

# When Goals Conflict With Values: Counterproductive Attentional and Oculomotor Capture by Reward-Related Stimuli

Mike E. Le Pelley, Daniel Pearson, Oren Griffiths, and Tom Beesley  
University of New South Wales

Attention provides the gateway to cognition, by selecting certain stimuli for further analysis. Recent research demonstrates that whether a stimulus captures attention is not determined solely by its physical properties, but is malleable, being influenced by our previous experience of rewards obtained by attending to that stimulus. Here we show that this influence of reward learning on attention extends to *task-irrelevant* stimuli. In a visual search task, certain stimuli signaled the magnitude of available reward, but reward delivery was not contingent on responding to those stimuli. Indeed, any attentional capture by these critical distractor stimuli led to a reduction in the reward obtained. Nevertheless, distractors signaling large reward produced greater attentional and oculomotor capture than those signaling small reward. This counterproductive capture by task-irrelevant stimuli is important because it demonstrates how external reward structures can produce patterns of behavior that conflict with task demands, and similar processes may underlie problematic behavior directed toward real-world rewards.

*Keywords:* visual attention, reward learning, reinforcement learning, eye movements, attentional capture

Attention describes the cognitive mechanisms used to allocate mental resources to the processing of certain aspects of sensory input. For example, when driving we might use controlled, goal-directed attention to prioritize processing of events on the road ahead, and to ignore conversation from the backseat. But a sudden bang from the car's rear will capture our attention in an automatic, stimulus-driven fashion (Yantis, 2000). Beyond these goal-directed and stimulus-driven processes, research (as described below) has demonstrated that the extent to which stimuli capture attention is also influenced by learning about the significance of those stimuli (for recent reviews, see Awh, Belopolsky, & Theeuwes, 2012; Chelazzi, Perlato, Santandrea, & Libera, 2013). This possibility of an interaction between attention and learning—wherein stimuli with meaningful consequences might “stand out”—is not a new idea. 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” (p. 393). While the idea of derived attention has been around for some time, convincing empirical support for the concept has arisen only relatively recently, and has emerged in parallel from two sources.

## Predictiveness-Driven Attentional Capture

Researchers working in the associative learning and animal conditioning tradition have tended to focus on how attention is influenced by learning about the *predictiveness* of stimuli (also referred to as validity or informativeness). A predictive stimulus is one that provides information regarding other events that will occur, or actions that should be taken. For example, a green traffic light provides the information that pulling out into an intersection is safe, and a red light provides the information that it is unsafe; hence these are predictive stimuli. The color of the car in the neighboring lane does not provide any information regarding whether pulling out is safe, and so is a nonpredictive stimulus. A large body of work in both humans and nonhuman animals has examined the influence of previous experience of the predictiveness of a stimulus on the rate of subsequent learning about that stimulus (for reviews, see Le Pelley, 2004, 2010; Pearce & Mackintosh, 2010). In humans at least, the typical finding is that predictive stimuli are learned about more rapidly in future than are nonpredictive stimuli. This finding is consistent with the suggestion that attention is allocated to stimuli as a function of learning about their predictiveness, on the assumption that the rate of learning about a stimulus provides a measure of attention to that stimulus. Recent studies go further, by demonstrating that the learned predictiveness of a stimulus influences the extent to which that stimulus automatically captures our attention (Le Pelley, Vadillo, & Luque, 2013; Livesey, Harris, & Harris, 2009). For example, Le Pelley et al. (2013) gave participants an initial training phase in which certain stimuli (say, colored squares) predicted which of two buttons would be the correct response on each trial, whereas other stimuli (say, sets of oblique lines) provided no information regarding the correct response and hence were nonpredictive. After many trials of training on this task, participants moved on to a test phase that involved a variant of the *dot probe* procedure (MacLeod, Mathews, & Tata, 1986). On each trial of this test phase, a target (a small white triangle) appeared either in a location cued by a stimulus that had been predictive in the training

---

This article was published Online First November 24, 2014.

Mike E. Le Pelley, School of Psychology, University of New South Wales; Daniel Pearson, School of Psychology, University of New South Wales; Oren Griffiths, School of Psychology, University of New South Wales; Tom Beesley, School of Psychology, University of New South Wales.

This work was supported by Australian Research Council Grant FT100100260. We thank Branka Spehar for help in creating the stimuli, and Tom Whitford for helpful comments on the manuscript.

Correspondence concerning this article should be addressed to Mike E. Le Pelley, School of Psychology, University of New South Wales, Sydney NSW 2052, Australia. E-mail: [m.lepelley@unsw.edu.au](mailto:m.lepelley@unsw.edu.au)

phase (a colored square), or in a location cued by a stimulus that had been nonpredictive (a set of oblique lines). Participants were required to press the spacebar as quickly as possible when the target appeared. Importantly, across these test trials, the location in which the target appeared was independent of the type of cueing stimulus (predictive or nonpredictive). Hence, participants gained no advantage by orienting their attention to the location of either type of cueing stimulus in advance of the target appearing. Indeed, participants were explicitly informed that in order to respond to the target as quickly as possible, their best strategy was to ignore the initially presented stimuli. Despite this instruction, responses were significantly faster when the target appeared in the location of the predictive stimulus than the location of the nonpredictive stimulus.

The implication is that the predictive stimulus captured participants' spatial attention and hence responses to events occurring in that location were initiated earlier (cf. Posner, 1980). This attentional capture occurred even though (a) it was not required by the task, (b) it was not adaptive with regard to that task (since target location was independent of the type of cueing stimulus), and (c) there was little time for participants to consciously process and respond to the cueing stimuli on each test trial. Le Pelley et al. (2013) demonstrated that providing more time for participants to consciously process the stimuli—by increasing the SOA on test trials to 1,000 ms—significantly *weakened* the influence of predictiveness on dot probe responding. This suggests that the pattern observed at short SOA was not a result of conscious, controlled processing but instead reflected a rapid and automatic effect of predictiveness on attentional capture. A long SOA then provides time for participants to use controlled processes to correct for the automatic attentional capture caused by presentation of the stimuli, returning attention to the center of the display (cf. Klauer, Roßnagel, & Musch, 1997).

### Value-Driven Attentional Capture

As noted above, predictiveness refers to the extent to which a stimulus provides information regarding subsequent events. In a parallel stream, researchers working within the perception and cognition traditions have demonstrated that attentional capture by a stimulus is also influenced by the value of those subsequent events (Anderson, Laurent, & Yantis, 2011a, 2011b; Anderson & Yantis, 2012; Libera & Chelazzi, 2009; Kiss, Driver, & Eimer, 2009; Rutherford, O'Brien, & Raymond, 2010; Theeuwes & Belopolsky, 2012). Specifically, if attending to a stimulus is consistently paired with a high-value reward, then that stimulus becomes more likely to capture attention than an equally salient stimulus paired with low-value reward.

The clearest demonstration of this *value-driven attentional capture* comes from studies of visual search (Anderson et al., 2011a, 2011b). On each trial of an initial training phase, a target (either a red or green circle) was presented together with distractors (similar circles rendered in other colors). Participants were required to respond according to the orientation of a line segment contained in the target circle. During this training phase, a particular participant might receive a large reward for making a rapid response when the target was, say, a red circle, and a small reward for making a rapid response when the target was a green circle. Hence for this participant, red was the "high-value color," and green was the "low-value color." After extensive training on this task—typically over

1,000 trials—participants moved on to a test phase, in which the target on each trial was now defined by shape (say, a diamond among circles; the diamond was never red or green). Critically, during the test phase, people were slower to respond to this diamond target if one of the nontarget circles was rendered in the high-value color than if it was rendered in the low-value color; that is, a nontarget in the high-value color was more distracting. This attentional capture by a stimulus previously paired with large reward occurred even though attending to color conflicted with current task demands in the test phase (respond to the diamond), suggesting that capture was automatic and involuntary.

### Derived Attention, Addiction, and Psychosis

The interaction between attention and learning that is implicated in predictiveness-driven and value-driven capture is notable because it demonstrates that processing of sensory input is not a fixed function of physical salience, but is instead malleable and based on our previous experiences. This may bring adaptive advantages by improving and speeding detection of informative and/or reward-related stimuli. But it may also create problems. For example, many drugs of abuse produce potent neural reward signals (Dayan, 2009; Hyman, 2005; Robinson & Berridge, 2001). Involuntary attentional capture by stimuli associated with these drug rewards (such as drug paraphernalia, or people and locations associated with drug supply) is known to predict relapse in recovering addicts (Cox, Hogan, Kristian, & Race, 2002; Marissen et al., 2006; Waters et al., 2003). Interestingly, Anderson, Faulkner, Rilee, Yantis, and Marvel (2013) have recently used the procedure described above to demonstrate that value-driven attentional capture is magnified in drug addicts, consistent with the idea of a link between addiction and attention to reward-related stimuli. It has also been argued that the psychotic symptoms of schizophrenia reflect a dysfunction of the relationship between reward learning and attention (Frank, 2008; Kapur, 2003; Morris, Griffiths, Le Pelley, & Weickert, 2013). On this account, abnormal reward learning results in patients attending to stimuli that would normally be ignored, and ignoring those to which they should attend. This in turn produces unusual sensory experiences (hallucinations) and cognitive efforts to make sense of those unusual experiences (delusions). A better understanding of the mechanisms underlying derived attention in humans has the potential to shed light on aspects of mental disorder, including addiction and psychosis, that implicate a dysfunction of these mechanisms.

