



Published in final edited form as:

Nebr Symp Motiv. 2012 ; 59: 91–116.

Reward and Attentional Control in Visual Search

Steven Yantis, Brian A. Anderson, Emma K. Wampler, and Patryk A. Laurent

Department of Psychological & Brain Sciences, Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218-2686, USA

Steven Yantis: yantis@jhu.edu

Abstract

It has long been known that the control of attention in visual search depends both on voluntary, top-down deployment according to context-specific goals, and on involuntary, stimulus-driven capture based on the physical conspicuity of perceptual objects. Recent evidence suggests that pairing target stimuli with reward can modulate the voluntary deployment of attention, but there is little evidence that reward modulates the involuntary deployment of attention to task-irrelevant distractors. We report several experiments that investigate the role of reward learning on attentional control. Each experiment involved a training phase and a test phase. In the training phase, different colors were associated with different amounts of monetary reward. In the test phase, color was not task-relevant and participants searched for a shape singleton; in most experiments no reward was delivered in the test phase. We first show that attentional capture by physically salient distractors is magnified by a previous association with reward. In subsequent experiments we demonstrate that physically inconspicuous stimuli previously associated with reward capture attention persistently during extinction—even several days after training. Furthermore, vulnerability to attentional capture by high-value stimuli is negatively correlated across individuals with working memory capacity and positively correlated with trait impulsivity. An analysis of intertrial effects reveals that value-driven attentional capture is spatially specific. Finally, when reward is delivered at test contingent on the task-relevant shape feature, recent reward history modulates value-driven attentional capture by the irrelevant color feature. The influence of learned value on attention may provide a useful model of clinical syndromes characterized by similar failures of cognitive control, including addiction, attention-deficit/hyperactivity disorder, and obesity.

Keywords

Attentional capture; Reward; Incentive salience; Visual search

Selective attention gates access to awareness. Attentional control therefore determines the contents of awareness and the starting point for almost any behavioral or cognitive act—perceiving, remembering, learning, or behaving. Attentional control has long been a core

issue in cognitive psychology and continues to drive a great deal of empirical and theoretical research.

Two broad domains of control have long been known to determine deployments of attention (Corbetta and Shulman 2002; Egeth and Yantis 1997; Yantis 2000). Voluntary or top-down attentional control is driven by current perceptual goals. When an individual is searching for a particular object or feature, or searching in a particular location, they can voluntarily direct overt attention (eye movements) or covert attention (without eye movements) to the task-relevant object, feature, or location. Such deployments of attention increase the speed and accuracy of behavioral responses (e.g., Pashler 1998) and evoke strong modulation of neural activity in the brain (e.g., Moran and Desimone 1985; Yantis 2008). A vast literature has documented the spatial and temporal properties of voluntary, goal-directed attentional control (see Egeth and Yantis 1997 and Pashler 1998 for reviews).

A second form of control is often referred to as bottom-up or stimulus-driven control (Itti and Koch 2001; Parkhurst et al. 2002; Theeuwes 1992, 2010; Yantis and Jonides 1984; Yantis 1993, 2000). When a salient, unexpected event occurs (e.g., the appearance of a new object, Christ and Abrams 2006; Yantis and Hillstrom 1994; or looming motion, Lin et al. 2009) the perceiver will often orient to that event, even if it may interfere with other ongoing cognitive operations. Here, too, a large literature has explored the limits of stimulus-driven attentional capture and its interactions with voluntary attentional control. Yantis and Jonides (1984) reported that the abrupt onset of a new perceptual object captures attention in visual search even when the onset does not reliably predict the target location (but top-down deployment of attention elsewhere can override this automatic attention response; Yantis and Jonides 1990). Theeuwes (1992) showed that when people search for a shape singleton target (that is, a unique shape in an array of otherwise identical shapes, e.g., a diamond among many circles), then the presence of an irrelevant color singleton (e.g., a red nontarget item when all the remaining items in the display are white) significantly slows search, an indication that it captured attention involuntarily.

A variety of intermediate or hybrid cases of attentional control have also been described. These are situations in which the deployment of attention is not entirely voluntary, but the effect appears to be at least partly a function of information held by the perceiver, either in the form of explicit goals or recent perceptual history. One example of this hybrid category is contingent attentional capture (Folk et al. 1992; Anderson and Folk 2010). In this scenario, subjects are instructed to search for a target defined by a salient feature (e.g., identify the red target in an array of otherwise white items, or identify the object that has an abrupt onset). Immediately preceding the array, a to-be-ignored “cue” appears (this is probably better labeled a “distractor”). The distractor consists of a set of salient items that appear surrounding one of the potential target locations; it occurs shortly before the search array appears and should always be ignored (by virtue of its location—surrounding a possible target location—and time of onset—before the main search array). The main finding is that if the distractor carries a feature that matches the target-defining feature (e.g., the target is red and the distractor is red), then the distractor captures attention (indexed by a faster RT when the distractor appears at the target location than when it appears in a nontarget location). When the distractor does not match the target defining feature (e.g., the

target is red and the distractor is an abrupt onset) then the distractor does not capture attention. So in this case, the content of the search set (e.g., red) guides attention to matching (or similar) features, even ones that are known to be irrelevant (because of their location, shape, or timing; Anderson and Folk 2010; Folk and Remington 1998).

In this chapter, we will explore an influence on attentional control that has received increasing scrutiny in recent years: attention to reward-related stimuli. Reward has long been known to play a key role in learning and cognition (Pessoa and Engelmann 2010; Schultz et al. 1997; Platt and Glimcher 1999; Sugrue et al. 2005). This is hardly surprising: organisms must procure food and water to survive; these primary rewards evoke powerful responses in the brain that lead to learning. For example, “where was I when I found this source of food, so I can find it again?” Saharan camels no doubt learn to associate a clump of palm trees with vital and rewarding water. In order to reproduce, animals must engage in sexual behavior, and sex is among the most powerfully motivating rewards we know. Here again, reward leads to learning, which can in turn increase the probability that the reward can be obtained again in the future.

Powerful learning mechanisms in the brain provide a way for stimulus-reward associations to be learned rapidly and persistently. For example, an animal may learn that a certain kind of tasty berry has a particular color, size, and shape, and so whenever they see that berry, they orient to it and consume it immediately. They may further learn that this kind of berry can be found near a particular species of tree, and so they learn to notice that tree (which is perhaps easier to spot in the dense jungle) and indeed when they see that tree they begin to anticipate the berries they soon will be eating.

A number of recent studies have investigated the role of reward on attention. They have shown that reward plays an important role in voluntary, deliberate deployments of attention in a variety of contexts. Raymond and O'Brien (2009) showed participants several novel faces during a training phase, and consistently followed each face with different amounts of positive or negative monetary reward. Following the training phase, participants carried out an attentional blink task in which two targets, each followed by a mask, are shown in rapid succession. The second target was one of the faces that had appeared during the training phase. In attentional blink tasks, the identification of the second target (termed T2) is typically impaired at short delays between the two targets. The probability of correctly recognizing a face was much greater when the face had been associated with large positive or negative rewards during training than if it had been associated with low or no reward. Indeed, faces associated with a large positive reward showed no evidence of an attentional blink. The authors concluded that learning to associate stimuli with high positive reward during learning enhances the degree to which those task-relevant stimuli can compete for attention.

Della Libera and Chelazzi (2009) had participants learn associations between complex shapes and reward in a matching task. Participants were presented with two superimposed colored shapes on one side of fixation, and a black shape on the other. They compared one of the two superimposed colored shapes to the black shape and indicated whether they were the same or different. A color cue at the beginning of each trial informed participants which

colored shape to use in the comparison. Monetary rewards were delivered that varied with the identity of the colored shapes; regardless of color, some shapes tended to predict more reward than others. Following performance of this task, which lasted several thousand trials over multiple days, participants were tested without rewards (i.e., in extinction) both in the same task and in a simple visual search task. The results showed that formerly reward-predictive shapes impaired performance as the to-be-ignored shape in the matching task, but did not impair performance as irrelevant distractors in the visual search task. However, visual search was facilitated for goal-relevant targets that were formerly predictive of high reward. This was taken as evidence that reward learning can have persistent effects on attentional selection.

Krebs et al. (2010) employed a Stroop task to investigate the effect of reward association on attention; certain color words were associated with the delivery of monetary reward when the color was correctly named. Participants were to name the ink color of color words; in some cases, the color words were different than the ink color, and these incompatible trials often slow responding due to competition between the required ink-color response and the automatic word-reading response. Overall, responses to rewarded ink colors were faster than those to unrewarded colors, consistent with increased attentional priority to the rewarded feature. When the to-be-ignored color name was reward-related (e.g., “blue” when the ink color blue predicted reward), it tended to magnify the usual slowing caused by color-word conflict (compared to when the color word was not reward-related). This study shows that stimuli associated with high reward in the current task context tend to draw attention even when those stimuli should be ignored.

