Generalization of value-based attentional priority

Brian A. Anderson, **Patryk A. Laurent**, and **Steven Yantis**
Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD, USA

Abstract

Attention is the mechanism by which important or salient stimuli are selected for perceptual and cognitive processing. Which stimuli are attended has important implications for effective goal-directed behaviour, survival, and well-being. A growing body of evidence suggests that reward-predicting stimuli capture attention involuntarily. In previous studies, value-based attentional priority has been observed only when the formerly reward-related stimuli themselves were presented as targets or distractors. Here we show that stimulus–reward associations learned in one task generalize to different stimuli that share a defining feature (colour) in another task. Our results reveal a broad and flexible role for reward learning in modulating attentional priority.

Keywords

The representational capacity of perception is limited. Attention plays a critical role in an organism's survival by selecting the sensory input that is required to identify objects in a scene, guide appropriate actions, and bring about rewarding outcomes (e.g., nourishment). By selectively attending to stimuli that predict the delivery of reward, an organism increases the likelihood that opportunities to acquire valuable resources will reach awareness and become available for action.

A growing body of research has revealed that stimuli associated with the delivery of reward have high attentional priority (Della Libera & Chelazzi, 2006, 2009; Della Libera, Perlato, & Chelazzi, 2011; Hickey, Chelazzi, & Theeuwes, 2010a, 2010b, 2011; Krebs, Boehler, & Woldorff, 2010; Peck, Jangraw, Suzuki, Efem, & Gottlieb, 2009; Raymond & O'Brien, 2009). Recently, we reported that stimuli associated with reward through learning (i.e., high-value stimuli) continue to capture attention involuntarily, even when they are not physically salient, no longer predict reward, and are irrelevant to the task (Anderson, Laurent, & Yantis, 2011b). We term this phenomenon *value-driven attentional capture*, a form of attentional control in which learned value has a direct influence on attentional priority.

In our previous demonstrations of value-driven attentional capture (Anderson et al., 2011a, 2011b), and in studies of the influence of reward learning on attention generally (Della Libera & Chelazzi, 2006, 2009; Della Libera et al., 2011; Hickey et al., 2010a, 2010b, 2011; Krebs et al., 2010; Peck et al., 2009; Raymond & O'Brien, 2009), the influence of reward value on attention has been observed only using the very same stimuli that were associated with reward during learning. Thus, the extent to which value-based attentional priority transfers across different stimuli is unknown. One possibility is that value-based attentional

priority will generalize to stimuli that share a defining feature with a stimulus that has acquired learned value. For example, a novel fruit that has a colour that signifies ripeness might capture attention more readily than one with a colour that appears unripe. Another possibility, however, is that the attentional priority that accompanies reward learning is restricted in scope, limited only to the same stimuli that were rewarded previously. Both mechanisms of value-based attentional priority are theoretically plausible: The former mechanism would exploit previous learning in the pursuit of rewards in novel contexts, but at the expense of suboptimal overgeneralization that would be mitigated by the latter mechanism.

Here we test between these two competing accounts of the learning that gives rise to value-based attentional priority. Participants learned probabilistic associations between coloured circles and reward-magnitude, as in the training phase of Anderson et al. (2011a, 2011b) and as described later. Following this training, the participants engaged in a flankers task that required spatially focused attention on a central letter (Eriksen & Eriksen, 1974); the irrelevant letter flankers were in some conditions rendered in the colour of formerly reward-predictive targets. This provided a means to determine whether value-based attentional priority generalized across different stimuli (from coloured geometric shapes to coloured letters) and across different cognitive contexts (from visual search to a focused-attention task).

EXPERIMENT 1

Experiment 1 consisted of two phases, a training phase and a test phase. In the training phase (Figure 1a), each of two target colours was associated with value via reward feedback. In the test phase, participants engaged in a flankers task during which no reward feedback was provided (Figure 1b). The to-be-ignored flanker letters were either compatible or incompatible with the response required by the central target letter, and they could be the same colour as the formerly high-reward target (high-value flanker), the formerly low-reward target (low-value flanker), or a former nontarget item (value-neutral flanker). To the extent that the irrelevant flankers are processed, compatible flankers will facilitate the correct response and incompatible flankers will compete with the correct response (Eriksen & Eriksen, 1974). If value-based attentional priority generalizes to different stimuli that share a defining feature, then the magnitude of the flanker compatibility effect should differ for high-, low-, and neutral-value flankers. If value-based attentional priority does not generalize across different stimuli, then the flanker colour should have no effect on performance.