### Is Task-Relevance Essential for Learned Attentional Capture?

All previous studies that have demonstrated learned attentional capture in humans (be that predictiveness- or value-driven) have one feature in common. In all cases, during the training phase that was used to establish differences in the predictiveness or value of the stimuli, these stimuli were *task-relevant* for participants; that is, they were the stimuli that participants were required to identify in order to perform the task. For example, in Le Pelley et al.'s (2013) study of predictiveness-driven capture, the predictive stimuli (colored squares in the example given above) were predictive precisely because they defined the correct response on each training trial. Hence during the training phase, participants needed to identify these stimuli in order to

make the correct response. In the initial training phase of Anderson et al.'s (2011a, 2011b) studies of value-driven capture, the stimuli that predicted reward magnitude (red and green circles) defined the targets to which participants were required to direct their responses. In each case, this raises the possibility that attentional capture by similar stimuli in the subsequent test phase was simply a carryover of an overlearned attentional orienting response to these stimuli, previously established in the training phase. For example, the training phase of Anderson et al.'s task essentially constitutes extensive training that orienting rapidly to (say) a red circle leads to a large reward. Reinforcement learning might therefore promote the extent to which red stimuli automatically elicit such a rapid orienting response in the future. If this automatic attentional orienting response persisted into the test phase, this would explain why red stimuli were more distracting during this test phase.<sup>1</sup>

In the real world, however, stimuli that signal rewards are not always the goals that produce those rewards. For example, an addict may typically take drugs in a particular room. This room signals the drug's rewarding effect, but has no instrumental relationship with achieving that reward: entering the room does not itself elicit a drug reward, and the drug would have a similar rewarding effect if ingested elsewhere. In this sense, the room is task-irrelevant with respect to the goal of achieving drug reward. This raises the question of whether such task-irrelevant stimuli can nevertheless support attentional capture. To investigate this possibility, the studies of value-driven capture reported here used training in which the critical stimuli were *never* task-relevant for participants. On every trial, participants searched for a diamond-shaped target among circles. One of these nontarget circles (the distractor) could be colored, and the distractor's color predicted the reward available on that trial. But crucially, reward was obtained by responding to the diamond, not to the distractor circle. Hence the stimuli predicting reward were not those to which participants were required to direct their responses or attention. In fact, we ensured that attending to reward-predictive distractors would, if anything, hinder participants' performance and hence the payoff they achieved. If value-driven attentional capture by distractors were nevertheless observed, it would imply the existence of an automatic attentional process that prioritizes reward-related stimuli even when capture by those stimuli is, and always has been, directly contrary to participants' goals.

A finding of value-driven capture by task-irrelevant stimuli would shed important light on the nature of the information that underlies learned attentional capture, because it would suggest that it is the simple correlation of stimuli with reward, rather than their functional role in obtaining that reward, that determines capture. Rephrased in the language of associative learning theory, this would imply that derived attention is a product of Pavlovian conditioning (i.e., learning about the extent to which a stimulus *signals* reward), rather than instrumental conditioning (i.e., learning about the relationship between a *response* and the reward that it produces). As noted above, the color of the distractor in our experiments signaled the size of the reward available on each trial. Hence, the high-value color was a Pavlovian signal of large reward, and the low-value color was a Pavlovian signal of small reward. If value-driven capture arises because signals of large reward are more likely to capture attention than signals of small reward, then we would expect greater capture by the high-value color than the low-value color. But as noted above, participants were not rewarded for *responding to* (or orienting attention

toward) the stimulus that was presented in the value-related color since this stimulus was, by definition, a nontarget. Hence there was no potential for instrumental conditioning to promote value-driven capture. In fact, since our experiments were arranged so that orienting toward the distractor resulted in loss of reward, participants would receive larger rewards when they successfully suppressed attention to this distractor (or oriented away from it). In particular, the largest rewards would occur when participants successfully refrained from orienting toward the high-value color. Consequently, the instrumental relationships in force in this experiment would, if anything, promote suppression of attentional orienting to the high-value color relative to the low-value color. This would result in less distraction by the high-value color; the opposite pattern to that anticipated by the Pavlovian account.

Below we describe three experiments investigating the influence of reward learning on attentional capture by task-irrelevant stimuli. In Experiments 1 and 2, the primary measure of capture is response time, on the assumption that attentional capture by the distractor will result in slowing of responses to the target. Experiment 3 goes on to use a gaze-contingent eye-tracking procedure to measure the extent to which reward-related distractors capture eye gaze.

## Experiment 1

### Participants

Participants in Experiment 1 were 27 University of New South Wales (UNSW) students, who received course credit and also performance-based payment ( $M = \$19.92$  AUD).

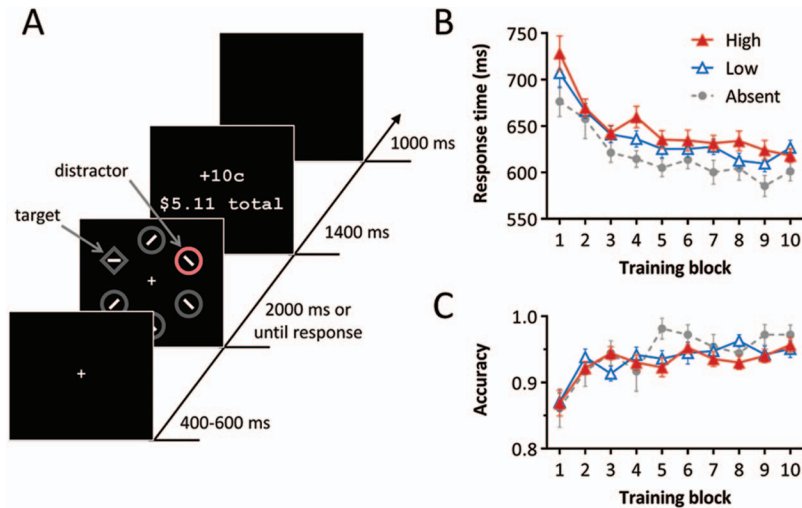
### Apparatus

Participants were tested individually using a standard PC with a 23-in. monitor ( $1,920 \times 1,080$  resolution, refresh rate 120 Hz), positioned  $\sim 60$  cm from the participant. Stimulus presentation was controlled by MATLAB using Psychophysics Toolbox extensions (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997).

### Stimuli

The experiment used a variant of the additional singleton paradigm (Theeuwes, 1991, 1992). Each trial consisted of a fixation display, a search display, and a feedback display (Figure 1a). All stimuli were presented on a black background. The fixation display consisted of a white cross (subtending 0.5 degrees of visual angle;

<sup>1</sup> It is worth mentioning here a recent study by Anderson et al. (2012). This used the same training procedure as Anderson et al. (2011a, 2011b) in which, say, red circles were the targets that predicted large reward and green circles were the targets predicting small reward. In this case, however, the subsequent test phase showed greater attentional capture by *letters* rendered in red rather than letters in green. Strictly, then, Anderson et al. (2012) demonstrated value-driven capture by stimuli (colored letters) that had not been task-relevant during the training phase—since they had not appeared during this training phase. Crucially, however, these stimuli shared the critical, task-relevant feature that defined targets in the training phase, namely red or green color. Hence, like the earlier results of Anderson et al. (2011a, 2011b), these findings reflect value-driven attentional capture by a *feature* (red or green color) that had been task-relevant. In contrast, in the current studies, the feature that predicted reward magnitude (again, color) was never task-relevant.



**Figure 1.** Sequence of trial events, Experiment 1 (a). Participants responded to the orientation of the line segment (horizontal or vertical) within the diamond (target). One of the nontarget circles could be a color singleton distractor (shown in red in online version, light gray in print version). Fast, correct responses to the target received monetary reward, depending on the distractor color. A high-value distractor color reliably predicted large reward; a low-value reliably predicted small reward; if no color singleton was present in the display (distractor-absent trial), then large and small reward were equally likely. Mean response time (b) and mean accuracy (proportion correct; c) across the 10 training blocks of Experiment 1, for trials with a high-value distractor, a low-value distractor, and distractor-absent trials. Error bars show within-subjects standard error of the mean (see Cousineau, 2005). Critically, response times were significantly slower on trials with a high-value distractor than trials with a low-value distractor, but no more accurate. See the online article for the color version of this figure.

dva) presented centrally. The search display comprised the fixation cross surrounded by six shapes ( $2.3 \times 2.3$  dva) positioned at equal intervals around an imaginary circle with diameter 10.1 dva (with the first position directly above the fixation cross). Five of these shapes were circles, and one was a diamond. The diamond and four of the circles were rendered in gray. The remaining circle (the distractor) was rendered either in red, blue, green, or the same shade of gray as the other shapes (Commission Internationale d'Elairage  $x, y$  chromaticity coordinates of .595/.360 for red, .160/.116 for blue, .300/.611 for green, .304/.377 for gray). The values of red, blue, and green had similar luminance ( $\sim 42.5$  cd/m<sup>2</sup>), which was higher than that of the gray (36.5 cd/m<sup>2</sup>). The target contained a white line segment (length 0.76 dva) oriented either vertically or horizontally. Each nontarget contained a similar line segment tilted 45° randomly to the left or right.

## Design

Two colors from the set of red, green and blue were randomly . . . assigned to act as *high-value* and *low-value* colors for each participant.<sup>2</sup> The experiment comprised 10 training blocks. Each block contained 18 trials with a distractor rendered in the high-value color, 18 trials with a distractor in the low-value color, and 4 distractor-absent trials on which there was no color singleton in the display. Trials occurred in random order.

Correct responses to the target that were faster than the participant's latency limit (see Procedure) were followed by feedback indicating reward. If that trial had a high-value distractor, reward was large (10¢); if it had a low-value distractor, reward was small

(1¢). If there was no distractor, the reward was equally likely to be small or large. Correct responses that were slower than the participant's latency limit received no reward, though participants were told how much they could have won. Errors resulted in loss of the amount that had been scheduled as a reward on that trial. Target location, distractor location, and target line segment orientation (vertical or horizontal) were randomly determined on each trial.

