewarded version of the task is a product of its learned value, and not merely because it has previously been a target. Caveats are needed here, though. Neither Anderson et al. nor Qi et al. tested whether capture by the former target was significantly greater in a rewarded procedure than in an unrewarded procedure, so these studies do not provide a robust test of the claim that reward is the critical variable. Wang et al. are alone in testing this critical interaction, which they did find to be significant. The majority of studies that rely on a comparison between high-value distractor versus no-distractor trials to claim an effect of learned value on attention do not include an unrewarded control condition at all, so evidence

from these studies should be regarded with caution (Anderson, 2015a, 2015b; Anderson, Faulkner, et al., 2013; Anderson, Laurent, et al., 2013, 2014; Anderson, Leal, et al., 2014; Anderson & Yantis, 2012; Laurent et al., 2015; Roper et al., 2014).

There is a simple solution to this problem. Recall that Anderson et al. (2011b) trained two colours in their training phase, one paired with high reward and the other with low reward. If, on test, the high-value colour were more likely to produce capture than the low-value colour, this *must* reflect an influence of reward prediction since this comparison controls for all other differences (e.g., both high- and low-value colours were equally likely to appear as targets during training and hence selection history is equal for both). This more stringent, high-versus-low difference has been reported in some recent studies.[5] These studies thus demonstrate conclusively that prior learning about the value of an outcome that is associated with a stimulus produces a change in the extent to which that stimulus will automatically capture attention in the future. This effect can be termed *value-modulated attentional capture* (VMAC). These experiments further suggest that: (1) High-value distractors are more likely to capture attention than low-value distractors even when the location of the target is known in advance (Wang et al., 2014: see also Munneke, Hoppenbrouwers & Theeuwes, 2015). (2) VMAC may be long-lived, lasting for months after the initial training phase (Anderson & Yantis, 2013). (3) Learned value can produce automatic capture of eye movements (known as oculomotor capture). Theeuwes and Belopolsky (2012) used a task similar to that of Anderson et al. (2011b), but with saccades to the target stimulus as the response. During the test phase, participants were more likely to make saccades towards a high-value distractor than a low-value distractor, even though these stimuli were task-irrelevant. (4) VMAC is not restricted to learning about reward. Two studies have found greater attentional capture by stimuli paired with large *loss* of money relative to stimuli paired with small (or no) loss (Wang et al., 2013; Wentura et al., 2014).

Wang et al. also reported greater capture by a stimulus paired with painful shock than by a stimulus that was never paired with shock (Van Damme, Crombez, Hermans, Koster & Eccleston, 2006 report a similar finding; for a further demonstration of a bias in spatial attention towards a stimulus paired with an aversive event [threat], see Field, 2006). These findings suggest the crucial determinant of VMAC is the *motivational significance* of the predicted outcome (large rewards and large punishments are more arousing events than are small rewards/punishments) rather than the *valence* of that outcome (whether it is affectively positive [appetitive] or negative [aversive]: cf. Kahnt, Park, Haynes & Tobler, 2014).

Later we shall consider further the nature of the VMAC effect implicated in these studies—whether it reflects the development of an 'attentional response habit' towards stimuli that produce reward, or an increase in the effective salience of those stimuli. First, though, we shall look at a related case, but in which reward-learning exerts a shorter-term effect on attentional selection.

Inter-trial priming procedures. In the studies described in the previous subsection, value relationships were constant over the training phase (e.g., responses to red targets were paired with large reward, and green targets with low reward, throughout training). The effect of these 'long-term' associations then lingered to influence attention in a test phase which occurred minutes (e.g., Anderson et al., 2011b), days (e.g., Della Libera & Chelazzi, 2009), or months (e.g., Anderson & Yantis, 2013) later. Other studies have instead looked at short-term effects of reward learning on attention. Imagine a search task in which the target on each trial is either red or green, chosen at random. Suppose that participants must respond to a red target on trial T; this will produce a performance benefit on the next trial (T+1) if it also features a red target, but a performance decrement if trial T+1 has a green target while a red stimulus appears as a distractor (Maljkovic & Nakayama, 1994, 1996, 2000). This short-term effect, wherein target choice on the previous trial influences performance on the current trial,

is known as *inter-trial priming*. Notably, several studies of inter-trial priming have provided rewards for responding to the target. Whether high or low reward is provided on each trial is decided randomly – unlike in the studies described in the previous subsection, there is no consistent relationship between a particular colour and reward magnitude. Crucially, if the response to (say) a red target on trial T is highly rewarded, the effect of inter-trial priming on trial T+1 is magnified relative to the case where the response to that same red target on trial T received low reward.[6] This pattern is again consistent with the idea that pairing a stimulus with high reward increases the extent to which that stimulus subsequently captures attention, making it easier to detect when it is a target and harder to ignore when it is a task-irrelevant distractor. Effectively this inter-trial priming procedure condenses the training and test phases of the studies in the previous subsection into a pair of trials, where trial T is the training phase, and trial T+1 is the test phase. In combination, then, it seems that learning of 'stable' reward associations can exert a long-term influence on attentional capture, while experience of individual and unpredictable stimulus–reward pairings can have a similar, but short-term, effect on capture.

Learned value influences capture by stimuli that have never been task-relevant. To reiterate, studies using distinct training and test phases, or inter-trial priming, share a common framework. Initially, responding to a stimulus is paired with high-value reward (either in a training phase or on trial T), and later that stimulus is found to be more likely to capture attention when it is task-irrelevant (in a test phase, or when it is presented as a distractor on trial T+1). Taking a step back, the fact that the reward-predictive stimulus was task-relevant during training (or on trial T)—i.e., it was the target that was responded to—means that participants needed to quickly orient to this stimulus in order to earn the reward on each training trial. As a result, it is possible that attentional capture by reward-related stimuli during the test phase (or on trial T+1) was a carryover of the attentional and/or oculomotor

orienting response that was initially trained. A fundamental principal of reward learning is that following an action with reward increases the likelihood that the action will occur in the future (the so-called Law of Effect: Thorndike, 1911). Perhaps then, it is not surprising that a rapid orienting response that was followed by large reward during training continues to occur in a subsequent test phase, even when the response is no longer relevant to the task goals. That is, perhaps the attentional and oculomotor orienting observed in these studies reflects a learned (conditioned) response that is automatically re-enacted whenever the relevant conditioned stimulus appears, i.e., an 'attentional habit'. Now we will consider evidence for another form of automatic attentional bias. Here, rather than training an attentional response to task-relevant cues, we examine whether cues paired with high-value outcomes can become more likely to capture attention—can increase in psychological salience—even though they were *never* task-relevant.

In a study by Le Pelley, Pearson, Griffiths and Beesley (2015), participants had to move their eyes as quickly as possible to a diamond-shaped target among circles, on every trial. A distractor circle could appear in either a high-value colour or a low-value colour (red or blue, counterbalanced); all other stimuli were grey. On trials with a distractor circle in the high-value colour, rapid saccades to the diamond earned a large reward. On trials with a low-value distractor, rapid saccades to the diamond earned a small reward. Thus while the distractor predicted reward value, it was never the stimulus to which people were required to respond (or direct their attention) to obtain that reward; the reward-predictive distractors were always task-irrelevant. Indeed, if at any point participants looked at the distractor circle, the reward on that trial was cancelled; these were termed *omission trials*. So attending to the distractor was counterproductive to people's goal of maximizing their payoff. The worst thing a participant can do under these conditions is to look at a distractor in the high-value colour, since that results in loss of a high-value reward. And yet that is exactly what people did: high-

value distractors produced significantly more omission trials than did low-value distractors. That is, people were more likely to look at high-value distractors than low-value distractors, even though doing so lost them the large reward that they would otherwise have obtained. This finding of counterproductive VMAC by task-irrelevant stimuli has proven robust (Failing, Nissens, Pearson, Le Pelley & Theeuwes, 2015; Pearson, Donkin, Tran, Most & Le Pelley, 2015: for similar findings from related procedures, see also Bucker, Belopolsky & Theeuwes, 2015; Buschschulte et al. 2014; Hopf et al. 2015; Mine & Saiki, 2015).

These findings show that reward learning can modulate attentional capture by stimuli even when responding to (or orienting attention towards) those stimuli has never been rewarded. The implication is that the influence of learned value on attentional capture (i.e., VMAC) is a function of the *signal-value* of stimuli, rather than their response-value. That is, stimuli which signal the availability of high-value outcomes become more likely to capture attention. In the terminology of conditioning research, these data suggest that VMAC reflects the operation of a Pavlovian, rather than instrumental, learning process.

Extensions of the basic procedure described above show that: (1) Counterproductive VMAC by task-irrelevant stimuli persists over extended training; even with extensive experience, participants did not come to show an adaptive pattern wherein they suppressed attention to the high-value distractor, which would have increased their payoff (Le Pelley et al., 2015). (2) The effect persists when participants are explicitly informed that looking at the distractor results in omission of reward, and are told every time this omission occurs (Pearson et al., 2015). This finding further demonstrates that people are unable to use controlled processes to overcome the effect, even when explicitly motivated to do so. (3) The effect occurs when the reward-predictive distractors are physically salient colour singletons (e.g., a single red circle among grey shapes: Le Pelley et al., 2015; Pearson et al., 2015), and when they are not physically salient (e.g., a red circle among other distinctly coloured shapes, as in

Figure 2a: Failing et al., 2015: see also Bucker et al., 2015). This last finding suggests that learned signal value can modulate the strength of capture that would otherwise by caused by physical salience, and drive capture by non-physically salient stimuli in its own right.

Learned value in nonspatial tasks. Studies described in the previous section investigated the effect of learned value on spatial attention; the extent to which attention is allocated to one location rather than another. But spatial attention is only one aspect of attentional selection. We can also prioritise detection of events that will occur in a known location; e.g., a driver at an intersection knows <u>where</u> the green light will appear, but not <u>when</u>. A handful of studies have shown that *non*spatial perceptual processing of task-irrelevant stimuli is also influenced by learned value; either via monetary reward (Della Libera & Chelazzi, 2009; Della Libera, Perlato & Chelazzi, 2011; Failing & Theeuwes, 2015), or punishment (an aversive loud noise: Smith, Most, Newsome & Zald, 2006). These studies show that stimuli previously paired with high-value outcomes (compared to low-value outcomes) interfere more strongly with processing of—and responding to—a target when they are presented in the same location as that target; either sequentially (Failing & Theeuwes, 2015; Smith et al., 2006),[7] or simultaneously (Della Libera & Chelazzi, 2009; Della Libera et al., 2011; Kelly & Forsyth, 2007). As for the spatial studies discussed earlier, this suggests that stimuli with high learned value captured attention and so reduced attentional resources available for processing the target. However here this capture effect must have been nonspatial, since all stimuli were presented in the same location.

Summary: Learned Value Affects Top-Down and Bottom-Up Attention

Overall, the pattern of data with respect to learned value is rich and interesting. Just as in the case of learned predictiveness, learning about outcome value appears to influence people's allocation of attention in a strategic, top-down way. However, the influence of learned value

on attention is not entirely under our control. Studies of value-modulated attentional capture (VMAC) show that stimuli that predict high-value outcomes (positively or negatively valenced) become more likely to capture our attention (both spatially and nonspatially) in an automatic, bottom-up fashion. This attentional bias could be due to a habitual orienting response (an 'attentional habit') or to a change in effective salience of stimuli brought about by changes in their signal value. It is not possible to say for certain which of these two potential mechanisms is responsible for the effects seen in the *training phase – test phase* and the *intertrial priming* procedures described above – both can explain the observed data. We can be more certain, however, of the mechanism at work in procedures in which reward-predictive stimuli are task-irrelevant throughout. Here, attentional orienting to the high-value cues is never rewarded and so it is not possible for an attentional habit to have developed. The attentional bias must, therefore, be a consequence of the signal value of the cues and hence their ability to capture attention. We now turn to potential theoretical models of all of the empirical data we have reviewed.