Method

Participants—Twenty-one participants were recruited from the Johns Hopkins University community. All were screened for normal or corrected-to-normal visual acuity and colour vision.

Apparatus and stimuli—A Mac Mini equipped with Matlab software and Psychophysics Toolbox extensions was used to present the stimuli on a Dell P991 monitor. The participants viewed the monitor from a distance of approximately 50 cm in a dimly lit room.

Training phase—Each trial in the training phase consisted of a fixation display, a search display, and a feedback display. The fixation display consisted of a white fixation cross (0.5° × 0.5° visual angle) presented in the centre of the screen against a black background, and the search display consisted of the fixation cross surrounded by six circles (2.3° × 2.3° visual angle) placed at equal intervals along an imaginary circle with a 5° radius. The six circles in the search display each had a different colour (red, green, blue, cyan, pink, orange, yellow,

or white). Targets were defined to be stimuli of either of two colours; exactly one of those colours was presented on each trial. Inside the target shape, a white line segment was oriented either vertically or horizontally, and inside each of the nontarget shapes, a white line segment was tilted at 45° to the left or to the right (randomly selected for each element). The feedback display informed participants of the reward earned on the trial, as well as total reward accumulated to that point. The selection of the two target colours from the set {red, green, blue} was manipulated as a between-subjects variable; each colour served as the high-reward and low-reward colour in one of three conditions (red and blue, green and red, and blue and green, respectively). The third colour (e.g., green for participants who searched for red and blue as targets) was always a nontarget colour during training.

Test phase—Trials in the test phase consisted of a fixation display, a flanker display, and a feedback display. The fixation display was identical to the fixation display in the training phase. The flanker display contained a white target letter presented in the centre of the screen (0.8° × 1.4° visual angle), flanked to the left and right by identical red, green, or blue letters of equal size (1.5° centre-to-centre). The central letter and flanking letters were never the same letter. The letters used for the target and flankers were "A", "B", "X", and "Y". The feedback display only informed participants if their previous response was incorrect.

Design and procedure—The experiment consisted of 240 training trials followed by 480 test trials. During the training phase, target identity and target location were fully crossed and counterbalanced, and trials were presented in a random order. During the test phase, target identity, flanker colour, and flanker compatibility were fully crossed and counterbalanced, and trials were presented in a random order.

Correct responses in the training phase were followed by visual feedback indicating monetary reward. High-reward targets were followed by high-reward feedback ($0.10) on 80% of trials and low-reward feedback ($0.02) on the remaining 20%; for low-reward targets, the percentages were reversed. One-third of the participants experienced each of the three target colour conditions. No reward feedback was provided during the test phase. Upon completion of the experiment, participants were given the cumulative reward they had earned (mean = $13.24).

Each trial began with the presentation of the fixation display for a randomly varying interval of 400, 500, or 600 ms. The search or flanker display then appeared and remained on screen until a response was made or the trial timed out. Trials timed out after 800 ms in the training phase and 1200 ms in the test phase. In the training phase, the search display was followed by a blank screen for 1000 ms, the reward feedback display for 1500 ms, and a 1000 ms intertrial interval (ITI). In the test phase, the flanker display was followed by a 1000 ms (nonreward) feedback display only if the participant had responded incorrectly, and then by a 500 ms ITI.

Participants made a forced-choice target identification by pressing the "z" and the "m" keys for the vertically and horizontally orientated targets in the training phase, respectively. In the test phase, participants responded with the "z" key to "X" and "Y" targets and with the "m" key to "A" and "B" targets. If the trial timed out, the computer emitted a 500 ms 1000 Hz feedback tone. Only correct responses were included in the analysis, and all RTs more than three standard deviations above or below the mean of their respective condition for each participant were excluded from the analysis.

Results and discussion

Trials during the training phase were categorized as containing either a high-reward target or a low-reward target. Response times (RTs) were slightly faster and more accurate to high-

reward targets (Table 1), although neither of these comparisons reached conventional levels of statistical significance, mean RT difference = 7.4 ms, $t(20) = 1.51$, $p = .146$, and mean error difference = 2.2%, $t(20) = 1.92$, $p = .069$, mirroring previous findings using this paradigm (Anderson et al., 2011a). Data from the test phase, however, were of primary interest.