Peck et al. (2009) measured neural responses in monkey area LIP, which is thought to contain a map-like representation of attentional priority (Bisley and Goldberg 2010). On each trial of the experiment, a target appeared in the display and the monkey had to make a rapid eye movement to the target's location. Almost a second before the target appeared, one of two cue shapes appeared in the display; the cue's location did not reliably predict the subsequent target location—in other words, it was not informative about the upcoming required eye movement. One of the two cue shapes indicated that the trial would be followed by a juice reward; the other shape indicated that no reward would be delivered. When the cue indicated that a reward was forthcoming, it evoked a much stronger neural response than if it indicated no reward would be delivered. Furthermore, behavioral responses were faster and more accurate when the subsequent target stimulus appeared in the location of a reward-predicting cue compared to a cue predicting no reward (indeed, the no-reward cues exhibited evidence of attentional repulsion). This behavioral effect persisted in second task in which well-learned cues no longer predicted reward, but only when the monkeys knew reward would not be available on the current trial. The latter finding suggests a persisting attentional effect of the former reward learning, but one that can be largely overcome by current reward-related goals.

To account for how physical salience and reward value might be combined to give rise to a representation of attentional priority, Navalpakkam et al. (2010) presented visual search displays to human participants; each display contained multiple nontarget bars all with the same tilt and two targets that differed from the nontargets—one with a tilt that differed only

slightly from the nontargets (low physical salience) and one with a very different tilt (high physical salience). The reward value of the two targets was varied across blocks so that various combinations of physical salience and reward value could be tested. They found that performance depended on both physical salience and relative amount of reward in a pattern that was best accounted for by a Bayesian ideal observer model in which expected value is maximized.

Several studies have shown that reward delivery gives rise to involuntary deployment of attention on the very next trial, a phenomenon that can be termed “reward priming” (Della Libera and Chelazzi et al. 2006; Hickey et al. 2010a, b). For example, Hickey et al. (2010a) conducted a visual search task based on the additional singleton paradigm introduced by Theeuwes (1992). Participants searched for a shape singleton (e.g., a diamond in an array of circles) and reported the orientation of a small line segment contained in the target. On some trials, all the shapes were rendered in the same color (red or green). On many trials, however, one of the nontarget shapes had a unique color (red among green or vice-versa). Each trial was followed by a feedback display containing the amount of reward that was received on that trial, either 1 point or 10 points, which translated into monetary reward at the end of the experiment. Reward was delivered randomly.

The key manipulation was whether the colors of targets and distractors swapped from one trial to the next. For example, if on trial N the target and most of the nontargets were red (and the color singleton was green) then on trial N + 1 the color assignment could be the same, or it could swap so that now the target and most of the distractors were green (and the color singleton was red). The main result was that on trials in which the colors did not swap, a high reward on trial N yielded faster responses on trial N + 1 than did a low reward on trial N. However, when the colors swapped, this pattern reversed: a high reward on trial N gave rise to slower responses on the following trial than did a low reward on trial N. This outcome indicates that when a particular color is rewarded, that color appears to draw attention to itself on the immediately following trial—even though color is not relevant to this task. Della Libera and Chelazzi (2006) reported a similar result in the context of a global/local judgment task.

Serences (2008) used a choice paradigm to study the effect of reward history on both the representation and selection of visual stimuli. Human participants selected one of two colored circles via a button press, and their selection either was or was not followed by the delivery of monetary reward. Throughout the course of the experiment, the relative probability of receiving a reward for selecting one color over the other was either 1:1, 1:3, or 3:1, which changed periodically during the experiment. The results showed that the recent reward history of each color predicted both stimulus selection and stimulus-evoked response in early visual areas as measured by functional magnetic resonance imaging (fMRI). Using a similar paradigm, Serences and Saproo (2010) extended this finding by showing that oriented gratings associated with larger rewards are represented with greater precision in early visual areas of the human brain. Additionally, Shuler and Bear (2006) found that when light flashes predicted reward, responses in rat area V1 reflected temporal expectations concerning reward delivery, with activity being either maximal or minimal at the time of

expected reward. Collectively, these results argue that current stimulus-reward associations bias perception, consistent with attentional priority to high-value features.

In the studies summarized above, positive effects of reward—usually, faster and/or more accurate responses, accompanied by increased brain activity, to high-reward stimuli—were observed when the task in question involved currently rewarded stimuli and/or stimuli that are currently task-relevant (e.g., they were the targets of search or they predicted reward). However, reward-related effects of task-irrelevant stimuli have only been observed in trials immediately following reward. Although investigators have in a few cases examined the effect of a to-be-ignored stimulus previously associated with reward, in most cases the previously reward-related stimuli failed to capture attention (Della Libera and Chelazzi 2009, Experiment 2; Krebs et al. 2010, Experiment 2; Raymond and O'Brien 2009, Experiment 2; for an exception, see Peck et al. 2009, Fig. 8). In these experiments, however, the stimuli in question were typically complex multi-feature or multidimensional objects such as complex shapes, words, or faces. This aspect of the stimuli may have precluded them from exerting a significant, persistent effect on observable behavior.

In this chapter, we describe several experiments that explore the degree to which otherwise neutral and elementary visual features (e.g., colors) can be associated with different amounts of monetary reward through training and feedback, and later come to capture attention involuntarily when they appear in contexts in which those stimuli should be ignored. The results of the experiments discussed in this chapter were originally reported in Anderson et al. (2011a, b, 2012).

Learned Value Modulates Salience-Based Attentional Capture

As noted earlier, when a person searches for a salient visual stimulus defined as a feature singleton in one dimension (e.g., a unique shape in an array of other shapes—diamond among circles or vice-versa), then the presence of a more salient feature singleton in another dimension captures attention and slows search (Theeuwes 1992). In the first experiment, we asked whether this form of stimulus-driven attentional capture can be modulated by reward associations.

Experiment 1 has several features in common with most of the experiments reported in this chapter. The experiment had two phases, a training phase and a test phase. In both phases, participants engaged in visual search for a target among five nontargets (Fig. 1). In the training phase, subjects searched for a red or a green target circle in an array of six circles, each rendered in a different color, presented on a black background. Exactly one item in each display was red or green (each color appeared as the target equally often in random order). Inside every nontarget circle was a white line segment that was oriented 45° to the left or right (randomly selected in each nontarget circle). The line segment inside the (red or green) target circle was either vertical or horizontal. The subject's task was to press one button if the line segment contained within the target circle was vertical and another button if it was horizontal. Following each correct response, a feedback display indicated that the participant had received a small monetary reward. The cumulative value of all rewards

earned so far was also displayed. Following an incorrect response, a feedback display indicated that no money had been earned on the trial.

The magnitude of the reward following correct responses was either large or small (in this experiment, 5¢ or 1¢). One color had a high probability ($p = 0.8$) of yielding a large reward and a low probability ($p = 0.2$) of yielding a small reward; this mapping was reversed for the other color. This partial reinforcement schedule, which is more resistant to extinction (Rescorla 1999), was modeled after Della Libera and Chelazzi (2009). The high- and low-reward colors were counterbalanced across subjects. The training phase was designed so that participants came to associate one color with high reward and the other color with low reward. In this experiment, the training phase included 1008 trials.

After a short break, the test phase began. In the test phase, subjects searched for a unique shape (diamond among circles or circle among diamonds, unpredictably). Once again, each shape contained a small line segment that was oriented at $\pm 45^\circ$, but the target shape contained either a vertically or horizontally oriented line segment (Fig. 1b). The subject was to identify the orientation of the line segment within the unique shape and press the corresponding button as rapidly as possible. The test phase consisted of 480 trials.

On most trials in the test phase, all the items were rendered in white on a black background. On a randomly-selected one-quarter of the trials, one of the nontarget shapes was rendered in red, and on one-quarter in green. Subjects were told to ignore color—the target shape was never the color singleton. Each response was followed by feedback about whether the response was correct or not—no reward feedback was provided during the test phase.

Many previous studies have demonstrated that the presence of a color singleton (red or green) captures attention and slows responses in the shape-search task (e.g., Theeuwes 1992). Thus, in this experiment, we expected to find that responses on trials containing either a red or a green distractor were significantly slowed compared to responses on trials without a color singleton distractor. The main question was whether the reward association established in the training phase affects the magnitude of this slowing. Eighteen participants completed the experiment in a single 2 h session, and earned between US\$ 21 and 28 (mean = US\$ 25.22).

We first checked the mapping of color to reward (i.e. red vs. green as the high-reward color) and found no interaction with the effect of reward on distraction ($F < 1$), so we collapsed across color in the remaining analyses. RT differed significantly in the three distractor conditions [Fig. 2a, $F(2,34) = 48.6, p < 0.001$]. Both the high-value and low-value distractor significantly slowed RT compared to the no-distractor condition [$t(17) = 8.45$ and 6.31 , respectively, both $p < 0.001$], which replicates the standard effect of a salient color singleton on search time.

More importantly, we found that the presence of a high-value distractor slowed search more than the presence of a low-value distractor [$t(17) = 3.37, p = 0.004$]. The difference in slowing due to reward history was not due to differences in physical salience (red slowed search more than green when it had been the high-reward color, and green slowed search more than red when it was the high-reward color). The effect of reward history on RT

declined over the course of the test phase, as revealed by a linear trend in the difference between RTs for high- and low-value distractor trials over trial epoch [Fig. 2a, $F(1,17) = 17.22, p = 0.001$].