Modelling Effects of Predictiveness and Value

In the previous section, we saw that there is now abundant evidence demonstrating that stimuli associated with high-value outcomes receive greater attention than those associated with low-value outcomes. Earlier in the article, we reviewed the substantial evidence in support of the predictiveness principle – stimuli that are more predictive of outcomes (i.e., that provide more information regarding which outcome will occur) tend to receive more attention than stimuli that are less predictive. This finding fits well with Mackintosh's (1975) model and its relatives. In contrast, the Pearce-Hall model, and the uncertainty principle on which it is based, did not receive strong support from the data. In this section we will consider how models of the predictiveness principle can be extended such that both

predictiveness and learned value might be reconciled within a simple, single framework.

Applying Mackintosh (1975) to Learned Value

As noted earlier, the suggestion of a relationship between learning and attention is not novel. William James described the possibility in 1890, and formal attentional models of associative learning have existed for over 50 years (Mackintosh, 1975, provides an early review). Most of the previous research on attentional learning in the associative tradition has tended to focus on learned predictiveness and uncertainty, rather than learned value, and the development of theories reflects this (e.g., Kruschke, 2001b; Le Pelley, 2004; Mackintosh, 1975; Pearce & Hall, 1980). Nevertheless, these theories can potentially also be applied to account for effects of learned value on attention. Consider Mackintosh's (1975) model, which implements the predictiveness principle. This model states that following each learning trial, the associative strength of each presented stimulus A (V_A) is updated according to the following equation:

$$\Delta V_A = S\alpha_A(\lambda - V_A) \qquad\qquad [1],$$

where S is a fixed learning-rate parameter. As discussed earlier, the *prediction error* ($\lambda - V_A$) represents the discrepancy between the actual magnitude of the outcome on a trial (λ), and the extent to which stimulus A predicts that outcome (associative strength of A, V_A). Critically, α_A is a variable representing the attention paid to stimulus A. Let X represent all other stimuli that are presented simultaneously with stimulus A on a given trial. According to Mackintosh's original model, α is determined by comparing how well the outcome is predicted by A (given by the absolute value of the prediction error for A, $|\lambda - V_A|$), with how well the outcome is predicted by X (given by $|\lambda - V_X|$). If A is a better predictor of the outcome than is X, then attention to A (α_A) should increase; if A is a poorer predictor, then α_A should decrease. However, beyond this general principle, Mackintosh gave no specific algorithm for determining changes in α. Following Le Pelley (2004), perhaps the simplest expression for updating α_A that is consistent with Mackintosh's principle is:

$$\Delta\alpha_A = \theta(|\lambda - V_X| - |\lambda - V_A|) \qquad [2],$$

where θ is a fixed rate parameter, and α_A is constrained to lie between a lower limit (representing minimal attention: here we use 0.1) and an upper limit (representing maximal attention: here we use 1). Further details and simulations of this model are provided by Le Pelley (2004).

In this model, attention to predictive stimuli will tend to increase towards the upper limit, regardless of exactly what outcome they predict. However, the *rate* of this increase depends on the value of the outcome, λ. This is because, early in training when V_A is small, a large value of λ will produce a large prediction error in Equation [1] and hence rapid learning. Thus cue A will rapidly become established as a better predictor of the outcome than are other simultaneously-presented stimuli, X (e.g., the experimental context), so attention to cue A will increase quickly according to Equation [2]. Consequently, at least early in training this model correctly anticipates greater attention to stimuli that predict a high-value outcome (e.g., a large reward) than those that predict a low-value outcome (Figure 3a). However, at asymptote the model anticipates that attention will depend on learned predictiveness (attention will be greater to predictive than nonpredictive cues) but not learned value (i.e., attention will not depend on the value of the outcome that a predictive cue predicts). That is, this model anticipates that the influence of learned value on attention will decrease as training proceeds. However, this pattern runs against the findings of Le Pelley et al. (2015), who found that the effect of learned value showed no sign of decreasing over extended training, and if anything tended to increase as training continued.

A Simpler Model: Attention Is Determined by Associative Strength

In fact, it is straightforward to modify the approach described above so that it is better-equipped to account for both learned value and the predictiveness principle, even after extended training. Rather than basing attention on a comparison of the predictiveness of

different stimuli (as in Equation [2]), an alternative approach has attention to a cue determined by the absolute associative strength of that cue:

$$\alpha_A = |V_A| \qquad [3],$$

Once again we set a lower limit for α of 0.1, representing minimal attention.[8] The resulting model still accounts for most demonstrations of an attentional advantage for predictive over nonpredictive stimuli, because the predictive stimuli in these studies typically have greater associative strength – since a predictive cue is consistently paired with the same outcome, it will develop a strong association with that outcome. Notably, in this alternative model, attention is also a direct function of learned value. This is because asymptotic associative strengths for stimuli paired with high-value outcomes will be greater than for stimuli paired with low-value outcomes (Figure 3b). (Formally: According to Equation [1], learning reaches asymptote when $V_A = \lambda$; since asymptotic V_A depends on outcome magnitude λ, then according to Equation [3], asymptotic α_A will also depend on λ). A more complex, and probably more representative, model implementing attentional learning along these lines has recently been developed by Esber and Haselgrove (2011).

The idea that attention towards a cue increases to the extent that it predicts a high-value outcome— attention is determined by associative strength—is very intuitive, and is consistent with the idea that attention goes to cues that are known to be significant (Mackintosh, 1975). It is also supported by studies from the learned predictiveness literature which indicate that attention is modulated by absolute predictiveness (e.g., how strongly do apples predict illness? How strongly do bananas predict illness?), as in Equation 3, rather than a comparison of the relative predictiveness of different cues presented on the same trial (which of apples or bananas is the better predictor of illness?), as in Equation 2 (Beesley & Le Pelley, 2010; Kattner, 2015; Le Pelley, Turnbull, et al., 2010; Livesey et al., 2011).

This simpler model, in which attention is proportional to associative strength, is

generally successful in accounting for empirical effects of both learned predictiveness and learned value. However, a challenge comes from another phenomenon – the inverse base rate effect. Here, there is some evidence that predictive cues command more attention than their less predictive competitors (i.e., that attention is determined by a comparison of relative predictiveness, as specified by Mackintosh, 1975) even when the less predictive cues have higher associative strength.

The inverse base rate effect. The inverse base-rate effect (IBRE: Medin & Edelson, 1988), and the analogous highlighting effect (Collins, Percy, Smith & Kruschke, 2011; Dennis & Kruschke, 1998; Kruschke, 1996a, 2009; Kruschke et al., 2005; Sewell & Lewandowsky, 2012), seem to conflict with the idea that attention can simply be equated with associative strength. In the IBRE, a compound of two cues predicts a commonly occurring outcome (AB→*common*), while an overlapping cue-compound predicts a rarer outcome (AC→*rare*). Cue B is thus a perfect predictor of the common outcome, C is a perfect predictor of the rare outcome, and A is an imperfect predictor. On test, when presented with compound BC participants typically report expecting the rare outcome (predicted by C). This is noteworthy, as both B and C are equally predictive of their respective outcomes, and since the outcome of B is more common, a rational person should predict the common outcome when presented with BC.

Of course, cue familiarity is not equated in the IBRE design, and so the relative novelty of cue C provides one possible explanation for why participants seem to place more weight on it. However, simply pairing one cue with its outcome more often than another cue does not produce the effect: the shared cue A is crucial. Thus, training on FD→*common* and GE→*rare* leads to the reverse of the IBRE effect. That is, participants select the common disease in response to DE (Kruschke, 2001a; Medin & Edelson, 1988; Medin & Robbins, 1971; Wills,

Lavric, Hemmings & Surrey, 2014). It is not, therefore, cue-familiarity that drives the IBRE.

Instead, choice of the rare outcome when presented with BC can be explained if it is assumed that some aspect of participants' experience has caused them to attend more to C than to B. Support for an attentional interpretation of the IBRE comes from eye-tracking findings showing greater overt attention to C than to B following $AB \rightarrow common$ and $AC \rightarrow rare$ training (Kruschke et al., 2005). Furthermore, using EEG, Wills et al. (2014) showed that a validated neural marker of selective attention was also modulated by IBRE training, consistent with greater attention to C than to B.

Why is attention to C greater than to B in the IBRE? It is at this point that our simple model, in which attention is determined by associative strength (Equation [3]) is found wanting. In some studies of the IBRE, as well as being tested with the BC compound, participants have also been presented with B alone and C alone during the test phase. In many (but not all) cases where an IBRE is observed, accuracy in predicting the common outcome for cue B alone numerically exceeds accuracy for predicting the rare outcome for cue C alone (Bohil, Markman & Maddox, 2005; Medin & Edelson, 1988; Wills et al., 2014; Winman, Wennerholm, Juslin & Shanks, 2005); in one case (Wills et al., 2014) this advantage for B over C was tested statistically, and shown to be significant. So the IBRE can be observed even when C has (if anything) a lower associative strength than B. As such, we cannot easily explain an attentional advantage for C over B in terms of our model which equates attention with associative strength.

In contrast, the attentional advantage for C over B emerges more naturally from an account in which attention is based on a comparison of the relative predictiveness of different cues presented together on a trial, as in Mackintosh's (1975) original model (Equation [2]). Briefly, the shared cue A is more predictive of the common outcome presented on AB trials than it is the rare outcome presented on AC trials. Hence cue A competes more effectively for

attention with B than it does with C, and so attention to C increases relative to B.

One possible way of reconciling the IBRE findings with our model equating attention with associative strength is to appeal to the role of the experimental context. On the basis of our model, the attentional advantage for C over B in the IBRE suggests that the B→*common* association is weaker than the C→*rare* association. If the context comes to be associated more strongly with the common outcome than the rare outcome, this could explain why people are nevertheless more likely to correctly choose *common* when presented with B alone than they are to choose *rare* when presented with C alone. This possibility remains to be tested empirically.

In conclusion, our attempt to develop a unified theory of learned predictiveness and learned value has met with some success. *Most of the data presented in this review are captured by the idea that learning increases attention to cues to the extent that those cues acquire associative strength.* The IBRE represents the single exception to this rule of which we are aware. It may be that future research will show how the IBRE can be reconciled with the simple model, or we may have to think along different lines in order to account for all of the data within a single model. We turn now to a final summary and concluding comments on the literature concerning associative learning and attention.

Summary and Conclusions

There is an abundance of evidence for attentional learning in humans. Attention is influenced both by learning about the predictiveness of a stimulus (how much information it provides about other, outcome events) and about its value (how significant—rewarding, punishing—those outcome events are). Furthermore, both of these phenomena manifest themselves as both top-down (controlled) and bottom-up (automatic) attentional biases. Taken together, these findings suggest a fundamental relationship between learning and attention, wherein

learning leads to prioritisation—both controlled and automatic—of those stimuli likely to be of greatest importance – cues that predict motivationally significant events. In turn, studies of associability show that the resulting attentional biases can themselves modulate the rate of future learning about the stimuli involved (predictive cues are subsequently learned about faster), such that the attention–learning relationship is truly interactive.

We propose that the effects of learned predictiveness and learned value are unlikely to be the consequence of entirely separate processes or systems. Intuitively, these two effects seem intimately related. So it would be desirable, if possible, to explain all of the observed phenomena within a single model. Our attempts above met with some success. A simple model in which a cue's associative strength modulates attention explains most of the available data. Some evidence for a dissociation between associative strength and attention can be seen, however, in the inverse base rate effect; in this case, attention seems to go to the target cue that has lower associative strength. More theoretical and empirical work is thus required to provide a single simple model of learned predictiveness and learned value.

Clinical Application

Beyond its theoretical importance, an improved understanding of the processes underlying the relationship between learning and attention has clinical significance. In particular, many drugs of abuse produce potent neural reward signals (Dayan, 2009; Hyman, 2005; Robinson & Berridge, 2001). The ideas advanced in this article might lead us to expect that drug users will show enhanced attention to stimuli that are associated with these drug rewards (such as drug paraphernalia, or people and locations associated with drug supply). This is exactly what is observed: for example, people who are substance-use dependent show particular difficulty in processing a nonsalient target stimulus in the presence of a drug-related distractor stimulus (for reviews, see Field & Cox, 2008; Field, Munafo & Franken, 2009).