Trials during the test phase were categorized as containing a high-value flanker colour, a low-value flanker colour, or a value-neutral flanker colour. The compatibility effect (i.e., the difference in RT on trials containing compatible and incompatible flankers) for each of these trial types on both RT and error rate was computed for each participant. Target colour assignment did not interact with the value of the flankers on RT compatibility effects, $F(4, 36) = 1.02$, $p = .412$, so further analyses collapse across target colour. A repeated measures analysis of variance (ANOVA) on RT revealed a main effect of flanker value (Figure 2a), $F(2, 40) = 3.86$, $p = .029$, $\eta_p^2 = .162$, a main effect of flanker compatibility, $F(1, 20) = 78.61$, $p < .001$, $\eta_p^2 = .797$, and a value by compatibility interaction, $F(2, 40) = 3.73$, $p = .033$, $\eta_p^2 = .157$. We report within-subjects standard errors in this paper (Loftus & Masson, 1994). Focused contrasts revealed that the compatibility effect for high-value flankers was significantly greater than that for low-value flankers (Figure 2b), $t(20) = 2.59$, $p = .017$, $d = .57$, but the compatibility effect for the value-neutral flankers did not differ from that for either the high-value or low-value flankers, $t(20) = 1.69$, $p = .106$, and $t(20) = -0.85$, $p = .405$, respectively. An ANOVA on error rate revealed only a main effect of compatibility (Table 2), $F(1, 20) = 36.59$, $p < .001$, $\eta_p^2 = .647$.

The difference in the flanker compatibility effect between high- and low-value flankers was slightly larger in the first half of the test phase compared to the second half, though not significantly so, mean difference = 7 ms, $t(20) = 0.63$, $p = .533$. This is consistent with our previous findings indicating that value-based attentional priority is persistent and can be evident even a week after learning (Anderson et al., 2011b).

The results of Experiment 1 reveal that learned associations between stimuli and reward have an involuntary influence on attentional priority that transfers across different stimuli and task contexts. Participants learned to associate the delivery of a large reward with one colour and the delivery of a smaller reward with another colour in the training phase. In the test phase, this learning resulted in larger compatibility effects for flanking letters rendered in the high-value colour compared to those rendered in the low-value colour, demonstrating that the high-value flankers exhibited increased attentional priority as a function of prior learning.

EXPERIMENT 2

The observed difference in the compatibility effect for high- and low-value flankers in Experiment 1 can only be explained in terms of a difference in learned value, because this is all that varied between these two conditions. However, it is perhaps surprising that the value-neutral flankers did not produce the smallest compatibility effect, as their colour had not been formerly associated with reward-predicting targets—that is, they were presumably less valuable than even the low-value distractors. One explanation for this result is that former nontarget colours, which were previously ignored, receive higher attentional priority than former target colours in a new task, independently of their reward value. This might arise from a novel orienting response—a bias to attend to less familiar stimuli (e.g., Horstmann & Ansorge, 2006; Johnston, Hawley, Plewe, Elliott, & DeWitt, 1990; Neo & Chua, 2006). In Experiment 2, we tested this possibility by removing the reward feedback component of the training phase, thereby isolating the effect of target colour history. We

compared the flanker compatibility effect for the two (familiar) former target colours to the compatibility effect for the (comparatively novel) former nontarget colour.

Method

Participants—Eighteen new participants were recruited from the Johns Hopkins University community. All were screened for normal or corrected-to-normal visual acuity and colour vision.

Apparatus and stimuli—The apparatus and stimuli were identical to Experiment 1, with the exception that the reward feedback display was removed from the training phase. If participants responded incorrectly, error feedback was inserted between the search display and the blank ITI, as in the test phase.

Design and procedure—The design and procedure were identical to Experiment 1, with the exception that participants received no task-contingent reward and were instead compensated with either $10 or course credit for participation.