Several previous studies had shown that reward influences both behavioral and neural responses (Platt and Glimcher 1999; Simen et al. 2009; Sugrue et al. 2005). In this experiment, there was not a robust difference in RT to high- and low-reward targets during the training phase, when rewards were being delivered in the form of monetary feedback (the mean difference in RT to the high-reward and low-reward color targets during the training phase was just 3 ms). However, there were substantial individual differences in this effect—some participants responded as much as 20 ms faster to the high-reward color than to the low-reward color on average, while others exhibited a negligible or even slightly negative difference. We wondered whether these individual differences could reflect the degree to which people differed in their sensitivity to reward. Such differences have been investigated in a variety of contexts and are thought to underlie differences in self-control in those contexts (Braver et al. 2010).

To examine this, we plotted the magnitude of value-driven attentional capture (that is, the mean difference in RT to high-value vs. low-value distractors) in the test phase as a function of reward-based speed of responding (that is the mean difference in RT to the high-reward vs. low-reward targets) in the training phase (Fig. 2b). The Pearson's correlation coefficient was significant ($r = 0.553, p = 0.017$), showing that individuals who responded more rapidly to high-reward targets during the training phase also responded more slowly in the presence of high-value distractors in the test phase.

We divided the 18 participants into two groups according to the mean difference in RT to high- and low-reward targets in the training phase (median split). Although the magnitude of value-driven attentional capture in the test phase task declines for the group as a whole (Fig. 2a), the slowing persisted longer for the subset of participants who exhibited greater reward sensitivity during the training phase (Fig. 3).

These results suggest that a salient distractor that was previously associated with high reward slows search more than a salient distractor previously associated with low reward. It is possible that this effect was merely the result of a persisting search set for the two color targets, and particularly—and for some participants—for the highly rewarded target, during the test phase, even though in the test phase color was not task relevant and was unrewarded. This possibility appears unlikely given recent evidence that people can adjust their deliberate search set rapidly and effectively: Lien et al. (2010) showed that participants could rapidly and flexibly adjust their search set on a trial-by-trial basis according to task requirements. However, some studies have shown that unrewarded former targets can capture attention under certain circumstances (Kyllingsbaek et al. 2001; Shiffrin and Schneider 1977). To definitively exclude the possibility of a persisting search set for the color distractors, eighteen new participants completed a control experiment that was very similar to Experiment 1, except that no reward feedback was provided during the training phase. Each participant in Experiment 2 was given a flat compensation of US\$ 25, approximately the mean amount earned by participants in Experiment 1.

In the training phase of Experiment 2, half the participants searched for red and blue targets in a multicolored array, and green was among the nontargets on half the trials (equally often with each color target); the other half of the participants searched for green and blue targets and red was among the nontargets half the time. The test phase was identical to that in Experiment 1: subjects searched for a unique shape (diamond among circles or vice-versa, unpredictably). All items were white on a black background, except that on one-quarter of all trials, one of the nontargets was red and on one-quarter of the trials one of the nontargets was green. Because participants had just spent 1008 training trials searching for red and always ignoring green (or vice-versa for half the participants), any residual tendency to continue to search for the previously prioritized target color should be evident in comparing trials in which a former target was the distractor vs. trials in which a former nontarget was the distractor.

The presence of a color singleton distractor significantly slowed search, as expected. However, there was no difference in RT on trials containing a former target vs. those containing a former nontarget [Fig. 2c, $t(17) = 0.34$, *n.s.*]. Furthermore, the amount of slowing caused by a high-value distractor in Experiment 1 was greater than that caused by a former target-colored distractor in Experiment 2 [mean difference = 27 ms, $t(34) = 2.29$, $p = 0.025$], confirming that learned associations between stimuli and reward indeed increase distraction beyond that produced by previously prioritized stimuli that are not associated with reward. We take this as evidence that slowing in the test phase in Experiment 1 was not due to a persisting search set for a former target.

Experiment 1 shows that when a salient but irrelevant item appears in a search array, the degree to which it slows search depends on its reward history. There are at least two possible mechanisms for this effect that can be considered. One is that reward history makes an item “more salient”—that is, the physical properties of the item that make it perceptually salient (e.g., local color contrast) are combined with the reward associations to magnify the conspicuity or pertinence (Bundesen 1990) of the item, increasing the probability that the item will be selected during search. In effect, the item competes more effectively for attention with the shape target.

A second possibility is that reward history has an effect only *after* the item has captured attention by virtue of its physical salience, by prolonging dwell time or increasing disengagement costs (Belopolsky et al. 2010; Duncan et al. 1994). On this account, the increased reward association causes an attended item to continue to “hold” attention when that item had previously been highly rewarded.

The present experiments cannot distinguish between these accounts. However, in the remaining experiments we ask whether reward history can evoke an involuntary shift of attention even when the item in question is not physically salient. If it can, then we may be able to conclude that reward history has an effect that is not strictly a consequence of increasing dwell time or disengagement costs.

Physically Inconspicuous Items Previously Associated With Reward Capture Attention

In the remainder of the chapter, we describe several experiments that investigate how otherwise inconspicuous and task-irrelevant stimuli associated with reward during training capture visual attention. Experiment 3 was similar in design to Experiment 1, except that the items in each trial of the test phase were rendered in multiple different colors, thereby ensuring that the previous target colors were not themselves physically salient (Fig. 4). The training phase was identical to that in Experiment 1: 1008 trials of visual search for a red or a green circle containing a vertical or horizontal line segment that specified the correct response. For half the participants, red targets were followed by high reward with probability $p = 0.8$ and by a low reward with $p = 0.2$, and green targets had the complementary contingencies; for the remaining participants, this mapping was reversed (Table 1).

In this experiment, because the red and green distractors are not physically salient, any slowing of RT they produce must be attributed to the value with which they were imbued in the training phase. Responses were indeed slowed by value: RT was significantly slower on trials containing a high-value distractor than when no value-related distractor was present; the mean RT in the presence of a low-value distractor was intermediate [ANOVA: $F(2,50) = 6.07$, $p = 0.004$; linear trend: $F(1,25) = 12.19$, $p = 0.002$; see Table 2). This result shows that an otherwise neutral stimulus feature captures attention when that feature has previously been associated with reward.

There were individual differences in susceptibility to value-driven attentional capture. We wondered if those differences were systematically related to other cognitive abilities or personality traits that have been shown in other studies to covary with attentional control. One such ability is visual working memory capacity. Fukuda and Vogel (2009) measured working memory capacity using a visual change detection task in which a memory array of colored squares is presented briefly, followed by a 900 ms retention interval, and then a test display containing a single probe square that is either the same or different in color than the square previously presented at the probed location. Performance generally declines with the number of items to be remembered during the delay interval. Using a simple formula, a parameter K , representing the individual's working memory capacity, can be estimated.

Fukuda and Vogel estimated visual working memory capacity, and then asked subjects to carry out a task requiring them to focus attention at a spatially cued location (without moving their eyes) to perform a visual discrimination task. Immediately following the visual display, a potentially distracting dot appeared either at the target location or at a nontarget location. The investigators reasoned that a person who was able to focus attention very well at the cued location should exhibit a much larger neural response to the dot when it appeared at the cued location than when it appeared elsewhere. A person who has less effective control over attention should exhibit a smaller difference in the magnitude of the neural response to probe dots appearing in the cued and uncued locations, respectively—in particular, the response to the dot should be larger when it appears in a to-be-ignored location particularly for low-capacity individuals compared to high-capacity individuals.

These two measures (effectiveness of focused attention and working memory capacity) were strongly correlated across individuals ($r = 0.73$). The authors argued that working memory requires effective attentional control in order to maintain information over a delay interval by minimizing the degree to which distracting information displaces to-be-remembered items. If this is correct, then we might expect that individuals with high working memory capacity should also be less vulnerable to value-driven attentional capture.

Individual differences in impulsivity and reward sensitivity have also been shown to be related to attentional control. Hickey et al. (2010b) showed that reward sensitivity (measured by a normed questionnaire that queries the degree to which the individual seeks reward in everyday life) is correlated with the degree to which a color distractor that was associated with high reward on the preceding trial of a search task would slow search on the current trial. Dickman and Meyer (1988) found that individual differences in impulsivity are associated with individual differences in the speed and accuracy in the performance of visual tasks. Impulsivity is thought to contribute to the comorbidity of ADHD and predisposition to substance abuse (Groman et al. 2008). We therefore measured trait impulsivity using the Barratt Impulsiveness Scale (Patton et al. 1995) as another covariate. Here again, we would expect individuals who are impulsive by self-report should also be more vulnerable to value-driven attentional capture.

We used visual working memory capacity and trait impulsivity as predictor variables in a simultaneous regression model of value-driven attentional capture. The model accounted for a significant proportion of the variance in value-driven capture ($R^2 = 0.355$, $p = 0.006$); the regression weights for both predictor variables were significantly greater than zero (est. $\beta = 0.378$, $p = 0.038$ for impulsivity; est. $\beta = -0.554$, $p = 0.004$ for WM capacity). Thus, individuals exhibiting greater impulsivity and lower visual WM capacity were more vulnerable to value-driven attentional capture. Figure 5 (diamonds and grey line) shows a scatter plot of WM capacity vs. value-driven capture in Experiment 3.