In fact, recent work using a value-modulated capture procedure suggests that addicts

also show a heightened attentional bias for stimuli associated with *non*drug (monetary) reward (Anderson, Faulkner, et al., 2013). This finding suggests that individual differences in susceptibility to substance abuse disorders may reflect, in part, differences in the level of a more general predisposition to develop attentional biases towards reward-related stimuli. Moreover, the extent of the attentional bias towards drug-related stimuli in recovering addicts is known to predict relapse (Cox, Hogan, Kristian & Race, 2002; Marissen et al., 2006; Waters et al., 2003). Such research raises the possibility that if such attentional biases could be reduced, the likelihood of relapse might also fall.

The issue of the automaticity of value-modulated attention is relevant here. We saw earlier that the influence of reward learning on attention reflects (at least in part) an automatic, involuntary process that is cognitively impenetrable. The implication is that mere instruction is likely to be ineffective in reducing the magnitude of maladaptive attentional biases. Instead, attentional re-training may be required for a full resolution. Consistent with this idea, a recent large-scale study of alcohol-dependent individuals has shown that a cognitive re-training procedure can indeed reduce relapse rates at one year by up to 10% (Eberl et al., 2013; see also Fadardi & Cox, 2009; Schoenmakers et al., 2010). Work has just begun to emerge investigating the effect of reward on attentional bias training for clinical disorders (Sigurjonsdottir, Bjornsson, Ludvigsdottir & Kristjansson, 2015), and this area represents an exciting potential application of the findings from the learned attention literature.

Conclusion

The survey above describes a research area that is rich in findings and also rich in potential. There are aspects of learned attention that we now understand well. The other areas, about which we know considerably less, will no doubt become the focus of intense investigation. There are four areas that we have identified as being (at least to us) especially interesting.

Firstly, there is abundant evidence for the predictiveness principle (in which attention is

preferentially allocated to cues that accurately predict their consequences), but relatively little support in human studies for the uncertainty principle (in which attention is preferentially allocated to cues whose outcome is currently unknown). These findings in humans stand in contrast to the literature on attentional learning in nonhuman animals, in which there have been many reports of data consistent with the uncertainty principle (see Pearce & Mackintosh, 2010, for a review; see Holland & Maddux, 2010, for a detailed examination of the neural basis of uncertainty-related effects in animals). While the predictiveness principle and the uncertainty principle are seemingly in opposition, they can be reconciled in hybrid attentional models (George & Pearce, 2012; Le Pelley, 2004, 2010; Pearce & Mackintosh, 2010). Future research should aim to establish the role (if any) of the uncertainty principle in humans, in order to determine whether such hybrid models must be considered in the context of human attentional learning. One suggestion (though there are other possibilities) is that evidence for the uncertainty principle might be obtained in learning environments that encourage exploration rather than exploitation (Beesley et al., 2015).

Secondly, we would be keen to know how far a simple model, in which attention is determined by associative strength, can go in explaining effects of both learned predictiveness and learned value. This might begin with further work looking at the inverse base rate effect, which provides possibly the strongest argument against the simple associative strength model.

There are also two further areas of exciting research that have really only just begun. One is the examination of the effects on attention of *outcome predictability* (Griffiths, Mitchell, Bethmont & Lovibond, 2015). There is some suggestion here that *outcomes* that have been predictable in the past might be learned about more easily than outcomes that were unpredictable. This emerging research suggests that there may be some symmetry between cues and outcomes with respect to learned predictiveness.

Finally, there are the fascinating findings concerning the automaticity of learned

attention. That is, when a participant has learned that a cue is predictive of a significant outcome (perhaps one of high value) they then pay attention to that cue in an automatic and involuntary manner. These findings have the potential to change the way in which psychologists think about attention. Often, controlled attention is referred to as *goal-directed* or endogenous, whereas automatic attention is referred to as *stimulus-driven* or exogenous (Corbetta & Shulman, 2002; Jonides, 1981; Posner, 1980; Posner & Cohen, 1984; Yantis, 2000). Goal-directed processes refer to subject-driven attention that encompasses a person's intentions. In contrast, stimulus-driven attentional processes relate to attention-grabbing characteristics that are intrinsic to the stimulus and which we cannot control; its brightness, onset, colour, etc. The automatic attention that participants pay to cues that predict valued outcomes is not goal-directed or under top-down control, but neither is it driven by the intrinsic, physical properties of the stimulus. Rather, the change in effective salience of the cue is the result of training; it comes about as a consequence of the participant's experience with, and memory of, that cue and its relationship to other events. It would appear that what has been uncovered is a form of attentional bias—*derived attention*—that does not fit the standard goal-directed/stimulus-driven distinction. While William James (1890/1983) introduced the concept of derived attention over 125 years ago, the importance of this interface between learning and other aspects of cognition is only now becoming apparent.

<center>**Footnotes**</center>

[1] (Buckley, Smith & Haselgrove, 2014; Dixon, Ruppel, Pratt & De Rosa, 2009; Eimas, 1966; George & Kruschke, 2012; Isaacs & Duncan, 1962; Kruschke, 1996b; Livesey & McLaren, 2007; Macho, 1997; Maes, Damen & Eling, 2004; Maes, Eling, Wezenberg, Vissers & Kan, 2011; Maes, Vich & Eling, 2006; Owen, Roberts, Polkey, Sahakian & Robbins, 1991; Van der Molen et al., 2012; Whitney & White, 1993: see also Ahmed & Lovibond, 2015; Slabosz et al, 2006; Suret & McLaren, 2003; Liu, Mercado, Church & Orduna, 2008)

[2] (Beesley & Le Pelley, 2010; Bonardi, Graham, Hall & Mitchell, 2005; Buckley, Smith & Haselgrove, 2015; Don & Livesey, 2015; Griffiths & Mitchell, 2008; Le Pelley, Beesley & Griffiths, 2011; Le Pelley, Calvini & Spears, 2013; Le Pelley, Oakeshott, Wills & McLaren, 2005; Le Pelley, Reimers, et al., 2010; Le Pelley, Schmidt-Hansen, Harris, Lunter & Morris, 2010; Le Pelley, Suret & Beesley, 2009; Le Pelley, Turnbull, Reimers & Knipe, 2010; Livesey, Thorwart, De Fina & Harris, 2011; Lucke, Lachnit, Koenig & Uengoer, 2013; Mitchell, Griffiths, Seetoo & Lovibond, 2012; Morris, Griffiths, Le Pelley & Weickert, 2013; Shone, Harris & Livesey, 2015: see also Kersten, Goldstone & Schaffert, 1998)

[3] Some studies of blocking have found evidence of greater overt attention to the pretrained cue A than to the blocked cue B (Ellis et al., 2014; Le Pelley et al., 2007). This finding is consistent with the predictiveness principle (since pretraining renders A a better predictor of the outcome), but not with the uncertainty principle (since pretraining means that the outcome on AB trials is, if anything, less surprising given the presence of A than the presence of B). However, the difference in predictiveness between A and B is confounded with a difference

in familiarity (pretraining means that A is experienced more times than has B). Given the possibility that familiarity might exert an independent effect on attention, the finding of greater attention to A than B provides somewhat ambiguous support for the predictiveness principle over the uncertainty principle.

[4] (Balcetis, Dunning & Granot, 2012; Brielmann & Spering, 2015; Bucker, Silvis, Donk & Theeuwes, 2015; Chapman, Gallivan & Enns, 2015; Chapman, Gallivan, Wong, Wispinski & Enns, 2015; Chelazzi et al., 2014; Chen, Mihalas, Niebur & Stuphorn, 2013; Della Libera & Chelazzi, 2009; Doallo, Patai & Nobre, 2013; Eckstein, Schoonveld, Zhang, Mack & Akbas, 2015; Field et al., 2011; Garland, Froeliger, Passik & Howard, 2013; Hardman, Scott, Field & Jones, 2014; Hickey & Peelen, 2015; Hogarth, Dickinson & Duka, 2005, 2009; Hogarth, Dickinson, Hutton, Elbers & Duka, 2006; Hogarth, Dickinson, Janowski, Nikitina & Duka, 2008; Hogarth, Dickinson, Wright, Kouvaraki & Duka, 2007; Kiss, Driver & Eimer, 2009; Kristjansson, Sigurjonsdottir & Driver, 2010; Le Pelley, Mitchell, et al., 2013; Lee & Shomstein, 2014; Li, Howard, Parrish & Gottfried, 2008; Luque et al., 2015; Marini, Marzi & Viggiano, 2011; Marx & Einhauser, 2015; O'Brien & Raymond, 2012; Pascucci, Mastropasqua & Turatto, 2015; Pool, Delplanque, et al., 2014; Raymond & O'Brien, 2009; Reinhart & Woodman, 2014; Rothkirch, Ostendorf, Sax & Sterzer, 2013; Schutz, Trommershauser & Gegenfurtner, 2012; Shomstein & Johnson, 2013; Stankevich & Geng, 2014, 2015; Wallis, Stokes, Arnold & Nobre, 2015; Wilbertz, van Slooten & Sterzer, 2014; Yokoyama, Padmala & Pessoa, 2015). Some of these studies (Chapman, Gallivan & Enns, 2015; Chapman, Gallivan, Wong, et al., 2015; O'Brien & Raymond, 2012; Raymond & O'Brien, 2009; Rothkirch et al., 2013; see also Rutherford, O'Brien & Raymond, 2010) conflated learned value and learned predictiveness: stimuli differed both in the nature of the

outcome they were paired with (gain or loss of money), and the probability with which that outcome occurred. Hence it is unclear which was responsible for the differences in attention observed during the test phase of these studies.

[5] (Anderson, Laurent & Yantis, 2011a, 2012; Anderson & Yantis, 2013; Failing & Theeuwes, 2014; Gong & Li, 2014; Hickey & Peelen, 2015; Infanti, Hickey & Turatto, 2015; Jiao, Du, He & Zhang, 2015; MacLean & Giesbrecht, 2015; Mine & Saiki, 2015; Theeuwes & Belopolsky, 2012; Wang, Duan, Theeuwes & Zhou, 2014; Wentura, Muller & Rothermund, 2014; see also Pool, Brosch, Delplanque & Sander, 2014)

[6] (Asgeirsson & Kristjansson, 2014; Braem, Hickey, Duthoo & Notebaert, 2014; Camara, Manohar & Husain, 2013; Hickey, Chelazzi & Theeuwes, 2010a, 2010b, 2011, 2014; Hickey, Kaiser & Peelen, 2015; Hickey & Los, 2015; Hickey & van Zoest, 2013; see also Itthipuripat, Cha, Rangsipat & Serences, 2015)

[7] In the study by Failing and Theeuwes (2015), the difference in reward history of the different categories during the test phase is confounded with a difference in their selection history: high-value stimuli were chosen more frequently as targets during the preceding training phase, raising the possibility that it is this greater selection history (rather than the difference in learned value) that drives greater capture by these stimuli on test.

[8] One problem that immediately arises from this change is that the resulting model is no longer able to account for cue competition effects such as overshadowing and blocking (described earlier). The simple solution is to amend Equation [1] to use a *summed prediction error*, such that changes in the associative strength of cue A are based on the summed

associative strength of all presented cues (ΣV), i.e. $\Delta V_A = S\alpha_A(\lambda - \sum V)$. This expression for changes in associative strength, which is the same as that of the well-known model by Rescorla and Wagner (1972), allows the model to account for cue competition effects such as overshadowing and blocking (see Le Pelley, 2004, for detailed discussion). For example, the model anticipates that blocked cues will gain less associative strength than control cues, and this will result in reduced attention (α) to these blocked cues by Equation [3], as observed empirically (and this reduced α will also slow further learning about these blocked cues based on the equation for ΔV_A above).

References

[a] marks articles retrieved from the ISI Web of Science database search for articles on learned predictiveness (Search 1 – see main text); [b] marks additional articles on learned predictiveness found via references in articles from this search. [c] marks articles retrieved from the ISI Web of Science database search for articles on learned value (Search 2 – see main text); [d] marks additional articles on learned value found via references in articles from this search.

Abad, M.J.F., Ramos-Alvarez, M.M., & Rosas, J.M. (2009). Partial reinforcement and context switch effects in human predictive learning. *Quarterly Journal of Experimental Psychology, 62*, 174-188. [a]

Ahmed, O., & Lovibond, P.F. (2015). The impact of previously learned feature-relevance on generalisation of conditioned fear in humans. *Journal of Behavior Therapy and Experimental Psychiatry, 46*, 59-65. [a]

Anderson, A.K., & Phelps, E.A. (2001). Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature, 411*, 305-309.