Results and discussion

Trials during the test phase were categorized as containing either a familiar (former target colour) or novel (former nontarget colour) flanker. The compatibility effect for these two trial types on both RT and error rate was measured for each participant. Target colour assignment did not interact with flanker condition on RT compatibility effects ($F<1$), so further analyses collapse across target colour. A paired-samples t-test revealed that the flanker compatibility effect on RT, shown in Figure 3, was larger for the novel flanker colour than for the familiar flanker colours, $t(17) = 4.22$, $p =.001$, $d = 1.00$. Like the effect of learned value from Experiment 1, this effect was consistent across trials and did not differ between the first and second half of the test phase, mean difference = 2 ms, $t(17) = 0.25$, $p =.804$. There was no difference in the flanker compatibility effect on error rates between the two conditions (Table 3), $t(17) = -1.21$, $p =.244$.

The results of Experiment 2 demonstrate that stimuli rendered in a former nontarget colour have higher attentional priority in a new task than stimuli rendered in a more familiar former target colour. In previous demonstrations of novel orienting, novelty is typically defined in terms of recent trial history and the effects of novelty are short-lived (e.g., Neo & Chua, 2006). The present results extend novel orienting to encompass novelty that is defined in terms of a history of attentional selection, suggesting a persistent bias to attend to stimuli about which comparatively less has been learned. More importantly, these results show that the compatibility effect in the neutral condition of Experiment 1, which was larger than that in the low-value condition, was likely magnified by the relative novelty of the neutral colour.

GENERAL DISCUSSION

Reward plays a critical role in the deployment of attention. Stimuli associated with the delivery of reward draw attention, and they come to capture attention involuntarily once their value has been learned, even when they are no longer rewarded (Anderson et al., 2011a, 2011b). In two experiments, we examined the extent to which the effects of reward learning on attentional priority generalize to newly encountered stimuli and tasks.

In Experiment 1, we found that learning to associate reward value with a stimulus defined by a particular feature (colour) resulted in a transfer of value-based attentional priority to different stimuli sharing that feature. This finding provides clear support for a general

principle: The attentional priority that accompanies reward learning is flexible and supports stimulus generalization, promoting the rapid application of former learning to newly encountered stimuli in different contexts. That the effect of reward learning generalizes across stimuli and task contexts attests to the robustness of value-based attentional priority.

Value-based attentional priority has been shown to influence both early and late components of stimulus selection. Valuable stimuli capture attention (Anderson et al., 2011b) and involuntarily draw eye movements (Anderson & Yantis, 2012) when they are not physically salient, implicating early stimulus selection. Learned value can also modulate attentional priority when selection is dominated by physical salience (Anderson et al., 2011a), suggesting a role for postselection stimulus processing as well. The flanker compatibility effect is known to arise at both early and late stages of stimulus processing (e.g., Casey et al., 2000; Mattler, 2006); the value-based modulation of performance in the present experiments could be operating at either or both of these stages.

We observed that stimuli expressing features of previously ignored (i.e., former nontarget) stimuli persistently evoke stronger flanker compatibility effects than stimuli expressing features of previously attended (but unrewarded) former targets in a new task; this indicates that former nontargets have greater attentional priority than former unrewarded targets in the test phase. This finding extends the scope of a previously documented form of novel orienting (e.g., Horstmann & Ansorge, 2006; Johnston et al., 1990; Neo & Chua, 2006), and demonstrates a bias to attend to stimuli about which less has already been learned. Thus, all other things being equal, this result suggests that previously ignored stimuli are attended more readily than familiar stimuli when they appear in a new context. Nevertheless, Experiment 1 showed that this novelty effect does not overshadow attentional priority to a feature that had been associated with high reward.