We performed a control experiment, as before, to examine the possibility that the effects observed here reflected a continuing top-down set for former targets, rather than an effect that depended critically upon the receipt of reward during the training phase. As noted in the description of Experiment 2, there is evidence that search sets can be adjusted rapidly in visual search, but we wished to confirm this in the present context. Experiment 4 was identical to Experiment 3, with ten naïve participants, except that no reward feedback was provided during the training phase. All participants were given a flat US\$ 25 payment for their participation, an amount that matched the average reward earned in Experiment 3. The absence of reward during training eliminated any slowing by former targets in the test phase [$t(9) = -10.39$, *n.s.*; see Table 3]. There was also no difference in error rate among the red, green, and no-distractor conditions [$F(2,18) = 2.30$, *n.s.*].

In Experiment 5, we reduced the amount of training and increased the delay between training and test to examine the robustness and persistence of value-driven attentional capture. The training and test phases of the experiment were identical to those in Experiment 3, with two exceptions: first, the number of trials in the training and test phases were reduced to 240 each (compared to 1008 and 480, respectively, in Experiment 3); second, the

magnitude of high and low rewards delivered in the training phase were increased to 10¢ and 2¢, respectively (compared to 5¢ and 1¢ in Experiment 3).

We observed significant value-driven attentional capture with this reduced training regimen that was nearly as large in magnitude as that observed in Experiment 3 [$F(2,46) = 5.17, p = 0.009$; see Table 2]. We also observed a significant negative correlation between visual working memory capacity and value-driven capture ($r = -0.468, p = 0.021$; Fig. 5, circles and black line) but the correlation with trait impulsivity, while positive, was not significantly greater than zero ($r = 0.093, n.s.$). We noted that the variance of impulsivity in Experiment 5 was quite low compared both to that observed in Experiment 3 and to that reported by Patton et al. (1995); either this restriction of range in this college-student population or the reduced number of training trials could have led to the absence of a significant correlation with impulsivity.

We invited the participants in this experiment back to the lab after several days had elapsed since training (4–21 days, mean = 8.8 days); 19 of the 24 participants agreed to return. These participants only completed the test phase of the experiment with no further training and no reward delivered. We continued to observe significant slowing due to formerly rewarded color distractors [$F(2,36) = 5.81, p = 0.007$; see Table 2].

The slowing caused by formerly rewarded distractor items in the test phase could have one of two possible sources. The first possibility is that the presence of a formerly highly rewarded item in the display caused a general slowing or filtering cost (Folk et al. 2009). The second possibility is that the formerly rewarded distractor caused a spatially specific, involuntary shift of attention to its location. To examine these possibilities, we took advantage of the spatially-specific effects of inhibition of return (IOR; Theeuwes and Godijn 2002). When attention is directed to an item and subsequently withdrawn, a residual inhibition of subsequent targets appearing in that location is reliably observed.

We examined response times on the subset of all trials that met two criteria: (a) they did not contain a low- or high-value distractor item, and (b) a high-value distractor appeared on the previous trial. We then separated these into those trials in which the target on trial N appeared in the same location as the high-value distractor on trial N–1, and those in which the target appeared in a different location than the preceding high-value distractor. We found RTs were 66 ms slower when the target appeared in the location of a distractor on the previous trial than when it appeared elsewhere [$t(23) = 3.13, p = 0.005$]. This shows that value-driven capture is a spatially-specific deployment of attention.

In Experiment 6, we tested whether learned associations between a stimulus feature (color) and reward can have an involuntary influence on attentional priority that extends to different stimuli and different task contexts. To this end, we had 21 new participants engage in a training phase similar to that employed in Experiment 5, with the exception that three different colors were used for the critical items (red, green, and blue). Two colors served as targets during training: one predicted high reward with probability $p = 0.8$ while the other predicted low reward with probability $p = 0.8$. The third color was always a nontarget (among other things, this manipulation allowed us to extend our findings to colors other than

red and green). In the 480-trial test phase, participants engaged in a flankers task (Eriksen and Eriksen 1974) in which they reported the identity of a centrally presented white letter while ignoring colored flanking letters on the left and right. These flanking letters could either be associated with the same response as the centrally presented letter (response compatible) or a different response (response incompatible), although the center and flanking letters were never physically identical. The compatibility effect, that is, the degree to which RTs are slowed on incompatible vs. compatible trials, can be taken as an index of the degree to which the irrelevant flankers have been processed. We observed larger compatibility effects for flankers that were the color of a formerly high-reward target compared to those that were the color of a formerly low-reward target [mean difference = 11.6 ms, $t(20) = 2.59$, $p = 0.017$], demonstrating value-driven attentional capture that transferred across stimuli (colored outline shapes to colored letters) and across tasks (visual search to a flankers task).

Reward at Test can Magnify Value-Driven Attentional Capture

In all the experiments described to this point, the search for shape in the test phase was unrewarded, except in the usual sense that participants experience an amorphous sense of satisfaction from carrying out the task they have been asked to do. In everyday life, however, multiple items may be associated with different amounts of reward, and these typically compete for attention—for example, when you open your refrigerator in order to find something non-specific to snack on, each food item competes for attention according to how rewarding its consumption has been in the past and according to your current motivational state.

In Experiment 7, we used a similar training regimen as in Experiment 5 (240 trials, 6¢ and 2¢, respectively, for high and low reward). The test phase lasted 240 trials, but now the two shape targets (diamond among circles and vice-versa) were each associated with reward: one shape was followed with high reward (3¢) on 80 % of the trials and low reward (1¢) on 20% of the trials; the other shape had the complementary mapping. The mapping of reward to shape was as always counterbalanced across subjects.

The first, simplest, question was whether the delivery of reward at test that is driven by the currently relevant target shapes would overpower and abolish value-driven capture by the formerly rewarded color distractors. The answer is no: value-driven capture by the color distractors was robust in the face of competition from currently rewarded shapes [$F(2,30) = 16.63$, $p < 0.001$], and indeed larger in magnitude than in the previous experiments, although the difference between high- and low-value distractors was no longer evident.

This experiment offers an opportunity to examine recent reward history on the magnitude of value-driven attentional capture. This is because, unlike all the previous experiments, rewards were delivered on each trial in the test phase, and prior research shows that participants are sensitive to the amount of reward received in the last few trials (e.g., Serences 2008). Thus, we can examine the degree to which the delivery of monetary reward on recent trials affects how strongly a formerly reward-associated color slows search on the current trial. One sensible possibility is that if a particular shape target has received high

reward on recent trials, search for the rewarded shape would be facilitated on the current trial, and thereby reduce the effect of value-associated color distractors.

We observed exactly the opposite outcome. We examined whether the magnitude of value-driven attentional capture (i.e., the difference in RT on trials containing a distractor vs. trials containing neither of the value-associated colors) depended on recent reward history for the current shape target. Figure 6a shows value-driven capture as a function of the mean reward obtained over the last five trials on which the current target shape was the target. The larger the reward the current target had received, the greater the magnitude of value-driven capture by a formerly rewarded color distractor [$F(1,15) = 6.97, p = 0.019$]. This strongly suggests that when a particular shape appears that has recently been associated with high reward, it causes the visual system to be particularly susceptible to capture by *any* stimulus that has been associated with reward. Of course, given the reward schedule, the shape associated with high reward was very often the high reward shape, and the shape associated with low reward was very often the low reward shape. Thus, this analysis reveals an effect of reward association.

A more interesting analysis involves examining not reward associated with the stimuli, but instead the effect of reward-prediction error on the previous trial. Reward prediction error is an extremely important driver of learning to associate particular stimuli with reward (Schultz et al. 1997). When a neutral stimulus appears, the animal expects no reward; if this low reward-prediction is then followed by the delivery of an unexpected reward, a strong reward-prediction error response can easily be measured in the ventral striatum following the reward (e.g., McClure et al. 2003; O'Doherty et al. 2003; Schultz et al. 1997). As learning proceeds, the appearance of the reward-predicting stimulus itself begins to evoke a reward prediction response, and the reward-prediction error signal at the delivery of the expected reward declines, because the reward is no longer unexpected. Similarly, if a reward-predicting stimulus appears but no reward is delivered, the negative reward-prediction error causes a reduction in the response of these neurons when the expected reward should have been delivered.

In the current situation, we can look at the magnitude of value-driven attentional capture when the immediately preceding trial involved different reward prediction errors. In Experiment 7, three types of reward-prediction error were possible: a target with a low probability of high reward followed by high reward (this is a positive reward-prediction error—the participant is pleasantly surprised); a target with a high probability of a high reward followed by low reward (this is a negative reward-prediction error—disappointment); and the reward following a given stimulus is as expected (zero reward-prediction error). When we examined the magnitude of value-driven capture following each of these three amounts of reward-prediction error, we found the greatest value-driven capture following positive reward-prediction error on trial N-1, and no significant value-driven capture following negative reward-prediction error on trial N-1 [Fig. 6b, ANOVA: $F(2,30) = 4.63, p = 0.018$]. This finding, which is not entirely independent of the finding reported earlier, also suggests that the degree to which a perceiver is susceptible to value-driven capture depends in part on recent reward learning. Large positive reward-prediction

error produces strong value-driven capture, as if the visual system is temporarily hyper-responsive to high-value stimuli.

