

Current Biology

Dissociable Components of Experience-Driven Attention

Highlights

- Associative learning biases eye movements toward a reward-predictive stimulus
- This bias is evident even when participants are paid to look away from the stimulus
- Pairing an eye movement with a stimulus results in a bias to repeat this behavior
- These two biases reflect dissociable consequences of selection history on attention

Authors

Haena Kim, Brian A. Anderson

Correspondence

hannah.kim@tamu.edu

In Brief

The control of attention is widely believed to be influenced by current task goals, the physical salience of stimuli, and selection history. In this study, Kim and Anderson demonstrate that selection history reflects two distinct learning processes that have dissociable consequences for the orienting of attention.



Dissociable Components of Experience-Driven Attention

Haena Kim^{1,2,*} and Brian A. Anderson¹¹Department of Psychological & Brain Sciences, Texas A&M University, 4235 TAMU, College Station, TX 77843, USA²Lead Contact*Correspondence: hannah.kim@tamu.edu<https://doi.org/10.1016/j.cub.2019.01.030>

SUMMARY

What we pay attention to is influenced by current task goals (goal-directed attention) [1, 2], the physical salience of stimuli (stimulus-driven attention) [3–5], and selection history [6–12]. This third construct, which encompasses reward learning, aversive conditioning, and repetitive orienting behavior [12–18], is often characterized as a unitary mechanism of control that can be contrasted with the other two [12–14]. Here, we present evidence that two different learning processes underlie the influence of selection history on attention, with dissociable consequences for orienting behavior. Human observers performed an antisaccade task in which they were paid for shifting their gaze in the direction opposite one of two color-defined targets. Strikingly, such training resulted in a bias to do the opposite of what observers were motivated and paid to do, with associative learning facilitating orienting toward reward cues. On the other hand, repetitive orienting away from a target produced a bias to repeat this behavior even when it conflicted with current goals, reflecting instrumental conditioning of the orienting response. Our findings challenge the idea that selection history reflects a common mechanism of learning-dependent priority and instead suggest multiple distinct routes by which learning history shapes orienting behavior. We also provide direct evidence for the idea that value-based attention is approach oriented, which limits the effectiveness of attentional bias modification techniques that utilize incentive structures.

RESULTS

Healthy human participants ($n = 30$) performed an experiment comprising a training phase and a test phase. The training phase involved an antisaccade task in which participants moved their eyes in the direction opposite that of a colored square (Figure 1). Correct responses resulted in a monetary reward when the square was one of two colors (high-value target), creating a difference in associated value while matching other aspects of selection history [19, 20] (see STAR Methods for details). Partic-

ipants were faster (main effect of value: $F = 10.29$, $p = 0.004$, $\eta^2_p = 0.284$) and more accurate (main effect of value: $F = 7.18$, $p = 0.013$, $\eta^2_p = 0.216$) in generating an antisaccade to the high-value target as the task progressed, demonstrating motivated behavior consistent with the reward structure (Figure 2).

Immediately following the training phase, participants completed the test phase, which involved making a prosaccade to a shape-defined (circle) target. A colored square was presented opposite the target, which served as a task-irrelevant distractor (Figure 1). We compared response time and errors (eye movements directed toward the distractor) based on the reward history of the target and distractor colors, focusing our comparisons on the colors that were matched for prior target history but differed in associated value. When the target was rendered in the previously high-value color, participants were faster ($t = 2.67$, $p = 0.013$, $d = 0.51$) and marginally more accurate ($t = 1.84$, $p = 0.077$, $d = 0.35$) compared to when the target was rendered in the previously unrewarded color (Figures 3 and S1). Likewise, when the distractor was rendered in the previously high-value color, participants were slower ($t = 2.60$, $p = 0.015$, $d = 0.50$) and more error-prone ($t = 2.51$, $p = 0.019$, $d = 0.48$) compared to when the distractor was rendered in the previously unrewarded color. Even though participants were more highly rewarded and thus more motivated to look opposite a stimulus rendered in the high-value color during training, such training produced the exact opposite bias when probed in the test phase. That is, participants found it more difficult to move their eyes opposite the stimulus they were previously paid to look away from compared to an unrewarded color matched for selection history and were in fact facilitated in looking toward this color when directed to do so.