Although the attentional priority of a newly encountered stimulus is magnified when it possesses a feature previously associated with high reward, the extent to which value-based attentional priority is principally feature-based or object-based remains unclear. For example, we might have observed larger interference effects for flankers that share more than one feature in common with a previously rewarded stimulus, particularly if reward was predicted by a specific conjunction of features. The present findings make it clear, however, that the transfer of value-based attentional priority is flexible enough to occur for objects that have themselves never been rewarded.
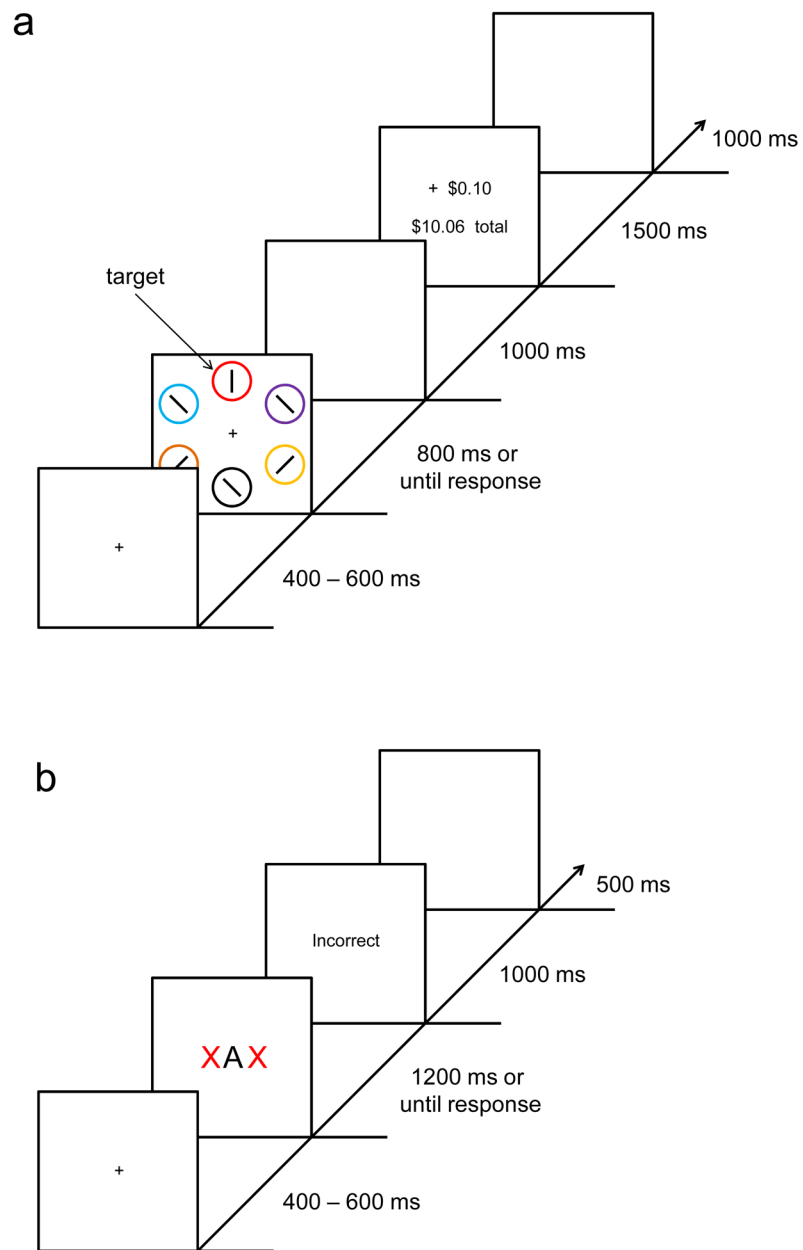
Our results contribute to a growing body of evidence that highlights a critical role for reward in the deployment of attention (Anderson et al., 2011a, 2011b; Della Libera & Chelazzi, 2006, 2009; Della Libera et al., 2011; Hickey et al., 2010a, 2010b, 2011; Krebs et al., 2010; Peck et al., 2009; Raymond & O'Brien, 2009). We show that the reward learning that underlies value-based attentional priority is flexibly applied to newly encountered stimuli and task contexts, supporting the rapid generalization of learning to facilitate the procurement of future rewards.

References

Anderson BA, Laurent PA, Yantis S. Learned value magnifies salience-based attentional capture. PLoS ONE. 2011a; 6(11):e27926.10.1371/journal.pone.0027926 [PubMed: 22132170]

Anderson BA, Laurent PA, Yantis S. Value-driven attentional capture. Proceedings of the National Academy of Sciences, USA. 2011b; 108:10367–10371.

Anderson BA, Yantis S. Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. 2012 Manuscript submitted for publication.

Casey BJ, Thomas KM, Welsh TF, Badgaiyan RD, Eccard C, Jennings JR, Crone EA. Dissociation of response conflict, attentional selection, and expectancy with functional magnetic resonance imaging. Proceedings of the National Academy of Sciences, USA. 2000; 97:8728–8733.