Discussion

In this chapter we have described a series of experiments that provide new evidence concerning the role of reward learning in the deployment of attention during visual search. The experiments used well-validated visual search tasks to show that a stimulus feature associated with increased reward during training slows visual search more than a feature associated with lower or no reward.

In the first experiment, physically salient color distractors prolonged visual search more during a shape singleton search task when it was previously associated with high reward than when it had been associated with low reward. Experiment 2 showed that this effect depended upon the delivery of reward during the training phase, and was not merely a persisting top-down set for former targets. Experiment 6, in which value-driven attentional capture generalized to a flankers task involving letters (rather than the circles that were used in training), showed that value-driven capture results from associations between prior reward and predictive stimulus features (in this case color), and does not reflect attentional capture that is specific to previously rewarded objects.

The slowing observed in Experiment 1 has at least two different interpretations. It could be that a color singleton formerly associated with high reward has a greater probability of capturing attention than one associated with low reward. Alternatively, it could be that the initial capture of attention by the color singleton was unaffected by reward history, but that the time required to disengage from the color distractor depended on whether it had previously been associated with high or low reward. Of course, both factors could play a role.

In order to determine whether reward history directly causes attentional capture, the remaining experiments used color distractors that were not physically salient—that is, they were always presented within arrays of items that varied in color. Experiment 3 showed that a physically inconspicuous nontarget that is rendered in a color that had been associated with high reward during the training phase slowed search more than one formerly associated with low reward or when neither color was present among the nontargets. This shows that the reward association itself gives rise to a unique mode of attentional capture, one that differs from the well-established stimulus-driven and goal-directed modes. Experiment 4 confirmed that the rewards delivered during training are essential to these observed attention effects.

Value-driven attentional capture can be evoked with even brief training, as shown in Experiment 5, and it persists for at least a week—indeed, in later experiments we have observed significant value-driven attentional capture when participants came back to the lab as long as 8 months after initial training.

We also found that individuals vary in the degree to which they are susceptible to value-driven attentional capture. These individual differences are correlated with visual working

memory capacity (low-capacity individuals exhibit more prolonged slowing due to value-driven capture) and with trait impulsivity (high impulsive individuals exhibit stronger value-driven capture). Individual differences in visual working memory capacity are thought to reflect variation in a general ability to resist distraction, such that individuals with high working memory capacity are better able to restrict selection to relevant items both in vision and in working memory (Fukuda and Vogel 2009, 2011). Our results support this idea by showing a consistently negative correlation with visual working memory capacity and susceptibility to distraction by valuable but task-irrelevant stimuli. Trait impulsivity is thought to reflect, in part, a measure of the degree to which an individual is able to inhibit a prepotent response (Dickman and Meyer 1988; Groman et al. 2008). Our results provide evidence that, at least with longer training regimens, more impulsive people are less able to avoid the tendency to select the valuable stimulus rather than the current target of visual search.

The effect of value-driven attentional capture is spatially specific: RT to targets appearing in a location occupied on the previous trial by a high-value distractor are especially slow—a manifestation of inhibition of return (Theeuwes and Godijn 2002). When a salient distractor captures attention, the subsequent active suppression of that item in order to withdraw and direct attention to the target of search gives rise to a persisting inhibitory signal at that location. Subsequent voluntary deployment of attention is slowed by this inhibition, and it is this slowing that we observed. This IOR-based signature provides strong evidence for a spatially-specific instance of involuntary attentional deployment.

The present findings may provide a model for impaired attentional capture in a variety of clinical syndromes that involve failures of cognitive control, including, for example, substance abuse. Consumption of alcohol, nicotine, cocaine, and other drugs of abuse cause the release of dopamine in the nucleus accumbens in the ventral striatum of the basal ganglia, and through repeated use can come to usurp the brain's reward circuitry, leading to compulsive craving (Robinson and Berridge 2003). Different accounts of the role of dopamine in addiction have been proposed. According to the hedonic account, dopamine is directly involved in the pleasurable experience evoked by reward (e.g., Koob and Le Moal 1997), and as tolerance to the drug increases, the homeostatic response to the drug leads to an unpleasant withdrawal state. However, drug-seeking persists well after the pleasurable effects of using have subsided; this undermines a purely hedonic account of substance abuse.

Another account is dopamine's role in learning associations between predictive cues, actions, and reward delivery. According to reinforcement learning theory, learned reward predictions are used for optimal action selection (Sutton and Barto 1998). Attended stimuli associated with reward during a training phase may therefore continue to capture attention in the test phase because the act of attending to them reliably preceded reward delivery. Everitt et al. (2001) suggested that exposure to addictive drugs can result in a transition from action-outcome learning in the ventral striatum to the formation of automatic stimulus-response habits mediated by the dorsal striatum. Robinson and Berridge (2003) argue, however, that over learned, automatic habits (e.g., tying your shoes) do not generally give rise to compulsive motivation. Some additional mechanism is needed.

Berridge and Robinson (1998); Robinson and Berridge (2003) review evidence for their proposal that the release of dopamine that accompanies the receipt of reward increases the *incentive salience* of reward-related stimuli. Incentive salience produces a motivation to want the reward-associated stimuli; in addition, wanting persists even when a drug no longer produces pleasure. This mechanism of motivated behavior is typically adaptive: animals are motivated to seek rewarding stimuli. However, when incentive salience overrides top-down intentions, this outcome can become debilitating.

Several of the studies reviewed above provide evidence that task-relevant stimuli that predict reward elicit enhanced behavioral and neural responses (Hickey et al. 2010a, b; Navalpakkam et al. 2010; Peck et al. 2009; Raymond and O'Brien 2009; Serences 2008). The present experiments demonstrate that reward learning can imbue stimuli with value that can override top-down intention and give rise to suboptimal behavior. Together the results of these experiments extend a growing understanding of how reward associations strongly influence attention. The attentional priority accorded to reward-related stimuli expands the landscape of attentional control beyond the well known stimulus-driven (Theeuwes 1992; Yantis and Jonides 1984) and goal-directed (Folk et al. 1992) modes of attentional capture.

The results reported in this chapter are broadly consistent with the notion that the representation of value and attentional priority are critically linked. Studies by Serences (2008), Serences and Saproo (2010), Shuler and Bear (2006) demonstrated that reward-associated features are represented more robustly in early visual cortex. It is also well established that stimuli that predict reward come to evoke the striatal response formerly associated with the reward itself (e.g., Schultz et al. 1997; Hollerman et al. 1998). Our findings suggest that such activity biases attention to reward-related stimulus features which have been imbued with incentive salience. Experiment 7 demonstrates that recent reward history can potentiate value-driven capture, which is consistent with this notion.

Because attention determines the content of perceptual experience and the resulting awareness of one's surroundings, it contributes to optimal behavior in all spheres of life. Disorders of attention and cognitive control accompany a variety of clinical syndromes. It is possible that some form of value-driven attentional capture plays a role in these syndromes, including drug addiction (Field and Cox 2008; Garavan and Hester 2007; Robinson and Berridge 2008), attention-deficit/hyperactivity disorder (Bush 2010), obsessive-compulsive disorder (Sheppard et al. 2010), and obesity (Davis 2010). These conditions tend to co-occur (Davis 2010; Sheppard et al. 2010), and correlations with individual differences in working memory capacity and impulsivity suggest that there may be common underlying mechanisms that make some individuals more susceptible to value-driven attentional capture and the disorders to which it may contribute.

Acknowledgments

We thank H. Egeth, J. Flombaum, L. Gmeindl, P. Holland, D. E. Meyer, and J. Serences for fruitful discussions and suggestions. The experiments reported here were supported by US National Institutes of Health grant R01-DA013165 to S.Y.