## Procedure

The single session began with a practice phase of 20 trials, with a yellow distractor on each and no reward feedback. The first two practice trials were discarded; the upper quartile of response latencies for the remaining correctly responded-to practice trials defined the latency limit for each participant. If response accuracy over practice trials was below 75%, participants repeated the practice phase. The mean latency limit was 964 ms (standard error of the mean [SEM] = 42 ms)

Instructions informed participants that for subsequent trials, they would earn a reward only for correct responses that were faster than their latency limit (this latency limit was given to them in

<sup>2</sup> The remaining color was assigned to act as a *rare* color. Each training block contained four trials with the distractor rendered in this rare color; reward magnitude was equally likely to be small or large on these trials. However, data from trials with the rare distractor are uninformative with regard to the critical issue of value-driven capture addressed in this article (indeed, the rare color was omitted from Experiments 2 and 3 for exactly this reason). Hence, for the sake of brevity we do not present or analyze these data here.



milliseconds), that such responses would earn either 1¢ or 10¢ “depending on the shapes presented on the trial,” and that errors would result in loss of the corresponding amount. Participants were not informed of the relationship between distractor color and reward magnitude.

Each trial began with presentation of the fixation display for a random period of 400, 500, or 600 ms. The search display then appeared until a response was made or the trial timed out (after 2 s). Participants responded to the orientation of the line segment in the diamond by pressing the *C* and *M* keys for horizontally and vertically oriented lines, respectively. For fast correct responses (or errors), the feedback display appeared for 1,400 ms (or 2,000 ms), showing the reward earned (or lost) on the previous trial, and total earnings so far. For correct responses slower than the latency limit, feedback displayed participants’ response time (RT) and the amount they would have received for a faster response. Intertrial interval was 1,000 ms. Participants took a short break every two blocks.

## Data Analysis

The first two trials, and the first two trials after each break, were discarded. Timeouts (0.06% of all trials) and trials with RTs below 150 ms (0%) were also discarded. RTs were then analyzed for correct responses only.

## Results

Figure 1b and 1c show RTs and accuracy across training. RTs were analyzed using a 3 (distractor type: high-value, low-value, absent)  $\times$  10 (block) analysis of variance (ANOVA). This revealed a main effect of block,  $F(9, 234) = 10.8, p < .001, \eta_p^2 = .29$ , with mean RT falling as participants became more practiced at the task. There was also a significant main effect of distractor type,  $F(2, 52) = 23.1, p < .001, \eta_p^2 = .47$ . The interaction was not significant,  $F(18, 468) = .70, p = .82, \eta_p^2 = .03$ .

Planned pairwise *t* tests, averaging across training blocks, were used to further analyze the main effect of distractor type. Each type of colored distractor slowed RT relative to distractor-absent trials—high-value versus absent:  $t(26) = 6.29, p < .001, d = 1.21$ , 95% CI of difference ( $CI_{diff}$ ) [19.8, 39.1]; low-value versus absent:  $t(26) = 5.05, p < .001, d = .97, CI_{diff}$  [11.6, 27.4]. Critically, RT on trials with the high-value distractor ( $M = 647$  ms) was significantly greater than on trials with the low-value distractor ( $M = 637$  ms),  $t(26) = 2.15, p = .041, d = .41, CI_{diff}$  [.42, 19.4]. Slower RTs on trials with the high-value distractor meant participants were more likely to miss out on the large rewards available on these trials: correct RTs were slower than the latency limit (and hence reward was omitted) on 4.3% of high-value distractor trials versus 3.6% of low-value distractor trials. This difference was significant,  $t(26) = 1.94, p = .032$  (one-tailed, since direction is anticipated by the RT difference),  $d = .37, CI_{diff}$  [.00, .013].

For the accuracy data, the omnibus 3  $\times$  10 ANOVA revealed a main effect of block,  $F(9, 234) = 8.98, p < .001, \eta_p^2 = .26$ , with accuracy increasing across training, but no significant main effect of distractor type,  $F(2, 52) = 1.55, p = .22, \eta_p^2 = .06$ , or interaction,  $F(18, 468) = .92, p = .55, \eta_p^2 = .03$ . Notably, averaging across blocks, accuracy on trials with the high-value distractor ( $M = 93.0\%$ ) was not significantly different from that

for trials with the low-value distractor ( $M = 93.6\%$ ),  $t(26) = 1.12, p = .27, d = .22, CI_{diff}$  [−.013, .004]. This suggests that the critical difference in RT observed on these trials did not reflect a speed–accuracy tradeoff.

## Discussion

The finding that each type of colored distractor slowed RT relative to distractor-absent trials replicates the well-established finding that search for a target defined by shape is slowed by the presence of a color singleton distractor (Theeuwes, 1992, 1994). The distractors were physically salient, because they were the only colored stimuli in the display. The implication is that these salient stimuli capture attention in a relatively automatic fashion (automatic, because attending to the distractors will impede the participant’s goal of responding to the target as fast as possible).

More importantly, responses were significantly slower (but no more accurate) for trials with a high-value distractor compared with a low-value distractor. This suggests that the high-value distractor was more likely to capture attention than the low-value distractor, even though these distractors were always task-irrelevant. That is, attentional capture by the distractors was modulated by the magnitude of the reward that they signaled. The implication, then, is that value-driven attentional capture is a consequence of Pavlovian conditioning (based on the extent to which distractor color provides a *signal* of reward value), rather than instrumental conditioning (based on the value of the reward that is produced by *responding* [i.e., attentional orienting] to the color). Notably, the greater capture by high-value distractors meant that participants missed more high-value than low-value rewards. Hence, this increased capture by the task-irrelevant high-value distractor was counterproductive to participants’ goal of maximizing their payoff.

## Experiment 2

Notably, however, mean correct RT in Experiment 1 (636 ms) was considerably shorter than the mean latency limit (964 ms). Given that the mean RT difference between high- and low-value distractor trials was numerically small ( $\sim 10$  ms), this meant that the greater attentional capture by the high-value target would rarely impact on the reward received. In other words, slower RTs on high-value distractor trials did not always result in reward omission; the relationship between RT and reward in Experiment 1 was relatively indirect. Experiment 2 therefore implemented a direct relationship between RT and reward, so that slower responses necessarily resulted in reduced reward. We also investigated whether the counterproductive effect observed in Experiment 1 would persist across extended training, or whether participants would eventually come to show a more adaptive pattern of faster responses on high-value trials. This possibility was based on data from conditioning studies with rats (e.g., Holland, 1979), which demonstrate that some overt behaviors (e.g., approaching a food magazine during an auditory stimulus that terminates in delivery of food) that are initially driven by Pavlovian conditioning can, with extended training, come under the control of instrumental conditioning; though other behaviors (e.g., the startle response elicited by the same auditory stimulus) remain under Pavlovian control even with extensive training. Finally,

Experiment 2 investigated the relationship between value-driven capture and explicit awareness of the color–reward contingencies.

## Method

**Participants.** Twenty-four people took part. Four were students participating for course credit, and 20 were recruited from the UNSW community in exchange for \$45. All participants also received performance-related payment ( $M = \$78.21$  AUD).

**Apparatus and stimuli.** Apparatus and stimuli were as for Experiment 1, except Experiment 2 had no green stimuli.

**Design.** The experiment comprised 36 training blocks of 48 trials, giving 1,728 trials in total. Each block contained 20 trials with the distractor in the high-value color, 20 trials in the low-value color, and 8 distractor-absent trials. These trials occurred in random order. Assignment of red and blue to high- and low-value colors was counterbalanced across participants.

Correct responses with RTs slower than 1,000 ms earned no reward. For faster correct responses, reward (in cents) was calculated as  $(1,000 - RT) \times 0.002 \times \text{bonus\_multiplier}$ , rounded to the nearest 0.01¢. For high-value distractor trials, *bonus\_multiplier* was always 10; for low-value distractor trials, *bonus\_multiplier* was always 1; for distractor-absent trials, *bonus\_multiplier* was equally often 1 and 10. Errors resulted in loss of the corresponding amount. Other aspects of design were as for Experiment 1.

**Procedure.** Participants completed three 1-hr sessions on consecutive days. Each session began with 10 practice trials featuring a yellow distractor, with no reward feedback. This was followed by 12 training blocks (giving 576 training trials per session). Participants were told they would earn 0.2¢ for every 100 ms their RT was below 1,000 ms, that on “10 × Bonus” trials this reward would be multiplied by 10, and that errors would result in loss of the corresponding amount. Instructions made no reference to the relationship between bonus trials and distractor color. Trials were as for Experiment 1, except that feedback was presented for 2,000 ms after correct responses and 2,500 ms after errors. On nonbonus trials, the feedback display showed the amount earned and total earnings; on bonus trials this was accompanied by a yellow box labeled “10 × Bonus trial!” (bonus trials were not signaled explicitly until after participants had made their response).

After the final training session, participants were told that bonus trials had been determined by the color in the display; that when certain colors appeared it would be a bonus trial, and when other colors appeared it would not be. In a final awareness test, they were then shown a red circle and a blue circle in random order, and for each selected whether it would be a bonus trial or not when this color of circle appeared in the display.

**Data analysis.** Data were analyzed as for Experiment 1. Timeouts (0.01% of all trials) and trials with RTs below 150 ms (0.03%) were discarded.

## Results

Figure 2A shows RTs across training. These were analyzed using a 3 (distractor type: high-value, low-value, absent) × 36 (block) ANOVA. A main effect of block,  $F(35, 805) = 67.1$ ,  $p < .001$ ,  $\eta_p^2 = .74$ , reflected the reduction in mean RT across blocks. There was a significant main effect of distractor type,  $F(2, 46) = 45.9$ ,  $p < .001$ ,  $\eta_p^2 = .67$ , and a Distractor Type × Block interaction,  $F(70, 1610) = 2.04$ ,  $p < .001$ ,  $\eta_p^2 = .08$ .