Anderson, B.A. (2015a). Value-driven attentional capture is modulated by spatial context. *Visual Cognition, 23*, 67-81. [c]

Anderson, B.A. (2015b). Value-driven attentional priority is context specific. *Psychonomic Bulletin & Review, 22*, 750-756. [c]

Anderson, B.A., Faulkner, M.L., Rilee, J.J., Yantis, S., & Marvel, C.L. (2013). Attentional bias for nondrug reward is magnified in addiction. *Experimental and Clinical Psychopharmacology, 21*, 499-506. [c]

Anderson, B.A., Laurent, P.A., & Yantis, S. (2011a). Learned value magnifies salience-based attentional capture. *Plos One, 6*, 6. [c]

Anderson, B.A., Laurent, P.A., & Yantis, S. (2011b). Value-driven attentional capture. *Proceedings of the National Academy of Sciences of the United States of America, 108*, 10367-10371. [c]

Anderson, B.A., Laurent, P.A., & Yantis, S. (2012). Generalization of value-based attentional priority. *Visual Cognition, 20*, 647-658. [c]

Anderson, B.A., Laurent, P.A., & Yantis, S. (2013). Reward predictions bias attentional selection. *Frontiers in Human Neuroscience, 7*, 6. [c]

Anderson, B.A., Laurent, P.A., & Yantis, S. (2014). Value-driven attentional priority signals in human basal ganglia and visual cortex. *Brain Research, 1587*, 88-96. [c]

Anderson, B.A., Leal, S.L., Hall, M.G., Yassa, M.A., & Yantis, S. (2014). The attribution of value-based attentional priority in individuals with depressive symptoms. *Cognitive Affective & Behavioral Neuroscience, 14*, 1221-1227. [c]

Anderson, B.A., & Yantis, S. (2012). Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. *Attention Perception & Psychophysics, 74*, 1644-1653. [c]

Anderson, B.A., & Yantis, S. (2013). Persistence of value-driven attentional capture. *Journal of Experimental Psychology-Human Perception and Performance, 39*, 6-9. [c]

Aron, A.R., Shohamy, D., Clark, J., Myers, C., Gluck, M.A., & Poldrack, R.A. (2004). Human midbrain sensitivity to cognitive feedback and uncertainty during classification learning. *Journal of Neurophysiology, 92*, 1144-1152. [a]

Asgeirsson, A.G., & Kristjansson, A. (2014). Random reward priming is task-contingent: The robustness of the 1-trial reward priming effect. *Frontiers in Psychology, 5*, 10. [c]

Austin, A.J., & Duka, T. (2010). Mechanisms of attention for appetitive and aversive outcomes in Pavlovian conditioning. *Behavioural Brain Research, 213*, 19-26. [a]

Austin, A.J., & Duka, T. (2012). Mechanisms of attention to conditioned stimuli predictive of

a cigarette outcome. *Behavioural Brain Research, 232*, 183-189. [a]

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.

Bach, D.R., Hulme, O., Penny, W.D., & Dolan, R.J. (2011). The known unknowns: Neural representation of second-order uncertainty, and ambiguity. *Journal of Neuroscience, 31*, 4811-4820. [a]

Bai, Y., Katahira, K., & Ohira, H. (2014). Dual learning processes underlying human decision-making in reversal learning tasks: Functional significance and evidence from the model fit to human behavior. *Frontiers in Psychology, 5*, 8. [a]

Balcetis, E., Dunning, D., & Granot, Y. (2012). Subjective value determines initial dominance in binocular rivalry. *Journal of Experimental Social Psychology, 48*, 122-129. [c]

Beck, S.B. (1963). Eyelid conditioning as a function of CS intensity, UCS intensity, and Manifest Anxiety scale score. *Journal of Experimental Psychology, 66*, 429-&.

Beesley, T., & Le Pelley, M.E. (2010). The effect of predictive history on the learning of sub-sequence contingencies. *Quarterly Journal of Experimental Psychology, 63*, 108-135. [a]

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. [a]

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. [b]

Behrens, T.E.J., Woolrich, M.W., Walton, M.E., & Rushworth, M.F.S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience, 10*, 1214-1221. [a]

Bohil, C.J., Markman, A.B., & Maddox, W.T. (2005). A feature-salience analogue of the

inverse base-rate effect. *Korean Journal of Thinking and Problem Solving, 15*, 17-28. [b]

Boll, S., Gamer, M., Gluth, S., Finsterbusch, J., & Buchel, C. (2013). Separate amygdala

subregions signal surprise and predictiveness during associative fear learning in humans.

European Journal of Neuroscience, 37, 758-767. [a]

Bonardi, C., Graham, S., Hall, G., & Mitchell, C. (2005). Acquired distinctiveness and

equivalence in human discrimination learning: Evidence for an attentional process.

Psychonomic Bulletin & Review, 12, 88-92. [a]

Bonardi, C., & Hall, G. (1996). Learned irrelevance: No more than the sum of CS and US

preexposure effects? *Journal of Experimental Psychology: Animal Behavior Processes,

22*, 183-191.

Bouton, M.E. (1993). Context, time, and memory retrieval in the interference paradigms of

Pavlovian learning. *Psychological Bulletin, 114*, 80-99.

Braem, S., Hickey, C., Duthoo, W., & Notebaert, W. (2014). Reward determines the context-

sensitivity of cognitive control. *Journal of Experimental Psychology-Human

Perception and Performance, 40*, 1769-1778. [c]

Brielmann, A.A., & Spering, M. (2015). Effects of reward on the accuracy and dynamics of

smooth pursuit eye movements. *Journal of Experimental Psychology-Human

Perception and Performance, 41*, 917-928. [c]

Broomfield, L., McHugh, L., & Reed, P. (2010). Factors impacting emergence of behavioral

control by underselected stimuli in humans after reduction of control by overselected

stimuli. *Journal of the Experimental Analysis of Behavior, 94*, 125-133. [a]

Bucker, B., Belopolsky, A.V., & Theeuwes, J. (2015). Distractors that signal reward attract

the eyes. *Visual Cognition, 23*, 1-24. [c]

Bucker, B., Silvis, J.D., Donk, M., & Theeuwes, J. (2015). Reward modulates oculomotor

competition between differently valued stimuli. *Vision Research, 108*, 103-112. [c]

Buckley, M.G., Smith, A.D., & Haselgrove, M. (2014). Shape shifting: Local landmarks interfere with navigation by, and recognition of, global shape. *Journal of Experimental Psychology-Learning Memory and Cognition, 40*, 492-510. [a]

Buckley, M.G., Smith, A.D., & Haselgrove, M. (2015). Learned predictiveness training modulates biases towards using boundary or landmark cues during navigation. *Quarterly Journal of Experimental Psychology, 68*, 1183-1202. [a]

Buschschulte, A., Boehler, C.N., Strumpf, H., Stoppel, C., Heinze, H.J., Schoenfeld, M.A., et al. (2014). Reward- and attention-related biasing of sensory selection in visual cortex. *Journal of Cognitive Neuroscience, 26*, 1049-1065. [c]

Callejas-Aguilera, J.E., & Rosas, J.M. (2010). Ambiguity and context processing in human predictive learning. *Journal of Experimental Psychology-Animal Behavior Processes, 36*, 482-494. [a]

Camara, E., Manohar, S., & Husain, M. (2013). Past rewards capture spatial attention and action choices. *Experimental Brain Research, 230*, 291-300. [c]

Cavanagh, J.F. (2015). Cortical delta activity reflects reward prediction error and related behavioral adjustments, but at different times. *Neuroimage, 110*, 205-216. [a]

Chapman, C.S., Gallivan, J.P., & Enns, J.T. (2015). Separating value from selection frequency in rapid reaching biases to visual targets. *Visual Cognition, 23*, 249-271. [c]

Chapman, C.S., Gallivan, J.P., Wong, J.D., Wispinski, N.J., & Enns, J.T. (2015). The snooze of lose: Rapid reaching reveals that losses are processed more slowly than gains. *Journal of Experimental Psychology-General, 144*, 844-863. [c]

Chelazzi, L., Estocinova, J., Calletti, R., Lo Gerfo, E., Sani, I., Della Libera, C., et al. (2014). Altering spatial priority maps via reward-based learning. *Journal of Neuroscience, 34*, 8594-8604. [c]

Chelazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual

selective attention. *Vision Research, 85*, 58-72. [c]

Chen, X.M., Mihalas, S., Niebur, E., & Stuphorn, V. (2013). Mechanisms underlying the influence of saliency on value-based decisions. *Journal of Vision, 13*, 23. [c]

Chun, M.M., & Potter, M.C. (1995). A two-stage model for multiple target detection in RSVP. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 109-127.

Collins, E.C., Percy, E.J., Smith, E.R., & Kruschke, J.K. (2011). Integrating advice and experience: Learning and decision making with social and nonsocial cues. *Journal of Personality and Social Psychology, 100*, 967-982. [a]

Cools, R., Rogers, R., Barker, R.A., & Robbins, T.W. (2010). Top-down attentional control in Parkinson's Disease: Salient considerations. *Journal of Cognitive Neuroscience, 22*, 848-859. [a]

Corbetta, M., & Shulman, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience, 3*, 201-215.

Cox, W.M., Hogan, L.M., Kristian, M.R., & Race, J.H. (2002). Alcohol attentional bias as a predictor of alcohol abusers' treatment outcome. *Drug and Alcohol Dependence, 68*, 237-243.

D'Ardenne, K., Lohrenz, T., Bartley, K.A., & Montague, P.R. (2013). Computational heterogeneity in the human mesencephalic dopamine system. *Cognitive Affective & Behavioral Neuroscience, 13*, 747-756. [a]

Dayan, P. (2009). Dopamine, reinforcement learning, and addiction. *Pharmacopsychiatry, 42 (Suppl. 1)*, S56-S65.

Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science, 20*, 778-784. [c]

Della Libera, C., Perlato, A., & Chelazzi, L. (2011). Dissociable effects of reward on

attentional learning: From passive associations to active monitoring. *Plos One, 6*, 6. [c]

Dennis, S., & Kruschke, J.K. (1998). Shifting attention in cued recall. *Australian Journal of Psychology, 50*, 131-138. [b]

Denton, S.E., & Kruschke, J.K. (2006). Attention and salience in associative blocking. *Learning & Behavior, 34*, 285-304. [a]

Deubel, H., & Schneider, W.X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Research, 36*, 1827-1837.

Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge: Cambridge University Press.

Dixon, M.L., Ruppel, J., Pratt, J., & De Rosa, E. (2009). Learning to ignore: Acquisition of sustained attentional suppression. *Psychonomic Bulletin & Review, 16*, 418-423. [a]

Doallo, S., Patai, E.Z., & Nobre, A.C. (2013). Reward associations magnify memory-based biases on perception. *Journal of Cognitive Neuroscience, 25*, 245-257. [c]

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

Dreher, J.C., Kohn, P., & Berman, K.F. (2006). Neural coding of distinct statistical properties of reward information in humans. *Cerebral Cortex, 16*, 561-573. [a]

Duncan, J. (1989). Boundary conditions on parallel processing in human vision. *Perception, 18*, 457-469.

Dunsmoor, J.E., Campese, V.D., Ceceli, A.O., LeDoux, J.E., & Phelps, E.A. (2015). Novelty-facilitated extinction: Providing a novel outcome in place of an expected threat diminishes recovery of defensive responses. *Biological Psychiatry, 78*, 203-209. [a]

Eberl, C., Wiers, R.W., Pawelczack, S., Rinck, M., Becker, E.S., & Lindenmeyer, J. (2013). Approach bias modification in alcohol dependence: Do clinical effects replicate and for whom does it work best? *Developmental Cognitive Neuroscience, 4*, 38-51.