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]

Della Libera C, Perlato A, Chelazzi L. Dissociable effects of reward on attentional learning: From passive associations to active monitoring. PLoS ONE. 2011; 6:e19460.10.1371/journal.pone. 0019460 [PubMed: 21559388]

Eriksen BA, Eriksen CW. Effects of noise letters upon the identification of a target letter in a nonsearch task. Perception and Psychophysics. 1974; 16:143–149.

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:e14087.10.1371/journal.pone.0014087 [PubMed: 21124893]

Hickey C, Chelazzi L, Theeuwes J. Reward has a residual impact on target selection in visual search, but not on the suppression of distractors. Visual Cognition. 2011; 19:117–128.

Horstmann G, Ansorge U. Attentional capture by rare singletons. Visual Cognition. 2006; 14:295–325.

Johnston WA, Hawley KJ, Plewe SH, Elliott JMG, DeWitt MJ. Attention capture by novel stimuli. Journal of Experimental Psychology: General. 1990; 119:397–411. [PubMed: 2148574]

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]

Loftus GR, Masson MEJ. Using confidence intervals in within-subject designs. Psychonomic Bulletin and Review. 1994; 1:476–490. [PubMed: 24203555]

Mattler U. Distance and ratio effects in the flanker compatibility paradigm are due to different mechanisms. Quarterly Journal of Experimental Psychology. 2006; 59:1745–1763.

Neo G, Chua FK. Capturing focused attention. Perception and Psychophysics. 2006; 68:1286–1296. [PubMed: 17378415]

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]

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]