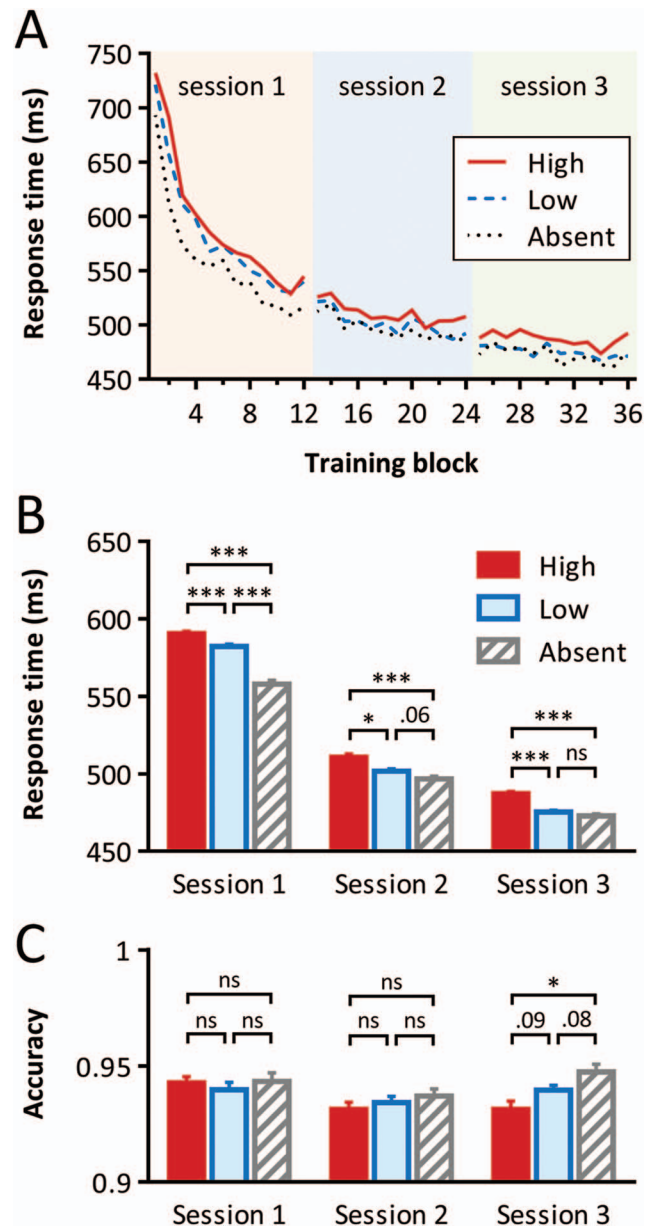


Figure 2. Mean response time across the 12 training blocks of each of the 3 sessions of Experiment 2, for trials with a high-value distractor, a low-value distractor, and distractor-absent trials (a). Point markers and error bars have been omitted for the sake of clarity. Mean response time (b) and accuracy (proportion correct; c) for each session of Experiment 2, for high-value, low-value, and distractor-absent trials. Error bars show within-subjects standard error of the mean (Cousineau, 2005). Response times were significantly slower on trials with a high-value distractor than trials with a low-value distractor in all sessions. See the online article for the color version of this figure. \*  $p < .05$ . \*\*  $p < .01$ . \*\*\*  $p < .001$ .

A follow-up  $2 \times 36$  ANOVA used only the data for high- and low-value distractors. The main effect of distractor type was significant,  $F(1, 23) = 25.7$ ,  $p < .001$ ,  $\eta_p^2 = .53$ , with slower responses on trials with high-value distractors than low-value distractors. The Distractor Type × Block interaction was not significant,  $F(35, 805) = 1.06$ ,  $p = .37$ ,  $\eta_p^2 = .044$ .

Figure 2b shows mean RTs for each session. Critically, the high-value distractor slowed RT relative to the low-value distractor in each session—Session 1:  $t(23) = 4.00, p < .001, d = .82, CI_{diff} [3.98, 12.5]$ ; Session 2:  $t(23) = 2.78, p = .011, d = .57, CI_{diff} [2.31, 15.8]$ ; Session 3:  $t(23) = 5.48, p < .001, d = 1.12, CI_{diff} [7.39, 16.4]$ . This did not reflect a speed–accuracy tradeoff. Figure 2c shows that mean accuracy was similar for high- and low-value distractors in Sessions 1 and 2,  $t < 1$  for each, and Session 3 showed a trend toward lower accuracy for high- than low-value distractors,  $t(23) = 1.75, p = .093, d = .36, CI_{diff} [-.002, .018]$ .

It is clear from Figure 2b that the RT difference between high-value and low-value distractor trials did not decrease with extended training; the numerical difference was roughly constant across Sessions 1 to 3 ( $M = 8.26, 9.05, \text{ and } 11.87$  ms, respectively), despite the reduction in baseline RTs across sessions. Indeed, when expressed as a proportion of baseline RT, the difference between high- and low-value trials is significantly greater in Session 3 than in Session 1,  $t(23) = 2.11, p = .046, d = .43, CI_{diff} [.000, .022]$  (and as noted above, this RT difference is accompanied by a trend-level difference in accuracy in Session 3 but not in Session 1).

In the final awareness test, seven participants showed no awareness of the color–reward contingencies, incorrectly selecting that the low-value color signaled bonus trials and that the high-value color did not. Across all trials, these “unaware” participants still showed significantly slower RTs on high-value distractor trials ( $M = 527$  ms) than on low-value distractor trials ( $M = 518$  ms):  $t(6) = 4.24, p = .005, d = 1.60, CI_{diff} [3.53, 13.2]$ ; alternatively Wilcoxon’s  $T(7) = 0, p = .016$ . For the remaining 17 participants, whose responses in the awareness test were consistent with them having become aware of the veridical relationships between the different colors and reward levels, RTs were also significantly slower on high-value distractor trials ( $M = 531$  ms) than on low-value distractor trials ( $M = 520$  ms):  $t(16) = 4.01, p = .001, d = .97, CI_{diff} [4.85, 15.7]$ . The magnitude of the value-driven capture effect (i.e., the difference in RT between high- and low-value distractor trials) did not differ significantly for “aware” and “unaware” participants,  $t(22) = .46, p = .65, d = .23, CI_{diff} [-10.7, 6.85]$ .

## Discussion

Experiment 2 replicated the key finding of Experiment 1: RTs were slower on trials with a high-value distractor rather than a low-value distractor. Once again, this suggests greater attentional capture by the high-value distractor. Importantly, because RTs directly determined reward in Experiment 2, this enhanced capture by the high-value distractor was directly counterproductive, because it meant that participants earned less on high-value trials than would otherwise have been the case. This counterproductive pattern persisted across extensive training: even after 2,000 trials, attentional capture remained under the control of the Pavlovian signal-value of the colors. Interestingly, this pattern was observed in participants who were unaware of the color–reward contingencies, suggesting that awareness is not necessary for value-driven attentional capture by task-irrelevant stimuli.

## Experiment 3

It is well-established that stimuli that capture attention also tend to capture eye movements, known as *oculomotor capture* (e.g., Anderson & Yantis, 2012; Ludwig & Gilchrist, 2002; Theeuwes & Belopolsky, 2012; Theeuwes, Kramer, Hahn, & Irwin, 1998; Theeuwes, Kramer, Hahn, Irwin, & Zelinsky, 1999). Experiment 3 used a gaze-contingent eye-tracking procedure to determine whether reward learning about task-irrelevant distractors influenced the extent to which those distractors elicited oculomotor capture. Critically, online monitoring of eye movements as participants performed the search task allowed us to ensure that oculomotor capture by distractors was *never* rewarded.

## Method

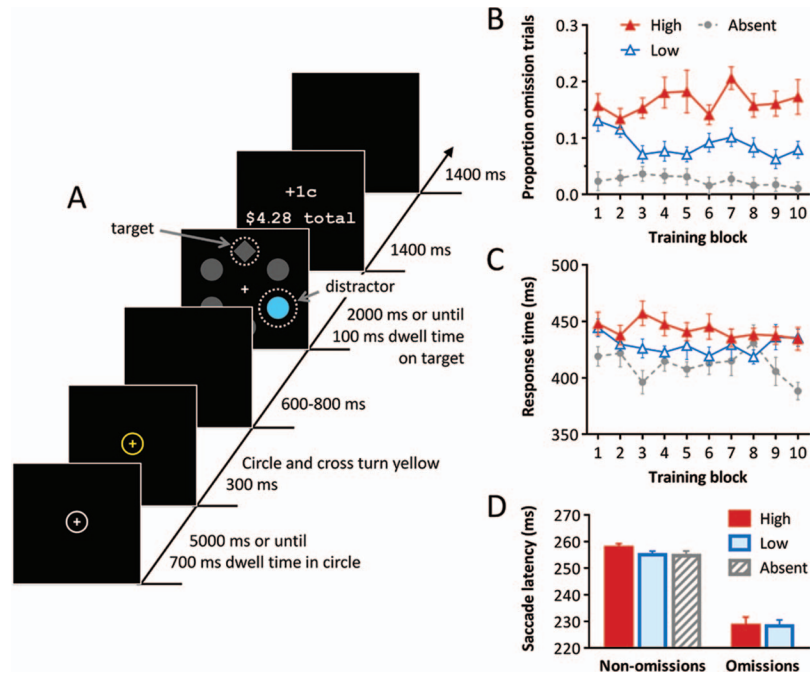
**Participants.** Twenty-four UNSW students participated for course credit and received performance-related payment ( $M = \$20.42$  AUD).