Eckstein, M.P., Schoonveld, W., Zhang, S., Mack, S.C., & Akbas, E. (2015). Optimal and human eye movements to clustered low value cues to increase decision rewards during search. *Vision Research, 113*, 137-154. [c]

Eickhoff, S.B., Pomjanski, W., Jakobs, O., Zilles, K., & Langner, R. (2011). Neural correlates of developing and adapting behavioral biases in speeded choice reactions-an fmri study on predictive motor coding. *Cerebral Cortex, 21*, 1178-1191. [a]

Eimas, P.D. (1966). Effects of overtraining and age on intradimensional and extradimensional shifts in children. *Journal of Experimental Child Psychology, 3*, 348-355. [b]

Eippert, F., Gamer, M., & Buchel, C. (2012). Neurobiological mechanisms underlying the blocking effect in aversive learning. *Journal of Neuroscience, 32*, 13164-13176. [a]

Ellis, N.C., Hafeez, K., Martin, K.I., Chen, L., Boland, J., & Sagarra, N. (2014). An eye-tracking study of learned attention in second language acquisition. *Applied Psycholinguistics, 35*, 547-579. [a]

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: Biological Sciences, 278*, 2553-2561.

Fadardi, J.S., & Cox, W.M. (2009). Reversing the sequence: Reducing alcohol consumption by overcoming alcohol attentional bias. *Drug and Alcohol Dependence, 101*, 137-145.

Failing, M., Nissens, T., Pearson, D., Le Pelley, M.E., & Theeuwes, J. (2015). Oculomotor capture by stimuli that signal the availability of reward. *Journal of Neurophysiology, 114*, 2316-2327. [d]

Failing, M.F., & Theeuwes, J. (2014). Exogenous visual orienting by reward. *Journal of Vision, 14*, 9. [c]

Failing, M.F., & Theeuwes, J. (2015). Nonspatial attentional capture by previously rewarded

scene semantics. *Visual Cognition, 23*, 82-104. [a,c]

Field, A.P. (2006). The behavioral inhibition system and the verbal information pathway to children's fears. *Journal of Abnormal Psychology, 115*, 742-752. [a]

Field, M., & Cox, W.M. (2008). Attentional bias in addictive behaviors: A review of its development, causes, and consequences. *Drug and Alcohol Dependence, 97*, 1-20.

Field, M., Hogarth, L., Bleasdale, D., Wright, P., Fernie, G., & Christiansen, P. (2011). Alcohol expectancy moderates attentional bias for alcohol cues in light drinkers. *Addiction, 106*, 1097-1103. [a]

Field, M., Munafo, M.R., & Franken, I.H. (2009). A meta-analytic investigation of the relationship between attentional bias and subjective craving in substance abuse. *Psychological Bulletin, 135*, 589-607.

Garland, E.L., Froeliger, B.E., Passik, S.D., & Howard, M.O. (2013). Attentional bias for prescription opioid cues among opioid dependent chronic pain patients. *Journal of Behavioral Medicine, 36*, 611-620. [a]

Gauntlett-Gilbert, J., Roberts, R.C., & Brown, V.J. (1999). Mechanisms underlying attentional set-shifting in parkinson's disease. *Neuropsychologia, 37*, 605-616. [b]

Gawronski, B., Rydell, R.J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. *Journal of Experimental Psychology-General, 139*, 683-701. [a]

George, D.N., & Kruschke, J.K. (2012). Contextual modulation of attention in human category learning. *Learning & Behavior, 40*, 530-541. [a]

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

George, D.N., & Pearce, J.M. (2012). A configural theory of attention and associative

learning. *Learning & Behavior, 40*, 241-254.

Glautier, S., & Shih, S.I. (2015). Relative prediction error and protection from attentional blink in human associative learning. *Quarterly Journal of Experimental Psychology, 68*, 442-458. [a]

Goldstone, R. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology-General, 123*, 178-200. [a]

Goldstone, R.L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General, 130*, 116-139. [b]

Gong, M.Y., & Li, S. (2014). Learned reward association improves visual working memory. *Journal of Experimental Psychology-Human Perception and Performance, 40*, 841-856. [c]

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

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

Griffiths, O., & Mitchell, C.J. (2008). Selective attention in human associative learning and recognition memory. *Journal of Experimental Psychology: General, 137*, 626-648. [a]

Griffiths, O., Mitchell, C.J., Bethmont, A., & Lovibond, P.F. (2015). Outcome predictability biases learning. *Journal of Experimental Psychology-Animal Learning and Cognition, 41*, 1-17. [a]

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

Hardman, C.A., Scott, J., Field, M., & Jones, A. (2014). To eat or not to eat. The effects of expectancy on reactivity to food cues. *Appetite, 76*, 153-160. [a]

Haselgrove, M., Le Pelley, M.E., Singh, N.K., Teow, H.Q., Morris, R.W., Green, M.J., et al. (2015). Disrupted attentional learning in high schizotypy: Evidence of aberrant salience. *British Journal of Psychology*. Advance online publication. doi: 10.1111/bjop.12175

Hauser, T.U., Iannaccone, R., Stampfli, P., Drechsler, R., Brandeis, D., Walitza, S., et al. (2014). The feedback-related negativity (FRN) revisited: New insights into the localization, meaning and network organization. *Neuroimage, 84*, 159-168. [a]

Hickey, C., Chelazzi, L., & Theeuwes, J. (2010a). Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience, 30*, 11096-11103. [c]

Hickey, C., Chelazzi, L., & Theeuwes, J. (2010b). Reward guides vision when it's your thing: Trait reward-seeking in reward-mediated visual priming. *Plos One, 5*, 5. [c]

Hickey, C., Chelazzi, L., & Theeuwes, J. (2011). Reward has a residual impact on target selection in visual search, but not on the suppression of distractors. *Visual Cognition, 19*, 117-128. [c]

Hickey, C., Chelazzi, L., & Theeuwes, J. (2014). Reward-priming of location in visual search. *Plos One, 9*, 9. [c]

Hickey, C., Kaiser, D., & Peelen, M.V. (2015). Reward guides attention to object categories in real-world scenes. *Journal of Experimental Psychology-General, 144*, 264-273. [c]

Hickey, C., & Los, S.A. (2015). Reward priming of temporal preparation. *Visual Cognition, 23*, 25-40. [c]

Hickey, C., & Peelen, M.V. (2015). Neural mechanisms of incentive salience in naturalistic human vision. *Neuron, 85*, 512-518. [c]

Hickey, C., & van Zoest, W. (2013). Reward-associated stimuli capture the eyes in spite of strategic attentional set. *Vision Research, 92*, 67-74. [c]

Hoffman, A.B., & Rehder, B. (2010). The costs of supervised classification: The effect of learning task on conceptual flexibility. *Journal of Experimental Psychology-General,*

139, 319-340. [a]

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. [a]

Hogarth, L., Dickinson, A., & Duka, T. (2005). Explicit knowledge of stimulus-outcome contingencies and stimulus control of selective attention and instrumental action in human smoking behaviour. *Psychopharmacology, 177*, 428-437. [a]

Hogarth, L., Dickinson, A., & Duka, T. (2009). Detection versus sustained attention to drug cues have dissociable roles in mediating drug seeking behavior. *Experimental and Clinical Psychopharmacology, 17*, 21-30. [a]

Hogarth, L., Dickinson, A., Hutton, S.B., Elbers, N., & Duka, T. (2006). Drug expectancy is necessary for stimulus control of human attention, instrumental drug-seeking behaviour and subjective pleasure. *Psychopharmacology, 185*, 495-504. [a]

Hogarth, L., Dickinson, A., Janowski, M., Nikitina, A., & Duka, T. (2008). The role of attentional bias in mediating human drug-seeking behaviour. *Psychopharmacology, 201*, 29-41. [a]

Hogarth, L., Dickinson, A., Wright, A., Kouvaraki, M., & Duka, T. (2007). The role of drug expectancy in the control of human drug seeking. *Journal of Experimental Psychology-Animal Behavior Processes, 33*, 484-496. [a]

Holland, P.C., & Maddux, J.-M. (2010). Brain systems of attention in associative learning. In C.J. Mitchell & M.E. Le Pelley (Eds.), *Attention and Associative Learning: From Brain to Behaviour* (pp. 305-349). Oxford, UK: Oxford University Press.

Holmes, N.M., & Harris, J.A. (2010). Latent inhibition. In C.J. Mitchell & M.E. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour* (pp. 99-130). Oxford, UK: Oxford University Press.

Hopf, J.M., Schoenfeld, M.A., Buschschulte, A., Rautzenberg, A., Krebs, R.M., & Boehler, C.N. (2015). The modulatory impact of reward and attention on global feature selection in human visual cortex. *Visual Cognition, 23*, 229-248. [c]

Huang, Y., & Yu, R.J. (2014). The feedback-related negativity reflects "more or less" prediction error in appetitive and aversive conditions. *Frontiers in Neuroscience, 8*, 6. [a]

Hyman, S.E. (2005). Addiction: A disease of learning and memory. *American Journal of Psychiatry, 162*, 1414-1422.

Infanti, E., Hickey, C., & Turatto, M. (2015). Reward associations impact both iconic and visual working memory. *Vision Research, 107*, 22-29. [c]

Isaacs, I.D., & Duncan, C.P. (1962). Reversal and nonreversal shifts within and between dimensions in concept formation. *Journal of Experimental Psychology, 64*, 580-585. [b]

Itthipuripat, S., Cha, K.X., Rangsipat, N., & Serences, J.T. (2015). Value-based attentional capture influences context-dependent decision-making. *Journal of Neurophysiology, 114*, 560-569. [c]

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

Jiao, J., Du, F., He, X.S., & Zhang, K. (2015). Social comparison modulates reward-driven attentional capture. *Psychonomic Bulletin & Review, 22*, 1278-1284. [c]

Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. In J.B. Long & A.D. Baddeley (Eds.), *Attention and Performance IX* (pp. 187-203). Hillsdale, NJ: Erlbaum.

Kahnt, T., Park, S.Q., Haynes, J.D., & Tobler, P.N. (2014). Disentangling neural representations of value and salience in the human brain. *Proceedings of the National Academy of Sciences of the United States of America, 111*, 5000-5005. [a]

Kalish, M.L., & Kruschke, J.K. (2000). The role of attention shifts in the categorization of

continuous dimensioned stimuli. *Psychological Research-Psychologische Forschung, 64*, 105-116. [a]

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

Kelly, M.M., & Forsyth, J.P. (2007). Observational fear conditioning in the acquisition and extinction of attentional bias for threat: An experimental evaluation. *Emotion, 7*, 324-335. [a]

Kenny, D.A. (2015). *Measuring model fit.* Retrieved April 28, 2016, from http://davidakenny.net/cm/fit.htm

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. [a]

Kiss, M., Driver, J., & Eimer, M. (2009). Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychological Science, 20*, 245-251. [c]

Klauer, K.C., Roßnagel, C., & Musch, J. (1997). List-context effects in evaluative priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 246-255.

Kopp, B., & Wolff, M. (2000). Brain mechanisms of selective learning: event-related potentials provide evidence for error-driven learning in humans. *Biological Psychology, 51*, 223-246. [a]

Kristjansson, A., Sigurjonsdottir, O., & Driver, J. (2010). Fortune and reversals of fortune in visual search: Reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Attention Perception & Psychophysics, 72*, 1229-1236. [c]

Kruschke, J.K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22-44.

Kruschke, J.K. (1996a). Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 3-26. [b]

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

Kruschke, J.K. (2001a). The inverse base-rate effect is not explained by eliminative inference. *Journal of Experimental Psychology-Learning Memory and Cognition, 27*, 1385-1400. [b]

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

Kruschke, J.K. (2006). Locally Bayesian learning with applications to retrospective revaluation and highlighting. *Psychological Review, 113*, 677-699.

Kruschke, J.K. (2009). Highlighting: A canonical experiment. In B. Ross (Ed.), *The Psychology of Learning and Motivation* (Vol. 51, pp. 153-185). San Diego, CA: Academic Press.

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

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

Kyllingsbaek, S., Schneider, W.X., & Bundesen, C. (2001). Automatic attraction of attention to former targets in visual displays of letters. *Perception & Psychophysics, 63*, 85-98.

Lashley, K.S. (1929). *Brain mechanisms and intelligence*. Chicago, IL: University of Chicago Press.