References

- Anderson BA, Folk CL. Variations in the magnitude of attentional capture: Testing a two-process model. *Attention, Perception, & Psychophysics*. 2010; 72:342–352.
- Anderson BA, Laurent PA, Yantis S. Value-driven attentional capture. *Proceedings of the National Academy of Sciences U S A*. 2011a; 108:10367–10371.
- Anderson BA, Laurent PA, Yantis S. Learned value magnifies salience-based attentional capture. *PLoS One*. 2011b; 6(11):e27926. [PubMed: 22132170]
- Anderson BA, Laurent PA, Yantis S. Generalization of value-based attentional priority. *Visual Cognition*. 2012; 20:647–658.
- Belopolsky AV, Schreij D, Theeuwes J. What is top-down about contingent capture? *Attention, Perception, & Psychophysics*. 2010; 72:326–341.
- Berridge KC, Robinson TE. What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Research Reviews*. 1998; 28:309–369. [PubMed: 9858756]
- Bisley JW, Goldberg ME. Attention, intention, and priority in the parietal lobe. *Annual Review of Neuroscience*. 2010; 33:1–21.
- Braver TS, Cole MW, Yarkoni T. Vive les differences! Individual variation in neural mechanisms of executive control. *Current Opinion in Neurobiology*. 2010; 20:242–250. [PubMed: 20381337]
- Bundesen C. A theory of visual attention. *Psychological Review*. 1990; 97:523–547. [PubMed: 2247540]
- Bush G. Attention-deficit/hyperactivity disorder and attention networks. *Neuropsychopharmacology*. 2010; 35:278–300. [PubMed: 19759528]
- Christ SE, Abrams RA. Abrupt onsets cannot be ignored. *Psychonomic Bulletin & Review*. 2006; 13:875–880. [PubMed: 17328388]
- Corbetta M, Shulman GL. Control of goal-directed and stimulus driven attention in the brain. *Nature Reviews Neuroscience*. 2002; 3:201–215.
- Davis C. Attention-deficit/hyperactivity disorder: associations with overeating and obesity. *Current Psychiatry Reports*. 2010; 12:389–395. [PubMed: 20632134]
- Della Libera C, Chelazzi L. Visual selective attention and the effects of monetary reward. *Psychological Science*. 2006; 17:222–227. [PubMed: 16507062]
- Della Libera C, Chelazzi L. Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*. 2009; 20:778–784. [PubMed: 19422618]
- Dickman SJ, Meyer DE. Impulsivity and speed-accuracy tradeoffs in information processing. *Journal of Personality & Social Psychology*. 1988; 54:274–290. [PubMed: 3346814]
- Duncan J, Ward R, Shapiro K. Direct measurement of attentional dwell time in human vision. *Nature*. 1994; 369:313–315. [PubMed: 8183369]
- Egeth HE, Yantis S. Visual attention: Control, representation, and time course. *Annual Review of Psychology*. 1997; 48:269–297.
- Eriksen BA, Eriksen CW. Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*. 1974; 16:143–149.
- Everitt BJ, Dickinson A, Robbins TW. The neuropsychological basis of addictive behaviour. *Brain Research Reviews*. 2001; 36:129–138. [PubMed: 11690609]
- Field M, Cox WM. Attentional bias in addictive behaviors: A review of its development, causes, and consequences. *Drug and Alcohol Dependence*. 2008; 97:1–20. [PubMed: 18479844]
- Folk CL, Remington RW. Selectivity in distraction by irrelevant featural singletons: Evidence for two forms of attentional capture. *Journal of Experimental Psychology: Human Perception & Performance*. 1998; 24:847–858. [PubMed: 9627420]
- Folk CL, Remington RW, Johnston JC. Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception & Performance*. 1992; 18:1030–1044. [PubMed: 1431742]
- Folk CL, Remington RW, Wu SC. Additivity of abrupt onset effects supports nonspatial distraction, not the capture of spatial attention. *Attention Perception & Psychophysics*. 2009; 71:308–313.

- Fukuda K, Vogel EK. Human variation in overriding attentional capture. *Journal of Neuroscience*. 2009; 29:8726–8733. [PubMed: 19587279]
- Fukuda K, Vogel EK. Individual differences in recovery time from attentional capture. *Psychological Science*. 2011; 22:361–368. [PubMed: 21310945]
- Garavan H, Hester R. The role of cognitive control in cocaine dependence. *Neuropsychological Review*. 2007; 17:337–345.
- Groman SM, James AS, Jentsch JD. Poor response inhibition: at the nexus between substance abuse and attention deficit/hyperactivity disorder. *Neuroscience & Biobehavioral Reviews*. 2008; 33:690–698. [PubMed: 18789354]
- Hickey C, Chelazzi L, Theeuwes J. Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience*. 2010a; 30:11096–11103. [PubMed: 20720117]
- Hickey C, Chelazzi L, Theeuwes J. Reward guides vision when it's your thing: Trait reward-seeking in reward-mediated visual priming. *PLOS One*. 2010b; 5:e14087. [PubMed: 21124893]
- Hollerman JR, Tremblay L, Schultz W. Influence of reward expectation on behavior-related neuronal activity in primate striatum. *Journal of Neurophysiology*. 1998; 80:947–963. [PubMed: 9705481]
- Itti L, Koch C. Computational modelling of visual attention. *Nature Reviews Neuroscience*. 2001; 2:194–203.
- Koob GF, Le Moal M. Drug abuse: Hedonic homeostatic dysregulation. *Science*. 1997; 278:52–58. [PubMed: 9311926]
- Krebs RM, Boehler CN, Woldorff MG. The influence of reward associations on conflict processing in the Stroop task. *Cognition*. 2010; 117:341–347. [PubMed: 20864094]
- Kyllingsbaek S, Schneider WX, Bundesen C. Automatic attraction of attention to former targets in visual displays of letters. *Perception & Psychophysics*. 2001; 63:85–98. [PubMed: 11304019]
- Lien MC, Ruthruff E, Johnston JV. Attentional capture with rapidly changing attentional control settings. *Journal of Experimental Psychology: Human Perception & Performance*. 2010; 36:1–16. [PubMed: 20121291]
- Lin JY, Murray SO, Boynton GM. Capture of attention to threatening stimuli without perceptual awareness. *Current Biology*. 2009; 19:1118–1122. [PubMed: 19523828]
- Moran J, Desimone R. Selective attention gates visual processing in the extrastriate cortex. *Science*. 1985; 229:782–784. [PubMed: 4023713]
- McClure SM, Berns GS, Montague PR. Temporal prediction errors in a passive learning task activate human striatum. *Neuron*. 2003; 38:339–346. [PubMed: 12718866]
- Navalpakkam V, Koch C, Rangel A, Perona P. Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences U S A*. 2010; 107:5232–5237.
- O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ. Temporal difference models and reward-related learning in the human brain. *Neuron*. 2003; 38:329–337. [PubMed: 12718865]
- Parkhurst D, Law K, Niebur E. Modeling the role of salience in the allocation of overt visual attention. *Vision Research*. 2002; 42:107–123. [PubMed: 11804636]
- Pashler, H., editor. *Attention*. London: Psychology Press; 1998.
- Patton JH, Stanford MS, Barratt ES. Factor structure of the barratt impulsiveness scale. *Journal of Clinical Psychology*. 1995; 51:768–774. [PubMed: 8778124]
- Peck CJ, Jangraw DC, Suzuki M, Efem R, Gottlieb J. Reward modulates attention independently of action value in posterior parietal cortex. *Journal of Neuroscience*. 2009; 29:11182–11191. [PubMed: 19741125]
- Pessoa L, Engelmann JB. Embedding reward signals into perception and cognition. *Frontiers in Neuroscience*. 2010; 4(17)10.3389/fnins.2010.00017
- Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. *Nature*. 1999; 400:233–238. [PubMed: 10421364]
- Raymond JE, O'Brien JL. Selective visual attention and motivation: The consequences of value learning in an attentional blink task. *Psychological Science*. 2009; 20:981–988. [PubMed: 19549080]
- Rescorla RA. Partial reinforcement reduces the associative change produced by nonreinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*. 1999; 25:403–414.

- Robinson TE, Berridge KC. Addiction. *Annual Review of Psychology*. 2003; 54:25–53.
- Robinson TE, Berridge KC. The incentive sensitization theory of addiction: some current issues. *Philosophical Transactions of the Royal Society: B Biological Sciences*. 2008; 363:3137–3146.
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
- Serences JT. Value-based modulations in human visual cortex. *Neuron*. 2008; 60:1169–1181. [PubMed: 19109919]
- Serences JT, Saproo S. Population response profiles in early visual cortex are biased in favor of more valuable stimuli. *Journal of Neurophysiology*. 2010; 104:76–87. [PubMed: 20410360]
- Sheppard B, Chavira D, Azzam A, Grados MA, Umaa P, Garrido P, Mathews CA. ADHD prevalence and association with hoarding behaviors in childhood onset OCD. *Depression & Anxiety*. 2010; 27:667–674. [PubMed: 20583294]
- Shiffrin RM, Schneider W. Controlled and automatic human information processing II: Perceptual learning, automatic attending, and general theory. *Psychological Review*. 1977; 84:127–190.
- Shuler MG, Bear MF. Reward timing in the primary visual cortex. *Science*. 2006; 311:1606–1609. [PubMed: 16543459]
- Simen P, Contreras D, Buck C, Hu P, Holmes P, Cohen JD. Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception & Performance*. 2009; 35:1865–1897. [PubMed: 19968441]
- Sugrue LP, Corrado GS, Newsome WT. Choosing the greater of two goods: neural currencies for valuation and decision making. *Nature Reviews Neuroscience*. 2005; 6:363–375.
- Sutton, RS.; Barto, AG. *Reinforcement learning: An introduction*. Cambridge: MIT Press; 1998.
- Theeuwes J. Perceptual selectivity for color and form. *Perception & Psychophysics*. 1992; 51:599–606. [PubMed: 1620571]
- Theeuwes J. Top-down and bottom-up control of visual selection. *Acta Psychologica*. 2010; 135:77–99. [PubMed: 20507828]
- Theeuwes J, Godijn R. Irrelevant singletons capture attention: Evidence from inhibition of return. *Perception & Psychophysics*. 2002; 64:764–770. [PubMed: 12201335]
- Yantis S. Stimulus-driven attentional capture. *Current Directions in Psychological Science*. 1993; 2:156–161.
- Yantis, S. Goal-directed and stimulus-driven determinants of attentional control. In: Monsell, S.; Driver, J., editors. *Attention and performance*. Vol. 18. Cambridge: MIT Press; 2000. p. 73–103.
- Yantis S. Neural basis of selective attention: Cortical sources and targets of attentional modulation. *Current Directions in Psychological Science*. 2008; 17:86–90. [PubMed: 19444327]
- Yantis S, Hillstrom AP. Stimulus-driven attentional capture: Evidence from equiluminant visual objects. *Journal of Experimental Psychology: Human Perception & Performance*. 1994; 20:95–107. [PubMed: 8133227]
- Yantis S, Jonides J. Abrupt visual onsets and selective attention: Evidence from visual search. *Journal of Experimental Psychology: Human Perception & Performance*. 1984; 10:350–374.
- Yantis S, Jonides J. Abrupt visual onsets and selective attention: Voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception & Performance*. 1990; 16:121–134. [PubMed: 2137514]