**Apparatus.** Experiment 3 used a Tobii TX300 eye-tracker, with 300 Hz temporal and  $0.15^\circ$  spatial resolution, mounted on a 23-in. widescreen monitor ( $1,920 \times 1,080$  resolution, refresh rate 60 Hz). Participants’ heads were positioned in a chinrest 60 cm from the screen. As participants performed the task, the program controlling stimulus presentation requested data from the eye-tracker every 10 ms. Because the eye-tracker recorded gaze location at 300 Hz (i.e., one recording every 3.3 ms), the 10 ms sample taken by the experiment program would typically contain three recordings of gaze location. Participants’ current gaze location was defined as the average of the locations contained in the most recent 10 ms sample. Invalid recordings (when the eye tracker failed to detect a gaze location) were not included in this average; if the current 10 ms sample contained no valid recordings, then the gaze location from the previous 10 ms sample was used instead. The eye-tracker was calibrated using a 5-point procedure prior to the practice phase, prior to the first training block, and after 6 training blocks.

**Stimuli.** Each trial consisted of a fixation display, a search display, and a feedback display (Figure 3a). The fixation display was a white cross surrounded by a white circle (diameter 3.0 dva). The search display was as for Experiment 2, except that: (a) all shapes were filled, (b) there were no line segments, (c) the fixation cross was absent, and (d) the shade of gray was darker than in previous experiments (luminance  $\sim 32$  cd/m<sup>2</sup>). The feedback display showed the reward earned on the previous trial and total earnings.

**Design.** The training phase comprised 10 blocks, which were structured as for Experiment 2. Target location and distractor location were randomly determined on each trial, with the constraint that the distractor could never appear adjacent to the target. On each trial, a small circular region of interest (ROI) with diameter 3.5 dva was defined around the diamond target; a larger ROI (diameter 5.1 dva) was defined around the distractor. A response was registered when participants had accumulated 100 ms of dwell time inside the target ROI. Responses with RTs slower than 600 ms earned no reward. If any gaze fell inside the distractor ROI prior to a response being registered, even for a single 10-ms period, the trial was recorded as an *omission trial* and no reward was delivered. On distractor-absent trials, one of the gray circles (that was not adjacent to the target) was chosen at random; gaze





**Figure 3.** Sequence of trial events, Experiment 3 (a). Participants responded by moving their eyes to the diamond target. One of the nontarget circles could be a color singleton distractor (shown in blue in online version, light gray in print version). Dotted lines (not visible to participants) indicate the region of interest around the target and distractor within which eye gaze was defined as falling on the corresponding stimulus. Fast, correct responses received monetary reward, depending on the distractor color. A high-value distractor color reliably predicted large reward; a low-value color reliably predicted small reward; on distractor-absent trials, large and small rewards were equally likely. If any gaze fell within the distractor region of interest (or, on distractor-absent trials, an equivalent region of interest positioned around a randomly chosen circle), the trial was deemed an omission trial, and no reward was delivered. Mean proportion of omission trials (b) and mean response times (c) across the 10 training blocks of Experiment 3, for high-value, low-value, and distractor-absent trials. Reward was more likely to be omitted, and response times were slower, on trials with the high-value distractor than trials with the low-value distractor. Mean saccade latencies on omission and non-omission trials, averaged across training blocks (d). Saccade latencies were generally slower for non-omission trials than omission trials. Latencies were also slower for non-omission trials featuring a high-value distractor than featuring a low-value distractor. All error bars show within-subjects standard error of the mean (Cousineau, 2005). See the online article for the color version of this figure.

falling inside an ROI around the selected gray circle caused an omission trial in exactly the same way as if the selected circle had been a distractor.<sup>3</sup>

If RT was faster than 600 ms and no gaze was registered in the distractor ROI, then a reward was delivered: 10¢ if the high-value distractor was present, 1¢ if the low-value distractor was present, and an equal likelihood of 10¢ or 1¢ on distractor-absent trials. Notably, in Experiment 3, participants never lost money (unlike in Experiments 1 and 2 where participants lost money for erroneous responses). Other design aspects were as for Experiment 2, except that in Experiment 3 the distractor never appeared adjacent to the target.

**Procedure.** The single session began with 8 practice trials with a yellow distractor and no rewards. Participants were then told that on subsequent trials they would earn either 0¢, 1¢, or 10¢, depending on “how fast and accurately you move your eyes to the diamond.”

Each trial began with the presentation of the fixation display. Participants’ gaze location was superimposed on this display as a

small yellow dot. Once participants had recorded 700 ms dwell time inside the circle surrounding the fixation cross, or if 5 s had passed, the cross and circle turned yellow, and the dot marking gaze location disappeared. After 300 ms, the screen blanked, and after a random interval of 600, 700, or 800 ms, the search display appeared. The trial terminated when a response was registered (see Design), or after 2 s (timeout). The feedback display then appeared for 1,400 ms. Intertrial interval was 1,400 ms.

**Data analysis.** As for previous experiments, the first two trials, and the first two trials after each break, were discarded.

<sup>3</sup> Allowing for omissions on distractor-absent trials is useful, because it permits a valid test of the influence of stimulus salience on oculomotor capture, by comparing the rate of omissions on trials featuring a salient distractor with the rate on distractor-absent trials. This comparison controls for causes of omission trials that are not related to distractor salience (e.g., inaccuracy in the recording of gaze location, random eye movements by the participant, etc.), since these will be equal on trials with a salient distractor and distractor-absent trials.



Timeouts (1.3% of all trials) were also discarded. Finally, we also excluded all trials on which valid gaze location was registered in less than 25% of the 10-ms samples between presentation of the search display and registering of a response (1.4% of all trials). For remaining trials, averaging across participants, valid gaze location was registered in 97.8% (SEM 0.8%) of samples, suggesting very high fidelity of the gaze data on these trials.

For analysis of saccade latencies, we used the raw data from the eye-tracker (sampled at 300 Hz, rather than the 100 Hz used for gaze-contingent calculations). For these analyses, in addition to the exclusions described above, we further excluded all trials on which no eye gaze was recorded within 5.1 dva (100 pixels) of the fixation point during the first 80 ms after presentation of the search display. Saccade latency was then found by identifying the first point at which five consecutive gaze samples lay more than 5.1 dva from the fixation point. Saccades faster than 80 ms were excluded from further analyses. The extra exclusions described in this paragraph resulted in loss of an additional 6.1% of trials.

## Results

Our primary measure in Experiment 3 was the proportion of omission trials across training, shown in Figure 3b. These data were analyzed using a 3 (distractor type: high-value, low-value, absent)  $\times$  10 (block) ANOVA. The main effect of block was not significant,  $F(9, 207) = 1.25, p = .26$ . There was a significant main effect of distractor type,  $F(2, 46) = 24.5, p < .001, \eta_p^2 = .52$ , and the Distractor Type  $\times$  Block interaction approached significance,  $F(18, 414) = 1.52, p = .079, \eta_p^2 = .062$ .

On the basis of this trend toward a change in the effect of the distractor type on proportion of omissions across the course of training, smaller ANOVAs were used to further analyze these data for the different distractor types (as opposed to  $t$  tests collapsing across training blocks). A 2 (distractor type)  $\times$  10 (block) ANOVA, using only the data for high-value and low-value distractors, revealed a main effect of distractor type,  $F(1, 23) = 11.4, p = .003, \eta_p^2 = .33$ ; that is, high-value distractors produced significantly more omission trials than low-value distractors (16.4% vs. 8.9%, collapsing across blocks). The Distractor Type  $\times$  Block interaction approached significance,  $F(9, 207) = 1.89, p = .055, \eta_p^2 = .076$ , with the difference in proportion of omission trials for high-value and low-value distractors tending to increase as training progressed (see Figure 3b). Similar ANOVAs comparing the data for high-value versus distractor-absent trials, and low-value versus distractor-absent, revealed a significant main effect of distractor type in each case: for high-value versus absent,  $F(1, 23) = 33.8, p < .001, \eta_p^2 = .59$ ; for low-value versus absent,  $F(1, 23) = 36.5, p < .001, \eta_p^2 = .61$ . Unsurprisingly, then, color singleton distractors generally produced more oculomotor capture than if no color singleton was present in the display; just 2.4% of distractor-absent trials produced omissions, collapsing across blocks. For the comparison of high-value and distractor-absent trials, the Distractor Type  $\times$  Block interaction was nonsignificant,  $F(9, 207) = .90, p = .52, \eta_p^2 = .04$ . For the comparison of low-value and distractor-absent trials, this Distractor Type  $\times$  Block Interaction approached significance,  $F(9, 207) = 1.87, p = .058, \eta_p^2 = .075$ . Figure 3b shows that the difference between low-value and distractor-absent trials in proportion of omissions reduced as training progressed.

Similar patterns were seen in RTs (Figure 3c; as noted earlier, a response was registered in this task when participants had accumulated 100 ms of dwell time inside the target ROI). For these data,  $3 \times 10$  ANOVA revealed a significant main effect of distractor type,  $F(2, 46) = 24.1, p < .001, \eta_p^2 = .51$ , but no main effect of block,  $F(9, 207) = .47, p = .89, \eta_p^2 = .02$ , or interaction,  $F(18, 414) = 1.22, p = .24, \eta_p^2 = .05$ . Follow-up  $t$  tests, averaging across training blocks, revealed that RTs were slower with high-value distractors ( $M = 442$  ms) than low-value distractors ( $M = 429$  ms,  $t(23) = 4.00, p = .001, d = .82, CI_{diff} [6.31, 19.80]$ ), and fastest on distractor-absent trials: for high-value versus absent,  $t(23) = 5.68, p < .001, d = 1.16, CI_{diff} [20.3, 43.5]$ ; for low-value versus absent,  $t(23) = 3.94, p = .001, d = .80, CI_{diff} [8.93, 28.70]$ .