Laurent, P.A., Hall, M.G., Anderson, B.A., & Yantis, S. (2015). Valuable orientations capture attention. *Visual Cognition, 23*, 133-146. [c]

Lawrence, A.D., Sahakian, B.J., Rogers, R.D., Hodges, J.R., & Robbins, T.W. (1999). Discrimination, reversal, and shift learning in huntington's disease: Mechanisms of impaired response selection. *Neuropsychologia, 37*, 1359-1374. [b]

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

Le Pelley, M.E. (2010). The hybrid modeling approach to conditioning. In N.A. Schmajuk (Ed.), *Computational Models of Conditioning* (pp. 71-107). Cambridge, UK: Cambridge 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. [a]

Le Pelley, M.E., Beesley, T., & Griffiths, O. (2014). Relative salience versus relative validity: Cue salience influences blocking in human associative learning. *Journal of Experimental Psychology-Animal Behavior Processes, 40*, 116-132. [a]

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, UK: Wiley.

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

Le Pelley, M.E., Calvini, G., & Spears, R. (2013). Learned predictiveness influences automatic evaluations in human contingency learning. *Quarterly Journal of Experimental Psychology, 66*, 217-228. [b]

Le Pelley, M.E., & McLaren, I.P.L. (2003). Learned associability and associative change in

human causal learning. *Quarterly Journal of Experimental Psychology Section B-Comparative and Physiological Psychology, 56*, 68-79. [a]

Le Pelley, M.E., Mitchell, C.J., & Johnson, A.M. (2013). Outcome value influences attentional biases in human associative learning: Dissociable effects of training and instruction. *Journal of Experimental Psychology: Animal Behavior Processes, 39*, 39-55. [c]

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. [a]

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

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. [c]

Le Pelley, M.E., Reimers, S.J., Calvini, G., Spears, R., Beesley, T., & Murphy, R.A. (2010). Stereotype formation: Biased by association. *Journal of Experimental Psychology-General, 139*, 138-161. [a]

Le Pelley, M.E., Schmidt-Hansen, M., Harris, N.J., Lunter, C.M., & Morris, C.S. (2010). Disentangling the attentional deficit in schizophrenia: Pointers from schizotypy. *Psychiatry Research, 176*, 143-149. [a]

Le Pelley, M.E., Suret, M.B., & Beesley, T. (2009). Learned predictiveness effects in humans: A function of learning, performance, or both? *Journal of Experimental Psychology-Animal Behavior Processes, 35*, 312-327. [a]

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. [a]

Le Pelley, M.E., Vadillo, M., & 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. [a,c]

Lee, J., & Shomstein, S. (2014). Reward-based transfer from bottom-up to top-down search tasks. *Psychological Science, 25*, 466-475. [c]

Leon, S.P., Abad, M.J.F., & Rosas, J.M. (2011). Context-outcome associations mediate context-switch effects in a human predictive learning task. *Learning and Motivation, 42*, 84-98. [a]

Li, J., Schiller, D., Schoenbaum, G., Phelps, E.A., & Daw, N.D. (2011). Differential roles of human striatum and amygdala in associative learning. *Nature Neuroscience, 14*, 1250-1252. [a]

Li, W., Howard, J.D., Parrish, T.B., & Gottfried, J.A. (2008). Aversive learning enhances perceptual and cortical discrimination of indiscriminable odor cues. *Science, 319*, 1842-1845. [c]

Lipp, O.V., Siddle, D.A.T., & Dall, P.J. (1993). Effects of miscuing on Pavlovian conditioned responding and on probe reaction-time. *Australian Journal of Psychology, 45*, 161-167. [a]

Liu, E.H., Mercado, E., Church, B.A., & Orduna, I. (2008). The easy-to-hard effect in human (Homo sapiens) and rat (Rattus norvegicus) auditory identification. *Journal of Comparative Psychology, 122*, 132-145. [a]

Liu, P.P., & Luhmann, C.C. (2013). Evidence that a transient but cognitively demanding process underlies forward blocking. *Quarterly Journal of Experimental Psychology, 66*, 744-766. [a]

Livesey, E.J., Harris, I.M., & Harris, J.A. (2009). Attentional changes during implicit

   learning: Signal validity protects a target stimulus from the attentional blink. *Journal of*

   Experimental Psychology-Learning Memory and Cognition, 35, 408-422. [a]

Livesey, E.J., & McLaren, I.P.L. (2007). Elemental associability changes in human

   discrimination learning. *Journal of Experimental Psychology-Animal Behavior*

   Processes, 33, 148-159. [a]

Livesey, E.J., Thorwart, A., De Fina, N.L., & Harris, J.A. (2011). Comparing learned

   predictiveness effects within and across compound discriminations. *Journal of*

   Experimental Psychology-Animal Behavior Processes, 37, 446-465. [a]

Lochmann, T., & Wills, A.J. (2003). Predictive history in an allergy prediction task. In F.

   Schmalhofer, R.M. Young & G. Katz (Eds.), *Proceedings of EuroCogSci 03* (pp. 217-

   222). Mahwah, NJ: Lawrence Erlbaum Associates.

Lovejoy, E. (1968). *Attention in discrimination learning.* San Francisco, CA: Holden-Day.

Lubow, R.E. (1989). *Latent inhibition and Conditioned Attention Theory.* Cambridge:

   Cambridge University Press.

Lucke, S., Lachnit, H., Koenig, S., & Uengoer, M. (2013). The informational value of

   contexts affects context-dependent learning. *Learning & Behavior, 41*, 285-297. [a]

Lucke, S., Lachnit, H., Stuttgen, M.C., & Uengoer, M. (2014). The impact of context

   relevance during extinction learning. *Learning & Behavior, 42*, 256-269. [a]

Luque, D., Morís, J., Rushby, J.A., & Le Pelley, M.E. (2015). Goal-directed EEG activity

   evoked by discriminative stimuli in reinforcement learning. *Psychophysiology, 52*, 238-

   248. [a,c]

Macho, S. (1997). Effect of relevance shifts in category acquisition: A test of neural networks.

   Journal of Experimental Psychology-Learning Memory and Cognition, 23, 30-53. [b]

Mackintosh, N.J. (1975). A theory of attention: Variations in the associability of stimuli with

reinforcement. *Psychological Review, 82*, 276-298.

Mackintosh, N.J., & Little, L. (1969). Intradimensional and extradimensional shift learning by pigeons. *Psychonomic Science, 14*, 5-6.

MacLean, M.H., & Giesbrecht, B. (2015). Irrelevant reward and selection histories have different influences on task-relevant attentional selection. *Attention Perception & Psychophysics, 77*, 1515-1528. [c]

Maes, J.H.R., Damen, M.D.C., & Eling, P. (2004). More learned irrelevance than perseveration errors in rule shifting in healthy subjects. *Brain and Cognition, 54*, 201-211. [a]

Maes, J.H.R., Eling, P., Wezenberg, E., Vissers, C., & Kan, C.C. (2011). Attentional set shifting in autism spectrum disorder: Differentiating between the role of perseveration, learned irrelevance, and novelty processing. *Journal of Clinical and Experimental Neuropsychology, 33*, 210-217. [a]

Maes, J.H.R., Vich, J., & Eling, P. (2006). Learned irrelevance and response perseveration in a total change dimensional shift task. *Brain and Cognition, 62*, 74-79. [a]

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition, 22*, 657-672.

Maljkovic, V., & Nakayama, K. (1996). Priming of pop-out: II. The role of position. *Perception & Psychophysics, 58*, 977-991.

Maljkovic, V., & Nakayama, K. (2000). Priming of pop-out: III. A short-term implicit memory system beneficial for rapid target selection. *Visual Cognition, 7*, 571-595.

Marini, F., Marzi, T., & Viggiano, M.P. (2011). "Wanted!" The effects of reward on face recognition: Electrophysiological correlates. *Cognitive Affective & Behavioral Neuroscience, 11*, 627-643. [c]

Marissen, M.A.E., Franken, I.H.A., Waters, A.J., Blanken, P., van den Brink, W., & Hendriks,

V.M. (2006). Attentional bias predicts heroin relapse following treatment. *Addiction, 101*, 1306-1312.

Marx, S., & Einhauser, W. (2015). Reward modulates perception in binocular rivalry. *Journal of Vision, 15*, 13. [c]

Mas-Herrero, E., & Marco-Pallares, J. (2014). Frontal theta oscillatory activity is a common mechanism for the computation of unexpected outcomes and learning rate. *Journal of Cognitive Neuroscience, 26*, 447-458. [a]

Maunsell, J.H.R. (2004). Neuronal representations of cognitive state: Reward or attention? *Trends in Cognitive Sciences, 8*, 261-265. [d]

Medin, D.L., & Edelson, S.M. (1988). Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General, 117*, 68-85. [b]

Medin, D.L., & Robbins, D. (1971). Effects of frequency on transfer performance after successive discrimination training. *Journal of Experimental Psychology, 87*, 434-436. [b]

Miller, R.R., & Matzel, L.D. (1988). The comparator hypothesis: A response rule for the expression of associations. *The Psychology of Learning and Motivation, 22*, 51-92.

Mine, C., & Saiki, J. (2015). Task-irrelevant stimulus-reward association induces value-driven attentional capture. *Attention Perception & Psychophysics, 77*, 1896-1907. [c]

Mitchell, C., & Hall, G. (2014). Can theories of animal discrimination explain perceptual learning in humans? *Psychological Bulletin, 140*, 283-307.

Mitchell, C.J., Griffiths, O., Seetoo, J., & Lovibond, P.F. (2012). Attentional mechanisms in learned predictiveness. *Journal of Experimental Psychology-Animal Behavior Processes, 38*, 191-202. [a]

Mitchell, C.J., Harris, J.A., Westbrook, R.F., & Griffiths, O. (2008). Changes in cue associability across training in human causal learning. *Journal of Experimental Psychology-Animal Behavior Processes, 34*, 423-436. [a]

Morris, R., Griffiths, O., Le Pelley, M.E., & Weickert, T.W. (2013). Attention to irrelevant cues is related to positive symptoms in schizophrenia. *Schizophrenia Bulletin, 39*, 575-582. [a]

Munneke, J., Hoppenbrouwers, S.S., & Theeuwes, J. (2015). Reward can modulate attentional capture, independent of top-down set. *Attention Perception & Psychophysics, 77*, 2540-2548. [d]

Myers, C.E., Oliver, L.M., Warren, S.G., & Gluck, M.A. (2000). Stimulus exposure effects in human associative learning. *Quarterly Journal of Experimental Psychology Section B-Comparative and Physiological Psychology, 53*, 173-187. [b]

Myers, C.E., Warren, S., Brawn, C.M., McGlinchey-Berroth, R., Monti, L., & Gluck, M.A. (2000). Latent learning in medial temporal amnesia: Evidence for disrupted representational but preserved attentional processes. *Neuropsychology, 14*, 3-15. [a]

Nelson, J.B., Lamoureux, J.A., & Leon, S.P. (2013). Extinction arouses attention to the context in a behavioral suppression method with humans. *Journal of Experimental Psychology-Animal Behavior Processes, 39*, 99-105. [a]

Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*, 39-57.