a

target

+ $0.10
$10.06 total

1000 ms

1500 ms

1000 ms

800 ms or
until response

400 – 600 ms

b

Incorrect

X A X

500 ms

1000 ms
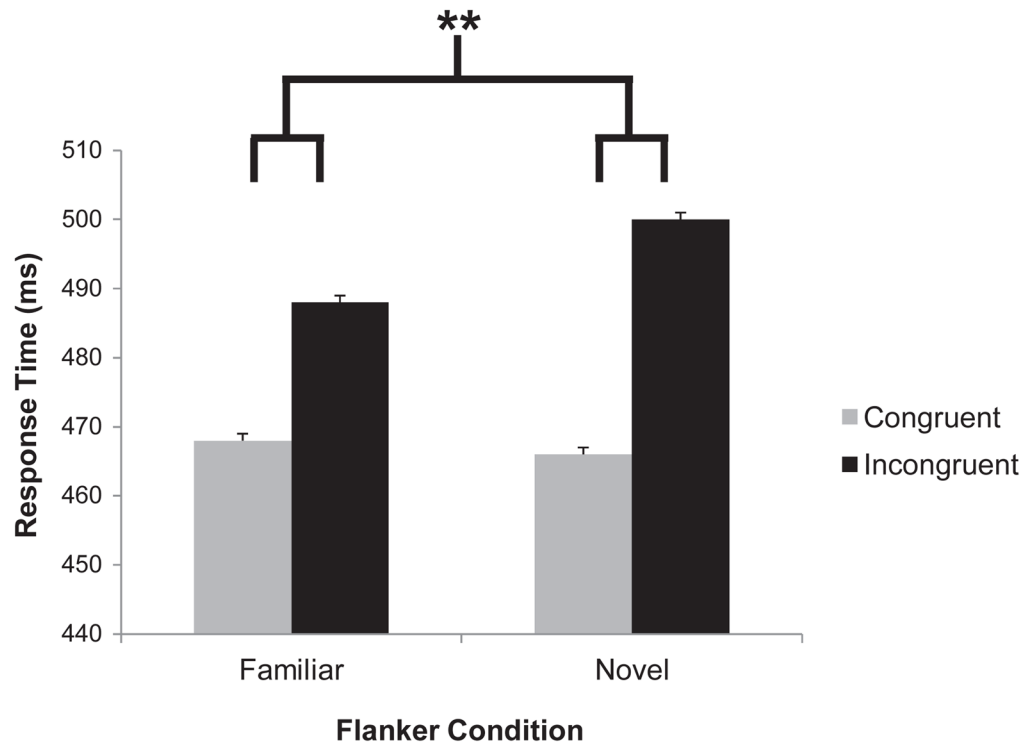
1200 ms or
until response

400 – 600 ms

Figure 1.
Sequence of events and time course for a trial during training (a) and at test (b) in
Experiment 1. To view this figure in colour, please see the online issue of the Journal.
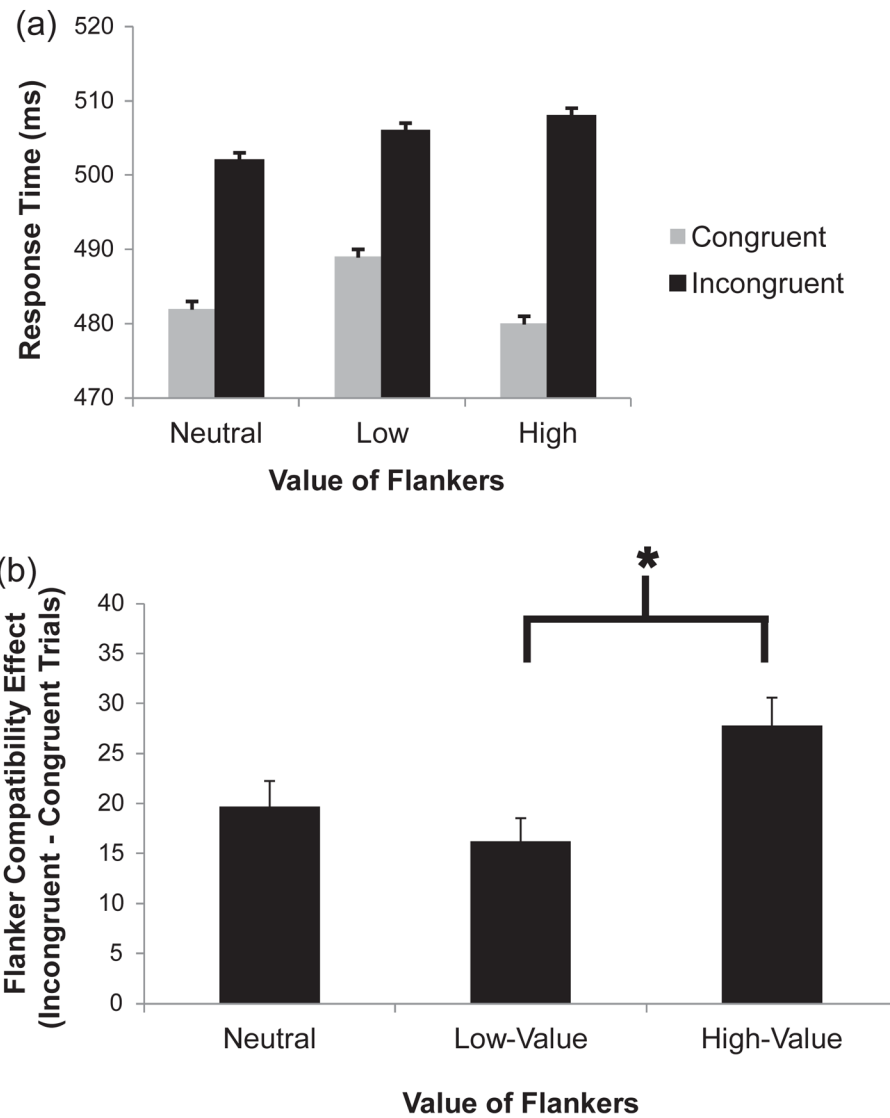
Figure 2.
Mean response time (a) and flanker compatibility effects (b) by condition in the test phase of Experiment 1 (*p<.05).

Figure 3.
Mean response time by condition in the test phase of Experiment 2 (**p<.01).

TABLE 1

Response times and error rates by condition in the training phase of Experiment 1; the error terms, in parentheses, reflect the standard error of the mean (SEM)

	Target condition	
	High-reward	**Low-reward**
Response time	548 (2.5)	555 (2.5)
Error rate	.102 (.006)	.124 (.006)

TABLE 2

Error rates by condition in the test phase of Experiment 1; the error terms, in parentheses, reflect the SEM

	Flanker condition		
	Value-neutral	**Low-value**	**High-value**
Compatible	.052 (.007)	.040 (.005)	.056 (.006)
Incompatible	.100 (.008)	.097 (.007)	.098 (.007)

TABLE 3

Error rates by condition in the test phase of Experiment 2; the error terms, in parentheses, reflect the SEM

	Flanker condition	
	Familiar	**Novel**
Compatible	.041 (.005)	.044 (.007)
Incompatible	.098 (.006)	.089 (.007)