Figure 3d shows saccade latencies on omission trials, and non-omission trials (i.e., trials on which participants did not look at the distractor), averaged across training blocks. Saccade latencies for distractor-absent omission trials are not shown—even though omissions could occur on these trials (see Design)—because there were so few of these trials (8 of 24 participants registered zero trials in this category, so mean saccade latencies could not be calculated for these participants). Saccade latency was generally shorter on omission trials than nonomission trials: this was true for trials with a high-value distractor,  $t(23) = 6.91, p < .001, d = 1.41, CI_{diff} [.021, .038]$ , and with a low-value distractor,  $t(23) = 8.72, p < .001, d = 1.78, CI_{diff} [.020, .033]$ . For nonomission trials, saccade latency was significantly longer for high-value than low-value trials,  $t(23) = 2.29, p = .031, d = .47, CI_{diff} [.0003, .0051]$ . While latency was numerically shortest for distractor-absent trials, the relevant differences failed to reach significance: for high-value versus absent,  $t(23) = 1.47, p = .15, d = .30, CI_{diff} [-.0012, .0072]$ ; for low-value versus absent,  $t(23) = .17, p = .87, d = .03, CI_{diff} [-.0035, .0041]$ . For omission trials, saccade latency did not differ significantly between high- and low-value trials,  $t(23) = .04, p = .97, d = .007, CI_{diff} [-.0088, .0091]$ .

Finally, the duration of total dwell time on the distractor on omission trials did not differ significantly between trials with high-value distractors ( $M = 112$  ms, SEM = 7.5 ms) and low-value distractors ( $M = 110$  ms, SEM = 8.2 ms),  $t(23) = .30, p = .77, d = .06, CI_{diff} [-.015, .020]$ .

## Discussion

Experiment 3 replicated the value-related effect on RTs observed in Experiments 1 and 2. Moreover, high-value distractors produced greater oculomotor capture than low-value distractors. This was counterproductive because if oculomotor capture occurred, reward was omitted. Experiment 3 thus provides an interesting example of reward learning promoting an oculomotor response that was never rewarded.

Mean saccade latency in Experiment 3 was generally shorter on omission trials (i.e., trials on which participants looked at the distractor before looking at the target) than on nonomission trials (i.e., trials on which participants did not look at the distractor prior to looking at the target). We interpret this to suggest that the salient distractor had a tendency to elicit rapid oculomotor capture in a stimulus-driven fashion, but that participants could use goal-directed (controlled) processes to overcome this tendency to make an initial saccade toward the distractor. The longer saccade latency on nonomission trials then reflects the cost of engaging this con-

trolled process. On this account, one would also expect longer saccade latencies on nonomission trials for distractor-present than distractor-absent displays. While the mean latencies showed this pattern numerically, the relevant differences did not reach significance. This could, however, be a consequence of the relatively low number of distractor-absent trials (80 over the whole experiment, as compared with 400 distractor-present trials), such that the mean latency on distractor-absent trials is likely to be a relatively poor estimate of the population mean.

Most interestingly, saccade latencies on nonomission trials featuring a high-value distractor were significantly longer than on nonomission trials featuring a low-value distractor. This implies that greater cognitive effort was required to suppress the tendency to saccade toward the high-value distractor than the low-value distractor.

On omission trials, saccade latency was not significantly different for high-value and low-value trials. This could be taken to suggest that—while the high-value distractor was more *likely* to produce oculomotor capture (Figure 3b)—when either distractor actually produced capture, it did so with the same degree of “force.” Also, the mean dwell time on high-value and low-value distractors on omission trials was not significantly different. This could be taken to suggest that, after capture, the ease of attentional *disengagement* did not depend on distractor value. However, the interpretations advanced in this paragraph are speculative. This is because both rely on null results; in particular, because the number of omission trials for each participant was relatively small, the experiment may well have lacked sensitivity to detect differences on these measures. Future studies could use more extensive training to generate more omission trials, which would allow these questions to be studied in greater depth.

It is perhaps worth noting that the omission contingency implemented in Experiment 3 (wherein if participants looked at the distractor, the reward was omitted) means that participants must have learned the signal-value of the distractor colors (high-value color signals high-value reward and low-value color signals low-value reward) on trials on which they did not look at the distractor. That is, participants must have encoded the presence of a particular distractor color in the array using peripheral vision, and this supported learning about the relationship between the presence of that color and the reward value obtained on that trial.

### General Discussion

Considered in the most general terms, our findings replicate recent demonstrations that involuntary attentional and oculomotor capture in visual search is influenced by reward learning (Anderson et al., 2011a, 2011b, 2012; Anderson & Yantis, 2012; Theeuwes & Belopolsky, 2012). More importantly, however—and unlike all previous studies of derived attentional capture in humans—the current experiments demonstrate value-driven attentional capture by stimuli that were never task-relevant for participants. That is, participants were never required to direct their responses (or attention) toward these stimuli. Indeed, if participants *did* direct attention toward the critical distractor stimuli, their payoff was reduced. Nevertheless, we still found evidence of greater capture by stimuli that predicted high-value rewards than those predicting low-value rewards. Moreover, in Experiment 2, this maladaptive

pattern of capture was observed in participants who were seemingly unaware of the stimulus–reward relationships.

Unlike in previous studies, then, the value-driven capture observed in the current experiments cannot be a hangover from participants’ previous experience of being rewarded for directing attention to the critical stimuli. The clearest evidence for this comes from Experiment 3, where the oculomotor capture produced by the distractor stimuli was never rewarded at any point in the experiment. These findings therefore imply that the crucial determinant of capture is not instrumental learning about the reward value produced by orienting attention to a stimulus (*response-value*). Instead capture seems dependent on Pavlovian learning about the reward value signaled by the presence of a stimulus (*signal-value*). Specifically, our findings suggest that signals of large reward become more likely to capture attention than signals of small reward.

The difference in behavior on high-value versus low-value trials must reflect learning about signal-value, since (across participants) this was the only difference between distractor stimuli. However, the size of this difference did not interact significantly with training block in any of the experiments (although the relevant interaction for the proportion of omission trials in Experiment 3 approached significance). While this null result may simply reflect noise in the block-by-block data, it suggests that the influence of reward learning on attentional capture developed early and did not change greatly over the course of training. Notably, Experiment 2 demonstrated that the maladaptive pattern of greater capture by high-value distractors persisted over extended training. Even with extensive experience, capture did not come under the control of response-value; that is, participants did not come to suppress attention to the high-value distractor (which would have increased their payoff). The implication is that value-driven attentional capture remains an effect of Pavlovian conditioning, and that instrumental learning about the relationship between orienting attention to distractors and the resulting reduction in payoff never comes to exert a comparable influence.

At this point, we should acknowledge a subtle issue of interpretation, and a caveat to some of the arguments set out above. In all of the current experiments, the value-predictive distractors were physically salient, since they were the only colored stimuli in the display. It is well-established that salient stimuli such as these “pop out,” and capture covert attention automatically (Theeuwes, 1992, 1994). Thus we might expect that covert attention will shift to the distractors on a certain proportion of trials due to their physical salience; we term this a distractor-shift. Now let us focus on the procedure of Experiment 1, in which participants were rewarded if their response to the target was faster than a certain criterion value (the latency limit). Suppose that, after a distractor-shift, participants are still generally able to respond to the target faster than the latency limit. Hence, the distractor-shift is followed by reward; more specifically, a distractor-shift to a high-value distractor is followed by high reward, and a distractor-shift to a low-value distractor is followed by low reward. Thus distractor-shifts to the different distractor types are differentially rewarded, even though shifting attention to the distractor plays no causal role in achieving that reward. This provides circumstances under which instrumental conditioning could differentially promote distractor-shifts to the high-value distractor in the future (in the terminology of learning theory, this would be an example of superstitious conditioning; cf.

Skinner, 1948). A similar argument could apply to the current Experiment 2 if, after a distractor-shift, participants were still able to respond to the target fast enough that the rewards achieved on high-value trials were (on average) greater than those on low-value trials.

The procedure of Experiment 3 mitigates against this line of reasoning to an extent. In this experiment, the overt attentional behavior that was measured (shifting eye gaze to the distractor) was *never* rewarded at any point in the experiment, since making such a shift would produce omission of reward. Consequently, it is less straightforward to see how the value-driven modulation of oculomotor capture observed in this study could be driven by instrumental conditioning. However, one possibility remains. As noted above, the distractors in Experiment 3 were physically salient and hence likely to produce automatic capture of covert attention. Suppose that there are a subset of trials on which (a) participants shift covert attention toward the distractor, *and* (b) they do not look at the distractor, *and* (c) their response to the target is still fast enough to fall within the 600-ms deadline for earning a reward. Under these circumstances, covert distractor-shifts could then be subject to superstitious instrumental conditioning as described in the previous paragraph. If we then assume that the greater likelihood of covert distractor-shifts to the high-value distractor occasionally translates into making overt (oculomotor) shifts to this distractor, then this could account for the observed pattern of more omission trials for high-value than low-value distractors. On this account, then, the oculomotor bias is a product of an instrumentally conditioned bias in covert attention. However, it is important to remember that, if participants *do* make an oculomotor shift to the distractor, then they lose a reward that would otherwise have been earned. Hence if instrumental conditioning of covert attentional shifts to the distractor builds up to such a level that it produces an oculomotor shift to this distractor, then the resulting loss of reward would immediately drive instrumental conditioning to reduce the tendency of making such oculomotor shifts in future. Essentially, the “covert bias produces oculomotor bias” argument is subject to a negative feedback loop, which reduces the likelihood that it is responsible for the effects observed in Experiment 3.