Nosofsky, R.M., Gluck, M.A., Palmeri, T.J., McKinley, S.C., & Glauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of shepard, hovland, and jenkins (1961). *Memory & Cognition, 22*, 352-369. [a]

O'Brien, J.L., & Raymond, J.E. (2012). Learned predictiveness speeds visual processing. *Psychological Science, 23*, 359-363. [c]

Orosz, A.T., Feldon, J., Gal, G., Simon, A., & Cattapan-Ludewig, K. (2007). Repeated measurements of learned irrelevance by a novel within-subject paradigm in humans. *Behavioural Brain Research, 180*, 1-3. [b]

Orosz, A.T., Feldon, J., Gal, G., Simon, A.E., & Cattapan-Ludewig, K. (2008). Deficient associative learning in drug-naive first-episode schizophrenia: Results obtained using a new visual within-subjects learned irrelevance paradigm. *Behavioural Brain Research, 193*, 101-107. [b]

Orosz, A.T., Feldon, J., Simon, A.E., Hilti, L.M., Gruber, K., Yee, B.K., et al. (2011). Learned irrelevance and associative learning is attenuated in individuals at risk for psychosis but not in asymptomatic first-degree relatives of schizophrenia patients: Translational state markers of psychosis? *Schizophrenia Bulletin, 37*, 973-981. [a]

Owen, A.M., Roberts, A.C., Hodges, J.R., Summers, B.A., Polkey, C.E., & Robbins, T.W. (1993). Contrasting mechanisms of impaired attentional set-shifting in patients with frontal-lobe damage or parkinsons-disease. *Brain, 116*, 1159-1175. [b]

Owen, A.M., Roberts, A.C., Polkey, C.E., Sahakian, B.J., & Robbins, T.W. (1991). Extra-dimensional versus intra-dimensional set shifting performance following frontal-lobe excisions, temporal-lobe excisions or amygdalo-hippocampectomy in man. *Neuropsychologia, 29*, 993-1006. [b]

Packer, J.S., Siddle, D.A.T., & Tipp, C. (1989). Restoration of stimulus associability, electrodermal activity, and processing resource allocation. *Biological Psychology, 28*, 105-121. [a]

Pascucci, D., Mastropasqua, T., & Turatto, M. (2015). Monetary reward modulates task-irrelevant perceptual learning for invisible stimuli. *PLoS One, 10*, 12. [c]

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

Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O'Doherty, J.P. (2013). The neural representation of unexpected uncertainty during value-based decision making. *Neuron, 79*, 191-201. [a]

Pearce, J.M., George, D.N., & Redhead, E.S. (1998). The role of attention in the solution of

conditional discriminations. In N.A. Schmajuk & P.C. Holland (Eds.), *Occasion setting: Associative learning and cognition in animals*. Washington, D.C.: American Psychological Association.

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

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.

Pearson, D., Donkin, C., Tran, S.C., Most, S.B., & Le Pelley, M.E. (2015). Cognitive control and counterproductive oculomotor capture by reward-related stimuli. *Visual Cognition, 23*, 41-66. [c]

Pitt, M.A., & Myung, I.J. (2002). When a good fit can be bad. *Trends in Cognitive Sciences, 6*, 421-425.

Ploghaus, A., Tracey, I., Clare, S., Gati, J.S., Rawlins, J.N.P., & Matthews, P.M. (2000). Learning about pain: The neural substrate of the prediction error for aversive events. *Proceedings of the National Academy of Sciences of the United States of America, 97*, 9281-9286. [a]

Pool, E., Brosch, T., Delplanque, S., & Sander, D. (2014). Where is the chocolate? Rapid spatial orienting toward stimuli associated with primary rewards. *Cognition, 130*, 348-359. [c]

Pool, E., Delplanque, S., Porcherot, C., Jenkins, T., Cayeux, I., & Sander, D. (2014). Sweet reward increases implicit discrimination of similar odors. *Frontiers in Behavioral Neuroscience, 8*, 8. [c]

Posner, M.I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology,*

32, 3-25.

Posner, M.I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D.

Bouwhuis (Eds.), *Attention and Performance X* (pp. 531-556). Hillsdale, NJ: Erlbaum.

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.

Preuschoff, K., Hart, B.M., & Einhauser, W. (2011). Pupil dilation signals surprise: evidence

for noradrenaline's role in decision making. *Frontiers in Neuroscience, 5*, 12. [a]

Prevost, C., McCabe, J.A., Jessup, R.K., Bossaerts, P., & O'Doherty, J.P. (2011).

Differentiable contributions of human amygdalar subregions in the computations

underlying reward and avoidance learning. *European Journal of Neuroscience, 34*, 134-

145. [a]

Qi, S.Q., Zeng, Q.H., Ding, C., & Li, H. (2013). Neural correlates of reward-driven

attentional capture in visual search. *Brain Research, 1532*, 32-43. [c]

Raymond, J.E., & O'Brien, J.L. (2009). Selective visual attention and motivation: The

consequences of value learning in an attentional blink task. *Psychological Science, 20*,

981-988. [a,c]

Raymond, J.E., Shapiro, K.L., & Arnell, K.M. (1992). Temporary suppression of visual

processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology:

Human Perception and Performance, 18*

Rehder, B., & Hoffman, A.B. (2005a). Eyetracking and selective attention in category

learning. *Cognitive Psychology, 51*, 1-41. [b]

Rehder, B., & Hoffman, A.B. (2005b). Thirty-something categorization results explained:

Selective attention, eyetracking, and models of category learning. *Journal of

Experimental Psychology: Learning, Memory, and Cognition, 31*, 811-829. [b]

Reinhart, R.M.G., & Woodman, G.F. (2014). High stakes trigger the use of multiple

memories to enhance the control of attention. *Cerebral Cortex, 24*, 2022-2035. [c]

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: Appleton-Century-Crofts.

Robinson, T.E., & Berridge, K.C. (2001). Incentive-sensitization and addiction. *Addiction, 96*,

103-114.

Roper, Z.J.J., Vecera, S.P., & Vaidya, J.G. (2014). Value-driven attentional capture in

adolescence. *Psychological Science, 25*, 1987-1993. [c]

Rosas, J.M., & Callejas-Aguilera, J.E. (2006). Context switch effects on acquisition and

extinction in human predictive learning. *Journal of Experimental Psychology-Learning

Memory and Cognition, 32*, 461-474. [a]

Rothkirch, M., Ostendorf, F., Sax, A.L., & Sterzer, P. (2013). The influence of motivational

salience on saccade latencies. *Experimental Brain Research, 224*, 35-47. [c]

Rutherford, H.J.V., O'Brien, J.L., & Raymond, J.E. (2010). Value associations of irrelevant

stimuli modify rapid visual orienting. *Psychonomic Bulletin & Review, 17*, 536-542. [c]

Sali, A.W., Anderson, B.A., & Yantis, S. (2014). The role of reward prediction in the control

of attention. *Journal of Experimental Psychology-Human Perception and Performance,

40*, 1654-1664. [c]

Schmajuk, N.A., & Moore, J.W. (1985). Real-time attentional models for classical

conditioning and the hippocampus. *Physiological Psychology, 13*, 278-290.

Schmidt-Hansen, M., & Le Pelley, M.E. (2012). The positive symptoms of schizophrenia and

latent inhibition in humans and animals: Underpinned by the same process(es)?

Cognitive Neuropsychiatry, 17, 473-505.

Schoenmakers, T.M., de Bruin, M., Lux, I.F.M., Goertz, A.G., Van Kerkhof, D.H.A.T., & Wiers, R.W. (2010). Clinical effectiveness of attentional bias modification training in abstinent alcoholic patients. *Drug and Alcohol Dependence, 109*, 30-36.

Schutz, A.C., Trommershauser, J., & Gegenfurtner, K.R. (2012). Dynamic integration of information about salience and value for saccadic eye movements. *Proceedings of the National Academy of Sciences of the United States of America, 109*, 7547-7552. [c]

Schutzwohl, A. (1998). Surprise and schema strength. *Journal of Experimental Psychology-Learning Memory and Cognition, 24*, 1182-1199. [a]

Sewell, D.K., & Lewandowsky, S. (2012). Attention and working memory capacity: Insights from blocking, highlighting, and knowledge restructuring. *Journal of Experimental Psychology-General, 141*, 444-469. [a]

Shapiro, K.L., Caldwell, J., & Sorensen, R.E. (1997). Personal names and the attentional blink: A visual "cocktail party" effect. *Journal of Experimental Psychology: Human Perception and Performance, 23*, 504-514.

Shomstein, S., & Johnson, J. (2013). Shaping attention with reward: Effects of reward on space- and object-based selection. *Psychological Science, 24*, 2369-2378. [c]

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. [a]

Siddle, D.A.T. (1985). Effects of stimulus omission and stimulus change on dishabituation of the skin-conductance response. *Journal of Experimental Psychology-Learning Memory and Cognition, 11*, 206-216. [b]

Siddle, D.A.T., Booth, M.L., & Packer, J.S. (1987). Effects of stimulus preexposure on omission responding and omission-produced dishabituation of the human electrodermal response. *Quarterly Journal of Experimental Psychology Section B-Comparative and*

Physiological Psychology, 39, 339-363. [b]

Sigurjonsdottir, O., Bjornsson, A.S., Ludvigsdottir, S.J., & Kristjansson, A. (2015). Money talks in attention bias modification: Reward in a dot-probe task affects attentional biases. *Visual Cognition, 23*, 118-132. [c]

Slabosz, A., Lewis, S.J.G., Smigasiewicz, K., Szymura, B., Barker, R.A., & Owen, A.M. (2006). The role of learned irrelevance in attentional set-shifting impairments in parkinson's disease. *Neuropsychology, 20*, 578-588. [a]

Smith, S.D., Most, S.B., Newsome, L.A., & Zald, D.H. (2006). An emotion-induced attentional blink elicited by aversively conditioned stimuli. *Emotion, 6*, 523-527. [d]

Stankevich, B.A., & Geng, J.J. (2014). Reward associations and spatial probabilities produce additive effects on attentional selection. *Attention Perception & Psychophysics, 76*, 2315-2325. [c]

Stankevich, B.A., & Geng, J.J. (2015). The modulation of reward priority by top-down knowledge. *Visual Cognition, 23*, 206-228. [c]

Stormer, V., Eppinger, B., & Li, S.C. (2014). Reward speeds up and increases consistency of visual selective attention: A lifespan comparison. *Cognitive Affective & Behavioral Neuroscience, 14*, 659-671. [c]

Suret, M., & McLaren, I.P.L. (2003). Representation and discrimination on an artificial dimension. *Quarterly Journal of Experimental Psychology Section B-Comparative and Physiological Psychology, 56*, 30-42. [a]

Sutherland, N.S., & Mackintosh, N.J. (1971). *Mechanisms of animal discrimination learning.* New York: Academic Press.

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

Theeuwes, J., & Belopolsky, A.V. (2012). Reward grabs the eye: Oculomotor capture by

rewarding stimuli. *Vision Research, 74*, 80-85. [c]

Thorndike, E.L. (1911). *Animal Intelligence*. New York, NY: Macmillan.

Trabasso, T.R., & Bower, G.H. (1968). *Attention in learning: Theory and research*. New York: Wiley.

Trick, L., Hogarth, L., & Duka, T. (2011). Prediction and uncertainty in human Pavlovian to instrumental transfer. *Journal of Experimental Psychology-Learning Memory and Cognition, 37*, 757-765. [a]

Turner, D.C., Aitken, M.R.F., Shanks, D.R., Sahakian, B.J., Robbins, T.W., Schwarzbauer, C., et al. (2004). The role of the lateral frontal cortex in causal associative learning: Exploring preventative and super-learning. *Cerebral Cortex, 14*, 872-880. [a]

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

Uengoer, M., Lachnit, H., Lotz, A., Koenig, S., & Pearce, J.M. (2013). Contextual control of attentional allocation in human discrimination learning. *Journal of Experimental Psychology-Animal Behavior Processes, 39*, 56-66. [a]

Van Damme, S., Crombez, G., Hermans, D., Koster, E.H.W., & Eccleston, C. (2006). The role of extinction and reinstatement in attentional bias to threat: A conditioning approach. *Behaviour Research and Therapy, 44*, 1555-1563. [a]

Van der Molen, M.J.W., Van der Molen, M.W., Ridderinkhof, K.R., Hamel, B.C.J., Curfs, L.M.G., & Ramakers, G.J.A. (2012). Attentional set-shifting in fragile X syndrome. *Brain and Cognition, 78*, 206-217. [a]

Vanni-Mercier, G., Mauguiere, F., Isnard, J., & Dreher, J.C. (2009). The hippocampus codes the uncertainty of cue-outcome associations: An intracranial electrophysiological study in humans. *Journal of Neuroscience, 29*, 5287-5294. [a]

Wallis, G., Stokes, M.G., Arnold, C., & Nobre, A.C. (2015). Reward boosts working memory encoding over a brief temporal window. *Visual Cognition, 23*, 291-312. [c]

Wang, L.H., Duan, Y.Y., Theeuwes, J., & Zhou, X.L. (2014). Reward breaks through the inhibitory region around attentional focus. *Journal of Vision, 14*, 7. [c]

Wang, L.H., Yu, H.B., & Zhou, X.L. (2013). Interaction between value and perceptual salience in value-driven attentional capture. *Journal of Vision, 13*, 13. [c]

Waters, A.J., Shiffman, S., Sayette, M.A., Paty, J.A., Gwaltney, C.J., & Balabanis, M.H. (2003). Attentional bias predicts outcome in smoking cessation. *Health Psychology, 22*, 378-387.