Selection history encompasses not only learning from the outcomes of past orienting responses but also the repetition of these orienting responses in the presence of a stimulus that signals the need to perform the action. In a second experiment, we manipulated the frequency of orienting away from different colors while removing any explicit rewards. A different group of participants ($n = 30$) completed four consecutive days of training in which they repeatedly performed an antisaccade in response to a frequent and infrequent color square. No reward feedback was provided (Figure 1). Participants improved in the speed and accuracy with which they executed the trained response across days (main effect of day: $F_s > 21.21$, $p_s < 0.001$, $\eta^2_p > 0.42$), similarly for the frequent and infrequent colors (main effect of frequency and interaction: $F_s < 0.83$, $p_s > 0.48$) (Figure 2). Then, on the fifth day of the experiment, participants completed



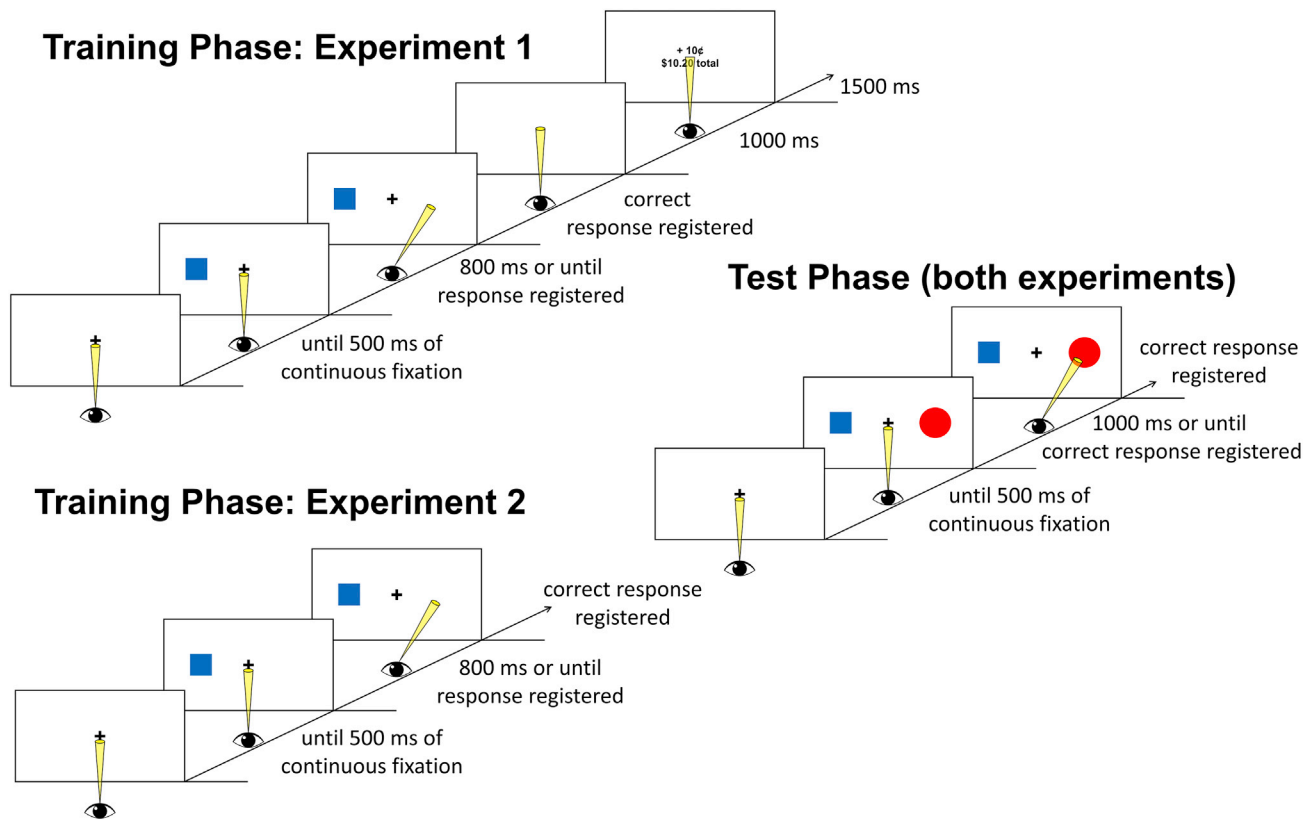


Figure 1. Time Course and Trial Events for Each Phase of the Experiments

During training, participants performed antisaccades in response to color targets. In experiment 1, one color predicted reward for a correct response, while another target color was never rewarded. In experiment 2, no explicit rewards were ever delivered, and one target color was more frequent than the other. The training phase comprised of a single session in experiment 1 and four sessions in experiment 2 (see STAR Methods for details). The test phase was the same for both experiments and consisted of a task in which participants made prosaccades to a circle target. The colors experienced during training were sometimes used for targets and distractors.

the same test phase as in the prior experiment. We focused our comparisons on the frequent former target color and the color that never served as a target previously, reflecting the presence versus absence of selection history (compare to the presence versus absence of reward history in experiment 1). Responses were numerically slower ($t = 1.10$, $p = 0.279$, $d = 0.20$) and significantly more error-prone ($t = 2.37$, $p = 0.025$, $d = 0.43$) when the target was rendered in the frequent color from training compared to when the target was a color that never served as a target previously (Figures 3 and S1). Likewise, responses were faster ($t = 2.12$, $p = 0.043$, $d = 0.39$) and less error-prone ($t = 2.32$, $p = 0.028$, $d = 0.42$) when the distractor was rendered in the frequent color from training compared to when it was rendered in a color that never served as a target previously. That is, participants found it more difficult to orient toward the stimulus they had frequently oriented away from during training and were facilitated in selecting a target when such selection also involved orienting away from this trained color. The pattern of results is directly opposite that of the first experiment, resulting in highly robust interactions between experiment and training history for all measures ($F_s > 7.67$, $p_s < 0.009$, $\eta_p^2 > 0.12$).


The infrequent target color produced an intermediate bias on all measures that did not differ significantly from either of the

other two colors ($p_s > 0.16$), which is consistent with the lack of a significant effect of frequency in the training data and suggests that both frequencies of repeated orienting across training days led to some learning. The interactions between experiments are maintained when comparing frequent and infrequent former targets in the current experiment ($F_s > 5.56$, $p_s < 0.023$, $\eta_p^2 > 0.09$), affirming the uniqueness of learned value in facilitating an approach-oriented bias. Consistent with a rapid influence of past orienting responses, there was some evidence for a bias to look away from the previously unrewarded color compared to the neutral color in experiment 1, particularly for the response-time measure (Figure S1), mirroring the pattern observed in experiment 2 for the frequent target color and suggesting that the influence of unrewarded selection history could emerge over a similarly brief timescale (see Figure S1 for a full breakdown of the data across all conditions).

DISCUSSION

What we pay attention to is biased by past experience or “selection history” [6–18]. Initially, selection history effects on attention were thought to reflect varieties of top-down guidance [21]. More recently, the influence of selection history on attention has been

Experiment 1



Experiment 2

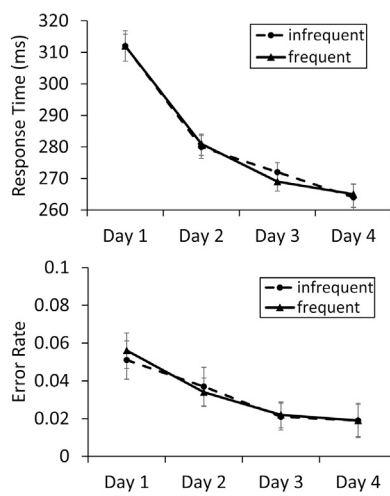


Figure 2. Response Time and Error Rate by Target Condition during the Training Phase of Experiments 1 and 2


Error bars reflect the within-subjects SEM. See also [Data S1](#).