More generally, previous research has demonstrated that learning about rewards modulates attentional capture by task-relevant stimuli, both when those stimuli are physically salient (Anderson et al., 2011a), and when they are not (Anderson et al., 2011b). The current experiments demonstrate that learning about rewards also modulates attentional capture by task-irrelevant stimuli, in the case in which they are physically salient. However, the physical salience of these stimuli opens up the possibility (at least in theory) that instrumental conditioning might contribute to the value-driven modulation of attention observed in these experiments. It is noteworthy, then, that a very recent study by Failing and Theeuwes (2014) has demonstrated value-driven capture by task-irrelevant distractors that are not distinguished by their physical salience. This experiment used a procedure based on the current Experiment 1, but crucially each of the outline shapes in the search display was uniquely colored (see Figure 4). Participants were required to report the letter (S or P) inside the shape singleton (the diamond in Figure 4) as rapidly as possible. As in the current experiments, for half of participants, the presence of a red distractor in the display signaled that a fast response to the target would receive high




Figure 4. Example search display from the study by Failing and Theeuwes (2014). Participants were required to report the letter (S or P) inside the shape singleton (the diamond in this example) as rapidly as possible. Each of the outline shapes in the display was uniquely colored (shown in different colors in online version, and different grayscale tones in print version). For half of participants, the presence of a red distractor in the display signaled that a fast response to the target would result in high reward and a blue distractor signaled low reward; for the other half of participants, this was reversed. See the online article for the color version of this figure.

reward, and a blue distractor signaled low reward; for the other half of participants this was reversed. As these critical value-predictive distractors were no longer color singletons (since all stimuli were colored), they would not be expected to elicit attentional capture based on their physical salience. Nevertheless, Failing and Theeuwes still found slower responses to the target on trials with a high-value distractor than on trials with a low-value distractor, implying that the high-value distractor became more likely to capture attention than the low-value distractor.

Although the current experiments demonstrate that reward learning can *modulate* attentional capture by salient stimuli, the data of Failing and Theeuwes (2014) suggest that reward learning is sufficient to *produce* attentional capture even by nonsalient stimuli. And since the study by Failing and Theeuwes used nonsalient stimuli, it rules out the account via instrumental conditioning advanced in the previous paragraphs. Taken together, our findings and those of Failing and Theeuwes therefore strongly suggest that value-driven capture by task-irrelevant stimuli is a product of Pavlovian, not instrumental, conditioning.

Finally in this section, we note an interesting complementary relationship between previous studies of value-driven capture and the current findings. Previous studies demonstrated value-driven capture by stimuli that had previously been task-relevant, but this capture was observed in a test phase in which these stimuli no longer predicted reward (e.g., Anderson et al., 2011a, 2011b). The current studies demonstrate value-driven capture by task-irrelevant stimuli, during a training phase in which rewards were provided throughout. The implication, then, is that the influence of value-



learning on attentional capture is quite general. Pairing with high-value reward increases the likelihood of capture by a stimulus that has been task-relevant but no longer predicts reward (Anderson et al., 2011a, 2011b), and by a stimulus that predicts reward but is task-irrelevant (current data and Failing & Theeuwes, 2014).

### A Previous Study of Capture by Distractors

One previous study by Della Libera and Chelazzi (2009) examined the influence of reward learning on attention to distractors in humans. In a complicated procedure, when critical stimuli appeared as distractors, they signaled reward magnitude (with 80% validity). Evidence from Libera and Chelazzi's Experiment 1 suggested that this training led to reduced capture by distractors that signaled large reward compared to small reward. This is the opposite of our findings, and suggests response-value was the critical variable in their case. The reason for this discrepancy remains unclear; however, we note that: (a) The effect for distractors in Libera and Chelazzi's Experiment 1 occurred on only one of two response measures, at  $p = .04$ , and did not replicate in Experiment 2. Our effect replicated across Experiments 1–3 with medium-to-large effect sizes, in both RTs and oculomotor capture. (b) Libera and Chelazzi had no consistent distinction between targets and distractors: a given stimulus acted as a target on some trials and as a distractor on others, but signaled reward magnitude only when it appeared in one of these roles. Thus all stimuli in their study were task-relevant on 50% of their appearances. In our experiments, colored stimuli only ever appeared as distractors, and so were task-irrelevant throughout. (c) Attentional capture by distractors in Libera and Chelazzi's procedure had no influence on rewards and hence was not maladaptive; in our experiments, capture resulted in reduced reward.

Leaving aside for a moment the specific issue of the task-relevance or task-irrelevance of stimuli, we would argue that our task has advantages over previous techniques (e.g., Anderson et al., 2011a, 2011b, 2012; Libera & Chelazzi, 2009; Theeuwes & Belopolsky, 2012) as a general procedure for investigating value-driven attentional capture in humans. All of these previous procedures require a lengthy training period during which differences in value are established, before a shorter, unrewarded test phase in which value-related differences in attentional capture might be observed. Since test phase trials are unrewarded, any value-related effects are liable to dissipate as reward learning extinguishes (e.g., see Anderson et al., 2011a). In contrast, our procedure involves only a single phase. Every trial is both a training trial (on which reward learning can occur) and a test trial (on which the effects of that learning on capture can be measured). Consequently, this procedure is more efficient, allows the influence of value-driven capture to be observed over prolonged training, and allows the development of this capture to be tracked online on a session-by-session or block-by-block basis.

### Neural Loci of Value-Driven Attentional Capture

The findings reported here are somewhat similar to those of a study by Peck, Jangraw, Suzuki, Efem, and Gottlieb (2009) using monkeys. On each trial of that study, a peripheral visual reward cue (RC) predicted whether the trial outcome would be juice reward (RC+) or no reward (RC−). However, to achieve this

outcome, monkeys were required to make a saccade to a target whose location was independent of the RC. Even though RCs had no operant role, the RC+ became more likely to attract attention and the RC− to repel attention (measured using eye tracking). This suggests that, as in the current experiments, attention was under the control of learning about the signal-value of the RC rather than its response-value. One difference is that in Peck et al.'s (2009) study, the independence of RC and target location meant that on 50% of trials, the target would appear in the same location as the RC. Hence, monkeys would often be rewarded for making a saccade toward a location in which the RC+ had recently appeared. In the current experiments, the distractor and target on a given trial never appeared in the same location, thus producing a clearer distinction between signal-value and response-value. Nevertheless, these findings suggest an interesting parallel between value-driven attention in humans and nonhuman animals. Using single-unit recording, Peck et al. showed that attentional modulation in their task was encoded in posterior parietal cortex, specifically in the lateral intraparietal area.

This latter finding is particularly interesting, because of its relation to previous data from humans. In a study in humans using task-relevant stimuli, Kiss et al. (2009) used electroencephalography to demonstrate that the magnitude of the reward that participants received for responding to target stimuli modulated event-related potential (ERP) signatures of attentional selection elicited by those stimuli. Specifically, the N2-posterior-contralateral (N2pc) ERP component occurred earlier, and had greater magnitude, for targets rendered in a high-value color than targets in a low-value color. The N2pc is an early, lateralized component emerging around 200 ms after display onset, and extensive study of singleton visual search has identified it as an important correlate of visual target selection (see Eimer, 1996; Woodman & Luck, 1999). Importantly, neural source analyses based on magnetoencephalography highlight the posterior parietal cortex as contributing to the N2pc induced by task-relevant items in visual search (e.g., see Hopf et al., 2000). Thus, we have two studies implicating the posterior parietal cortex in value-driven attentional capture: Kiss et al.'s (2009) study (in humans) using task-relevant stimuli, and Peck et al.'s (2009) study (in monkeys) using task-irrelevant stimuli. Although the species and procedures used in these studies are clearly very different, these findings are at least consistent with the possibility that the same brain regions—and potentially the same neural mechanisms—underlie value-driven attentional capture by both task-relevant and task-irrelevant stimuli. Future work in humans using the procedures developed in this article will probe this question more directly, by comparing neural signatures of value-driven capture by task-relevant and task-irrelevant stimuli within a single species, and with a single procedure.

### Sign-Tracking and Goal-Tracking of Attention

The contrast between signal- and response-value of stimuli that is made by the current experiments relates to animal conditioning studies distinguishing between *sign-tracking* and *goal-tracking* behavior (Boakes, 1977). Suppose a neutral cue (e.g., insertion of a lever into a conditioning chamber) is repeatedly paired with food. A rat showing goal-tracking behavior learns to approach the location of food delivery when the lever is inserted. In contrast, a rat showing sign-tracking behavior will approach the lever itself and

grasp, lick, or gnaw it, indicating that this signal of reward has become attractive in its own right. Sign-tracking thus involves the neutral cue taking on “incentive salience” and hence coming to exert control over behavior, even though it has no operant value in producing reward. The parallel with the current procedure is clear. Specifically, our data indicate that reward learning promotes attentional capture by the distractor (signal) more than by the target (goal). These findings therefore provide evidence of a sign-tracking process modulating attentional capture in humans.

Notably, it has been argued that the attribution of salience to predictive cues that is implicated in sign-tracking, and the subsequent ability of those cues to gain inordinate control over behavior, bears striking similarities to prominent symptoms of drug abuse (Flagel, Akil, & Robinson, 2009; Tomie, Grimes, & Pohorecky, 2008; Uslaner, Acerbo, Jones, & Robinson, 2006). In support of this, animal studies suggest that individual differences in the tendency to engage in sign-tracking versus goal-tracking behavior may confer vulnerability or resistance to compulsive behavioral disorders, including addiction (Flagel, Watson, Akil, & Robinson, 2008; Flagel, Watson, Robinson, & Akil, 2007). This raises the possibility that our procedure might be used to assay the extent to which an individual’s attentional capture is under the control of sign-tracking, and hence to identify people who are predisposed toward showing maladaptive patterns of attentional bias that are associated with addiction and other compulsive disorders.