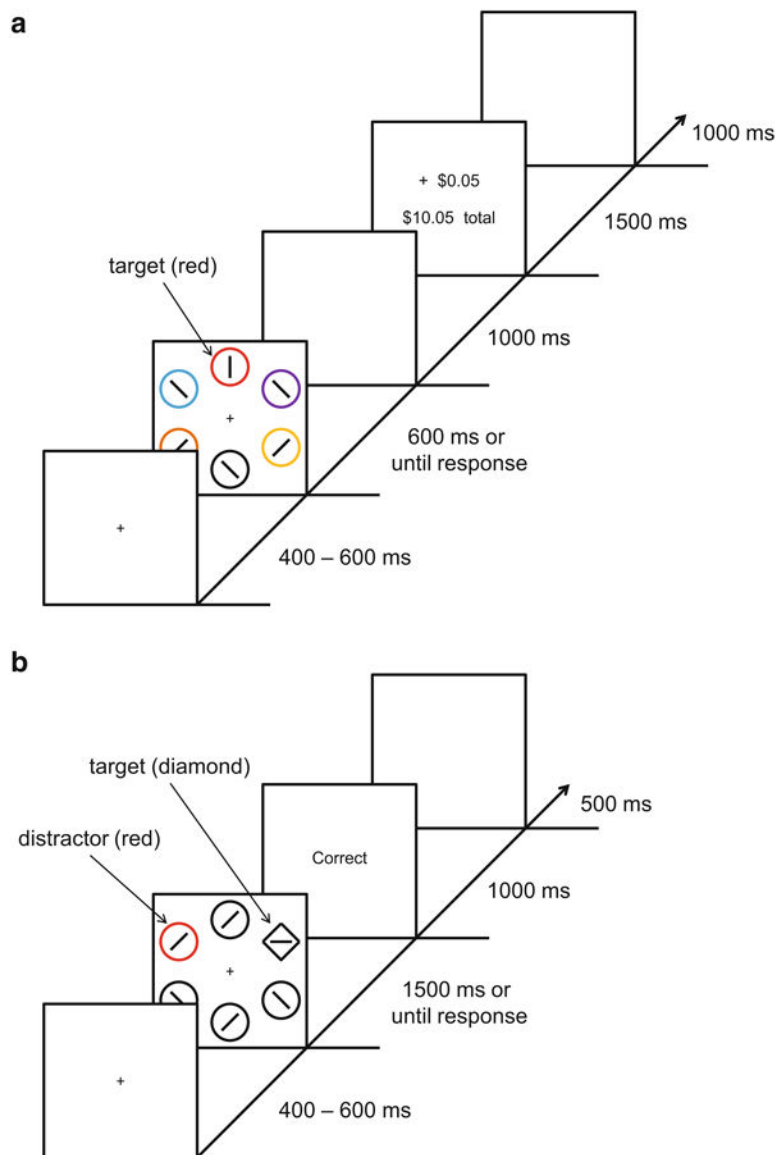


Fig. 1. Sequence of trial events in Experiment 1. **a** During training, an array of six *circles* each rendered in a different color appeared on a *black* background; one of these was either *red* or *green*. The participant determined the orientation of the line segment within the *red* or *green* circle and pressed a corresponding button. A subsequent feedback display indicated the amount of reward obtained on that trial and the cumulative reward earned so far in the experiment. The high-reward target color was followed on 80 % of the trials by a high reward (5¢) and on 20 % of the trials by a low reward (1¢); this mapping was reversed for the other target color. **b** At test, participants were to determine the orientation of the line segment within the shape singleton (*diamond* among *circles* or *circle* among *diamonds*, unpredictably). On half the trials, all the items were *white*. On one-quarter of the trials one of the nontargets was *red* and on one-quarter one of the nontargets was *green*. No reward was delivered during the test phase

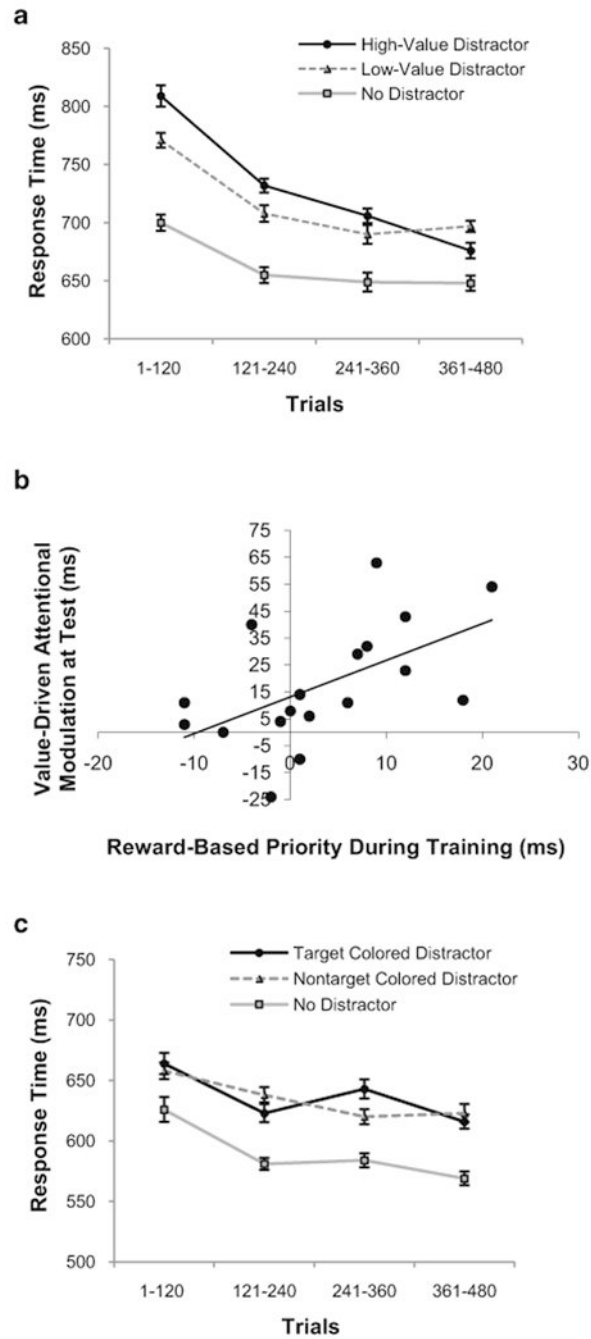


Fig. 2. Results for experiments 1 and 2. **a** Mean response time (in milliseconds) for each condition in the test phase of Experiment 1. **b** Correlation between reward-based performance during training (defined as the mean RT difference to low- and high-reward targets), and the magnitude of value-driven attentional modulation at test (defined as the mean difference in RT on trials containing high- and low-value distractors, respectively) in Experiment 1. **c** Mean response time (in milliseconds) for each condition in the test phase of Experiment 2. Error bars represent \pm within-subjects SEM