Welham, A.K., & Wills, A.J. (2011). Unitization, similarity, and overt attention in categorization and exposure. *Memory & Cognition, 39*, 1518-1533. [a]

Wentura, D., Muller, P., & Rothermund, K. (2014). Attentional capture by evaluative stimuli: Gain- and loss-connoting colors boost the additional-singleton effect. *Psychonomic Bulletin & Review, 21*, 701-707. [c]

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

Wilbertz, G., van Slooten, J., & Sterzer, P. (2014). Reinforcement of perceptual inference: Reward and punishment alter conscious visual perception during binocular rivalry. *Frontiers in Psychology, 5*, 9. [c]

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. [a]

Wills, A.J., Lavric, A., Hemmings, Y., & Surrey, E. (2014). Attention, predictive learning, and the inverse base-rate effect: Evidence from event-related potentials. *Neuroimage, 87*, 61-71. [a]

Wills, A.J., & Pothos, E.M. (2012). On the adequacy of current empirical evaluations of formal models of categorization. *Psychological Bulletin, 138*, 102-125.

Winman, A., Wennerholm, P., Juslin, P., & Shanks, D.R. (2005). Evidence for rule-based processes in the inverse base-rate effect. *Quarterly Journal of Experimental Psychology Section a-Human Experimental Psychology, 58*, 789-815. [b]

Wolfe, J.M., Butcher, S.J., Lee, C., & Hyle, M. (2003). Changing your mind: On the contributions of top-down and bottom-up guidance in visual search for feature singletons. *Journal of Experimental Psychology: Human Perception and Performance, 29*, 483-502.

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

Yantis, S., & Jonides, J. (1984). Abrupt visual onsets and selective attention: Evidence from visual search. *Journal of Experimental Psychology: Human Perception and Performance, 10*, 601-621.

Yokoyama, T., Padmala, S., & Pessoa, L. (2015). Reward learning and negative emotion during rapid attentional competition. *Frontiers in Psychology, 6*, 10. [c]

Zeaman, D., & House, B.J. (1963). The role of attention in retardate discrimination learning. In N.R. Ellis (Ed.), *Handbook of mental deficiency: Psychological theory and research* (pp. 378-418). New York: McGraw-Hill.

Figure 1. Example sequence of events on a trial in the study by Livesey, Harris and Harris (2009). The initial asterisk signals the beginning of the trial. A series of letters is then presented rapidly (~100 ms per item). Two of these letters (the targets, T1 and T2) appear in red. The final mask (hash symbol) is immediately followed by a response cue – either the left or right circle lights up, and participants must press a corresponding button as quickly as possible. After making this response, participants are asked to report T1 and T2 in turn. Certain letters that appeared as target T2 predicted the correct left/right response that would be required by the response cue; participants were better at detecting these predictive target letters than other, nonpredictive target letters.

Figure 2. Example stimulus displays from the training phase (a) and test phase (b) of the study by Anderson, Laurent and Yantis (2011a). In the training phase, the target was a red or green circle; in the test phase, the target was a shape singleton (either a diamond among circles, as here, or a circle among diamonds). Participants responded to the orientation (horizontal or vertical) of the line segment contained in the target. During training, rapid correct responses received reward, with the amount of reward depending on the colour of the target circle. A high-value distractor colour typically predicted large reward; a low-value colour typically predicted small reward.

Figure 3. Simulation results using variants of Mackintosh's (1975) attentional theory of associative learning. Simulations comprised 100 trials on which cues A and X were together paired with an outcome (AX+), alternated with 100 trials on which X alone was presented without the outcome (X-). Thus A represents a reliable predictor of the outcome, while X represents a nonpredictive stimulus. Upper panels show the associative strength of A (V_A)

across training, and lower panels show attention to A (α_A). Blue lines show simulation results for a high-value outcome ($\lambda = 0.8$), and red lines show results for a low-value outcome ($\lambda = 0.3$). (a) Attention calculated based on a comparison of relative predictiveness (Equation [2]). Since A is the most predictive stimulus regardless of outcome magnitude, α_A increases to the upper limit of 1 in both cases. However, it approaches this limit more rapidly when outcome magnitude is large ($\lambda = 0.8$) than when it is small ($\lambda = 0.3$). So this model anticipates an influence of learned value on attention early in training, but not at asymptote. [Other parameters: $S = \theta = 0.2$]. (b) Attention determined by absolute associative strength (Equation [3]). As A develops associative strength, α_A increases for both $\lambda = 0.8$ and $\lambda = 0.3$. However, since attention is determined by associative strength, which is in turn limited by λ, asymptotic attention is greater when outcome magnitude is larger than when it is small. So this simpler model anticipates a persistent influence of learned value on attention. [Other parameter: $S = 0.3$]. (Figure taken from Le Pelley, Beesley & Griffiths, in press)
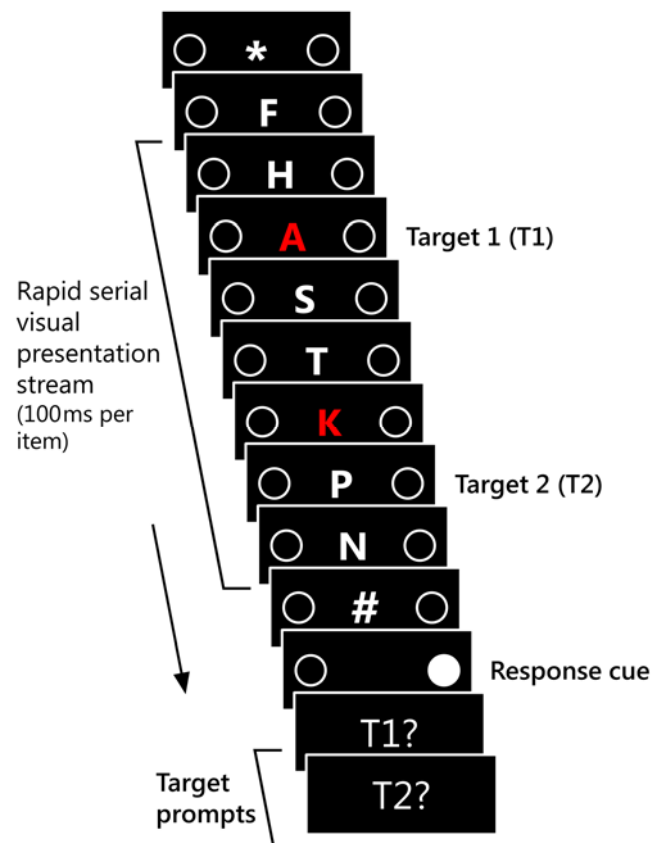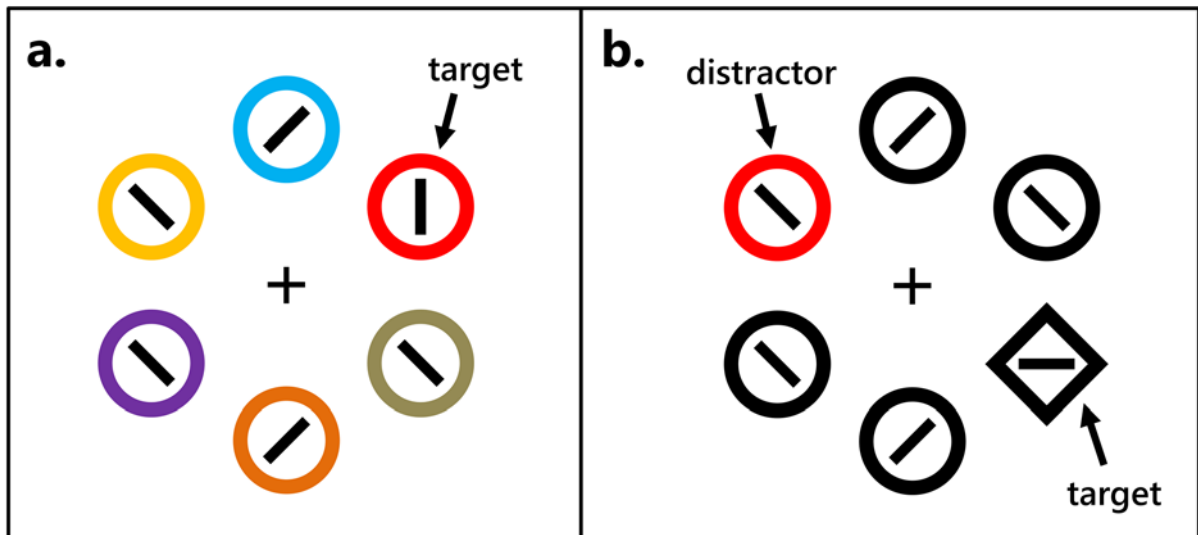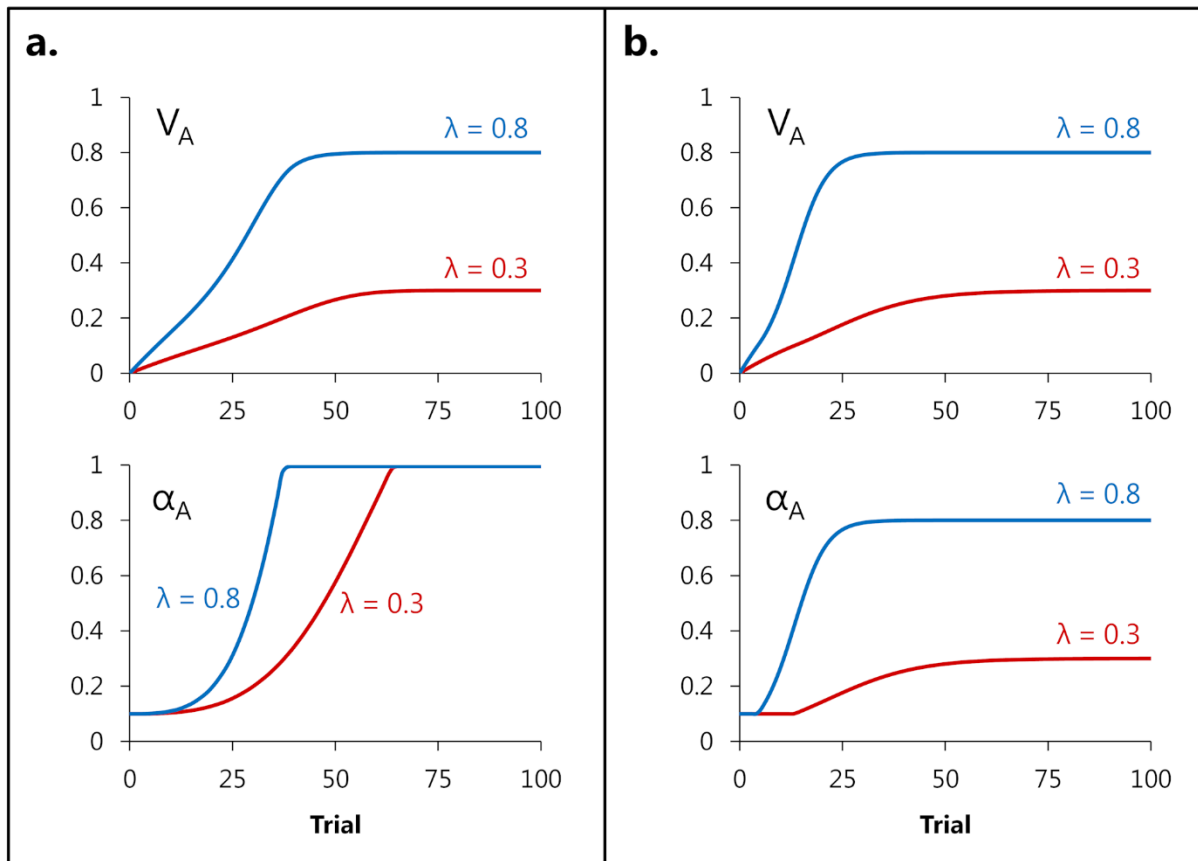
Figure 1

Figure 2

Figure 3