### Conclusions

As noted in the introduction, research on derived attention has proceeded largely independently and in parallel in the associative learning and perception–cognition literatures. We hope that the research reported here might go some way toward bringing these literatures closer together. Our results demonstrate value-driven attentional capture by task-irrelevant stimuli. More generally, however, these studies show that the well-established techniques of perceptual–cognitive research (such as visual search and eye tracking) can be used to shed light on the associative learning mechanisms underlying derived attentional capture. And as a complement, a consideration of associative learning theory (such as the distinction between Pavlovian and instrumental conditioning) can shed light on the cognitive information structures that support changes in capture of visual attention. This improved understanding of the psychology of derived attention can then be used to investigate the brain mechanisms supporting these processes, and the mechanisms by which some dysfunctional cognitive features (e.g., heightened salience of drug paraphernalia in the case of addiction) are sustained.

### References

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. <http://dx.doi.org/10.1037/a0034575>

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011a). Learned value magnifies salience-based attentional capture. *PLoS ONE*, *6*, e27926. <http://dx.doi.org/10.1371/journal.pone.0027926>

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. <http://dx.doi.org/10.1073/pnas.1104047108>

Anderson, B. A., Laurent, P. A., & Yantis, S. (2012). Generalization of value-based attentional priority. *Visual Cognition*, *20*, 647–658. <http://dx.doi.org/10.1080/13506285.2012.679711>

Anderson, B. A., & Yantis, S. (2012). Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. *Attention, Perception, & Psychophysics*, *74*, 1644–1653. <http://dx.doi.org/10.3758/s13414-012-0348-2>

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. <http://dx.doi.org/10.1016/j.tics.2012.06.010>

Boakes, R. A. (1977). Performance on learning to associate a stimulus with positive reinforcement. In H. Davis & H. Hurwitz (Eds.), *Operant-Pavlovian interactions* (pp. 67–97). Hillsdale, NJ: Erlbaum.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436. <http://dx.doi.org/10.1163/156856897X00357>

Chelazzi, L., Perlatto, A., Santandrea, E., & Libera, C. D. (2013). Rewards teach visual selective attention. *Vision Research*, *85*, 58–72. <http://dx.doi.org/10.1016/j.visres.2012.12.005>

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

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. [http://dx.doi.org/10.1016/S0376-8716\(02\)00219-3](http://dx.doi.org/10.1016/S0376-8716(02)00219-3)

Dayan, P. (2009). Dopamine, reinforcement learning, and addiction. *Pharmacopsychiatry*, *42*, S56–S65. <http://dx.doi.org/10.1055/s-0028-1124107>

Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*, *20*, 778–784. <http://dx.doi.org/10.1111/j.1467-9280.2009.02360.x>

Eimer, M. (1996). The N2pc component as an indicator of attentional selectivity. *Electroencephalography & Clinical Neurophysiology*, *99*, 225–234. [http://dx.doi.org/10.1016/0013-4694\(96\)95711-9](http://dx.doi.org/10.1016/0013-4694(96)95711-9)

Failing, M. F., & Theeuwes, J. (2014). *Don’t let it distract you: How information about the availability of reward affects attentional selection*. Manuscript submitted for publication.

Flagel, S. B., Akil, H., & Robinson, T. E. (2009). Individual differences in the attribution of incentive salience to reward-related cues: Implications for addiction. *Neuropharmacology*, *56*, 139–148. <http://dx.doi.org/10.1016/j.neuropharm.2008.06.027>

Flagel, S. B., Watson, S. J., Akil, H., & Robinson, T. E. (2008). Individual differences in the attribution of incentive salience to a reward-related cue: Influence on cocaine sensitization. *Behavioural Brain Research*, *186*, 48–56. <http://dx.doi.org/10.1016/j.bbr.2007.07.022>

Flagel, S. B., Watson, S. J., Robinson, T. E., & Akil, H. (2007). Individual differences in the propensity to approach signals vs goals promote different adaptations in the dopamine system of rats. *Psychopharmacology*, *191*, 599–607. <http://dx.doi.org/10.1007/s00213-006-0535-8>

Frank, M. J. (2008). Schizophrenia: A computational reinforcement learning perspective. *Schizophrenia Bulletin*, *34*, 1008–1011. <http://dx.doi.org/10.1093/schbul/sbn123>

Holland, P. C. (1979). Differential effects of omission contingencies on various components of Pavlovian appetitive conditioned responding in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*, 178–193. <http://dx.doi.org/10.1037/0097-7403.5.2.178>

Hopf, J. M., Luck, S. J., Girelli, M., Hagner, T., Mangun, G. R., Scheich, H., & Heinze, H. J. (2000). Neural sources of focused attention in visual search. *Cerebral Cortex*, *10*, 1233–1241. <http://dx.doi.org/10.1093/cercor/10.12.1233>

Hyman, S. E. (2005). Addiction: A disease of learning and memory. *The American Journal of Psychiatry*, *162*, 1414–1422. <http://dx.doi.org/10.1176/appi.ajp.162.8.1414>

- James, W. (1890/1983). *The principles of psychology*. Cambridge, MA: Harvard University Press.
- Kapur, S. (2003). Psychosis as a state of aberrant salience: A framework linking biology, phenomenology, and pharmacology in schizophrenia. *The American Journal of Psychiatry*, *160*, 13–23. <http://dx.doi.org/10.1176/appi.ajp.160.1.13>
- 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. <http://dx.doi.org/10.1111/j.1467-9280.2009.02281.x>
- 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. <http://dx.doi.org/10.1037/0278-7393.23.1.246>
- Kleiner, M., Brainard, D. H., & Pelli, D. G. (2007). What's new in Psychtoolbox-3? *Perception*, *36*, ECVF Abstract Supplement.
- Le Pelley, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, *57*, 193–243. <http://dx.doi.org/10.1080/02724990344000141>
- Le Pelley, M. E. (2010). Attention and human associative learning. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour* (pp. 187–215). New York, NY: Oxford University Press.
- 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. <http://dx.doi.org/10.1037/a0033700>
- 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. <http://dx.doi.org/10.1037/a0014525>
- Luck, S. J., & Woodman, G. F. (1999). Electrophysiological measurement of rapid shifts of attention during visual search. *Nature*, *400*, 867–869. <http://dx.doi.org/10.1038/23698>
- Ludwig, C. J. H., & Gilchrist, I. D. (2002). Stimulus-driven and goal-driven control over visual selection. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 902–912. <http://dx.doi.org/10.1037/0096-1523.28.4.902>
- MacLeod, C., Mathews, A., & Tata, P. (1986). Attentional bias in emotional disorders. *Journal of Abnormal Psychology*, *95*, 15–20. <http://dx.doi.org/10.1037/0021-843X.95.1.15>
- 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. <http://dx.doi.org/10.1111/j.1360-0443.2006.01498.x>
- 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. <http://dx.doi.org/10.1093/schbul/sbr192>
- 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). New York, NY: Oxford University Press.
- Peck, C. J., Jangraw, D. C., Suzuki, M., Efem, R., & Gottlieb, J. (2009). Reward modulates attention independently of action value in posterior parietal cortex. *The Journal of Neuroscience*, *29*, 11182–11191. <http://dx.doi.org/10.1523/JNEUROSCI.1929-09.2009>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442. <http://dx.doi.org/10.1163/156856897X00366>
- Posner, M. I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology*, *32*, 3–25. <http://dx.doi.org/10.1080/00335558008248231>
- Robinson, T. E., & Berridge, K. C. (2001). Incentive-sensitization and addiction. *Addiction*, *96*, 103–114. <http://dx.doi.org/10.1046/j.1360-0443.2001.9611038.x>
- 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. <http://dx.doi.org/10.3758/PBR.17.4.536>
- Skinner, B. F. (1948). Superstition in the pigeon. *Journal of Experimental Psychology*, *38*, 168–172. <http://dx.doi.org/10.1037/h0055873>
- Theeuwes, J. (1991). Cross-dimensional perceptual selectivity. *Perception & Psychophysics*, *50*, 184–193. <http://dx.doi.org/10.3758/BF03212219>
- Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*, 599–606. <http://dx.doi.org/10.3758/BF03211656>
- Theeuwes, J. (1994). Endogenous and exogenous control of visual selection. *Perception*, *23*, 429–440. <http://dx.doi.org/10.1068/p230429>
- Theeuwes, J., & Belopolsky, A. V. (2012). Reward grabs the eye: Oculomotor capture by rewarding stimuli. *Vision Research*, *74*, 80–85. <http://dx.doi.org/10.1016/j.visres.2012.07.024>
- Theeuwes, J., Kramer, A. F., Hahn, S., & Irwin, D. E. (1998). Our eyes do not always go where we want them to go: Capture of the eyes by new objects. *Psychological Science*, *9*, 379–385. <http://dx.doi.org/10.1111/1467-9280.00071>
- Theeuwes, J., Kramer, A. F., Hahn, S., Irwin, D. E., & Zelinsky, G. J. (1999). Influence of attentional capture on oculomotor control. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 1595–1608. <http://dx.doi.org/10.1037/0096-1523.25.6.1595>
- Tomie, A., Grimes, K. L., & Pohorecky, L. A. (2008). Behavioral characteristics and neurobiological substrates shared by Pavlovian sign-tracking and drug abuse. *Brain Research Reviews*, *58*, 121–135. <http://dx.doi.org/10.1016/j.brainresrev.2007.12.003>
- Uslaner, J. M., Acerbo, M. J., Jones, S. A., & Robinson, T. E. (2006). The attribution of incentive salience to a stimulus that signals an intravenous injection of cocaine. *Behavioural Brain Research*, *169*, 320–324. <http://dx.doi.org/10.1016/j.bbr.2006.02.001>
- 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. <http://dx.doi.org/10.1037/0278-6133.22.4.378>
- 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.

Received April 13, 2014

Revision received September 24, 2014

Accepted October 7, 2014 ■