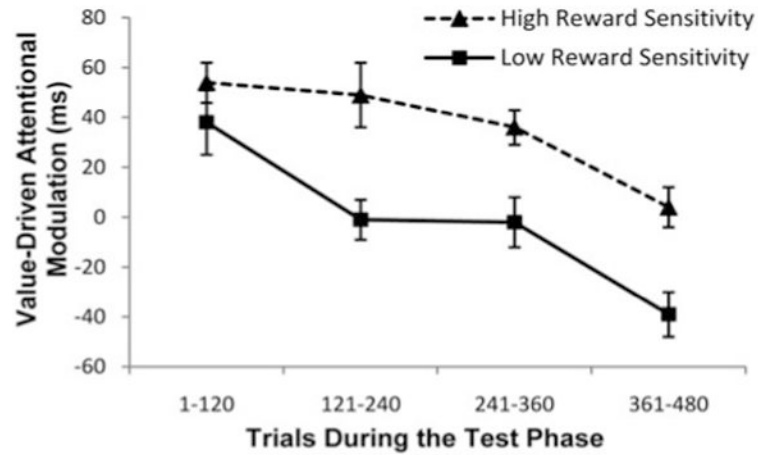
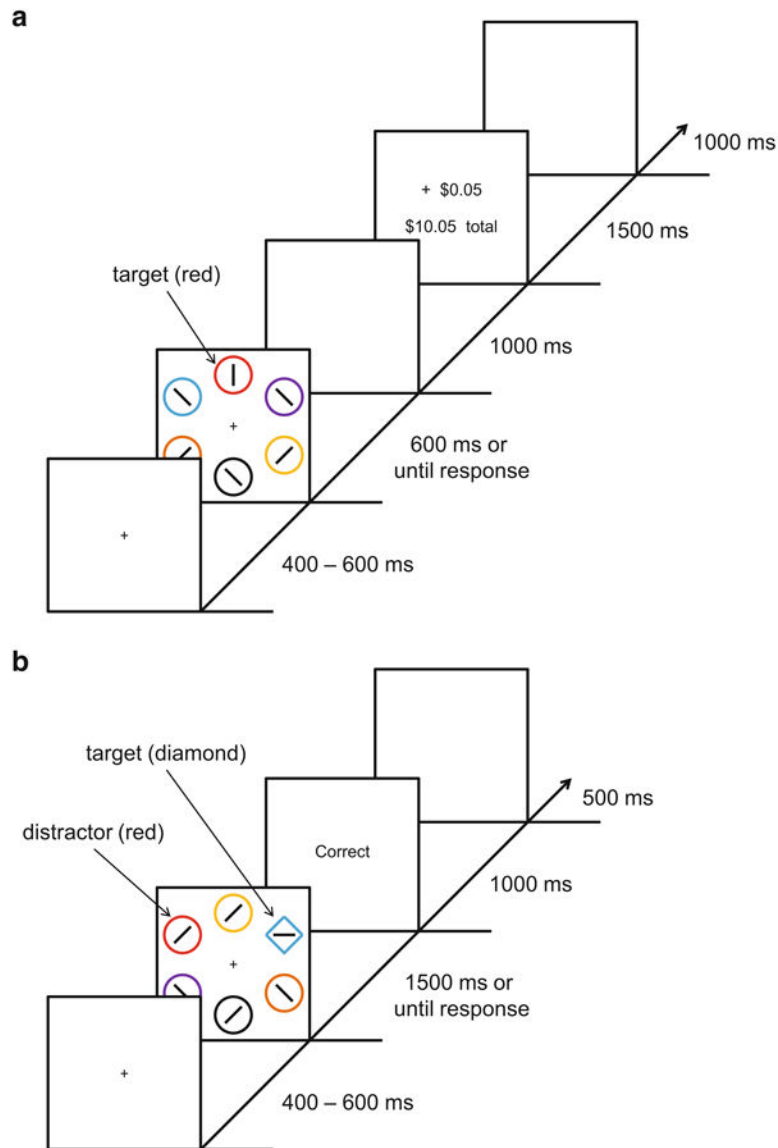


Fig. 3.

Mean difference in response time (in milliseconds) for high- and low-value distractors over the course of the test phase of Experiment 1, plotted separately for two subgroups of participants: those who exhibited the largest difference in RT to high- and low-reward targets in the training phase and those who exhibited the smallest difference (median split). The effect of reward at test is larger and persists longer for participants who were more sensitive to reward during training

**Fig. 4.**

Sequence of trial events in experiment 3. **a** During training, an array of six *circles* each rendered in a different color appeared on each trial; one of these was either *red* or *green*. The participant determined the orientation of the line segment within the *red* or *green* circle and pressed a corresponding button. Text indicating monetary reward feedback was presented after each correct response. **b** During the test phase, participants searched for a unique shape (*diamond* among *circles* or *circle* among *diamonds*, unpredictably); no reward feedback was provided. The items were rendered in six different colors. On 25 % of the trials, one of the nontarget items was red and on 25 % of the trials, one of the nontarget items was green; on the remaining trails none of the items were *red* or *green*




Fig. 5. Scatter plot shows the magnitude of value-driven attentional capture (RT in the presence of a high-value distractor minus RT when neither distractor was present) vs. visual working memory capacity in Experiment 3 (long training, *diamonds, grey line*) and Experiment 5 (short training, *circles, black line*). Best-fitting regression lines are shown

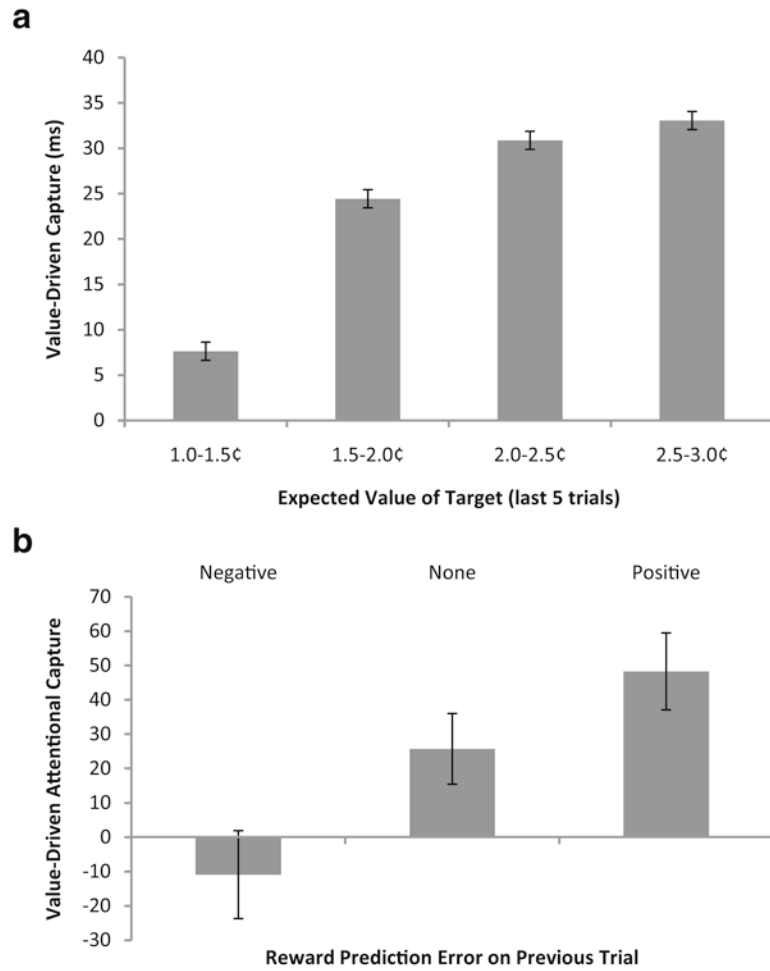


Fig. 6. Results for Experiment 7. **a** Value-driven attentional capture (defined here as the mean RT difference between distractor present and distractor absent trials) as a function of the average reward that the current target has received over the last five trials on which it was followed by reward. **b** Value-driven attentional capture as a function of the reward prediction error on the previous trial

Table 1

Error rates by distractor condition for Experiments 1 and 2

	Distractor condition in Experiment 1		Distractor condition in Experiment 2	
	Low-value	High-value	None	Target colored
None	0.09 (0.003)	0.10 (0.004)	0.10 (0.005)	0.11 (0.003)
Target colored			0.13 (0.005)	0.13 (0.005)

The error terms, in parentheses, reflect the within-subjects SEM

Table 2

Mean response time (in milliseconds) and error rate, respectively, in the test phase of Experiments 3 and 5 in which reward was delivered for each of three training conditions: Experiment 3, long training (1008 trials) with low and high reward of 2¢ or 10¢ per trial; Experiment 5, brief training (240 trials) with rewards of 2¢ and 10¢ per trial, and Experiment 5, brief training followed by a delay of 4–21 days

Training phase	Distractor condition in the test phase		
	None	Low-value	High-value
1008 trials	665 (2.8)	673 (2.8)	681 (2.6)
	0.11 (0.004)	0.10 (0.004)	0.11 (0.004)
240 trials	667 (2.0)	675 (3.0)	682 (2.9)
	0.12 (0.005)	0.12 (0.006)	0.12 (0.006)
4–21 days ago	614 (1.8)	624 (2.7)	630 (3.3)
	0.06 (0.004)	0.07 (0.006)	0.08 (0.005)

The error terms, in parentheses, reflect the within-subjects SEM

Table 3
Mean response time (in milliseconds) and error rate, respectively, in the test phase of
Experiment 4 in which no reward was delivered

Training phase	Distractor condition in the test phase		
	None	Red	Green
1008 trials (Unrewarded)	602 (3.9) 0.14 (0.004)	606 (2.1) 0.17 (0.006)	593 (3.9) 0.15 (0.005)
	0.14 (0.004)	0.17 (0.006)	0.15 (0.005)

The error terms, in parentheses, reflect the within-subjects SEM