described as reflecting a dedicated mechanism of control by which priorities are updated based on the outcomes of prior selection [12–18], constituting a third distinct factor in the control of attention [12–14]. Consistent with this idea, it has been proposed that correct task performance generates an internal reward signal that serves as a teaching signal to the perceptual system [22–25], mimicking the effects of learning from extrinsic rewards, which could explain why training with explicit rewards is not

When a stimulus predicts reward, the learning of this predictive relationship results in the stimulus gaining competitive priority in the visual system, biasing selection in its favor. This is in spite of the fact that participants were only ever rewarded for orienting away from the stimulus, providing powerful evidence for the role of associative learning in the orienting of attention [26–29]. On the other hand, repeatedly performing an orienting response in the presence of a particular stimulus facilitates the future

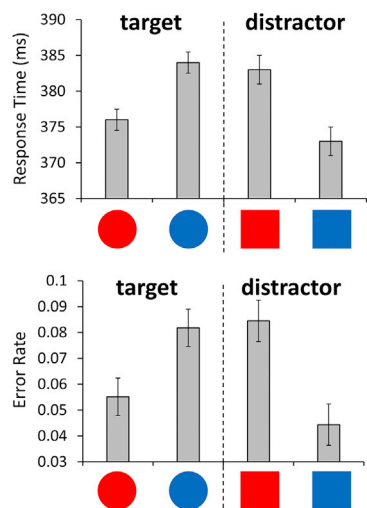
necessary to observe attentional capture by former target features [15].

Our findings provide clear evidence for two dissociable mechanisms underlying experience-driven attention and argue against a common mechanisms framework. Using an antisaccade training task, in one experiment we varied whether a stimulus predicted reward while matching the need to orient away from the stimulus, while in a different experiment we varied the need to orient away from a stimulus while removing any reward feedback.



Experiment 1

Legend for Experiment 1:
■ previously rewarded
■ previously unrewarded



Experiment 2

Legend for Experiment 2:
■ prior target (frequent)
■ no selection history

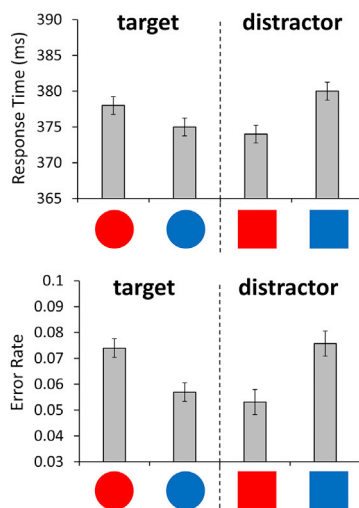


Figure 3. Response Time and Error Rate as a Function of the Selection History of the Target and Distractor Colors during the Test Phase of Experiments 1 and 2

In each experiment, three colors were used (previously high-value/frequent, previously unrewarded/infrequent, and never presented during training). Comparisons focus on the colors with and without reward history, matched for history as a former target (experiment 1), and the colors with and without history as a former target (experiment 2). Error bars reflect the within-subjects SEM. See also [Figure S1](#) and [Data S1](#).

engagement of this orienting response to the point of habit formation, reflecting a role for instrumental conditioning in the control of attention.

In a survey of the visual search literature, prior selection and reward were identified as two important factors governing the control of attention [30]. Our findings are consistent with this distinction and further indicate that their influence is fundamentally dissociable. Our findings are also consistent with the idea that the effects of statistical learning on attention, which is related to history as a former target, are distinctly habit based [31], which can be contrasted with the mechanisms of Pavlovian reward learning hypothesized to underlie value-driven attention [26–29]. However, a direct comparison of these hypothesized mechanisms has been lacking. By pitting these two sources of learning against each other, our study provides evidence for a distinctly reward-dependent mechanism of learning and a mechanism that is reward independent, with different components of the same experience independently shaping attentional priority.

When individuals are rewarded for orienting to a stimulus [6, 7, 19, 20], these two learning mechanisms will have the same consequence on future selection, with the resulting bias reflecting some combination of the two. This is the case, for example, in the context of addiction-related attentional biases resulting from drug use [31–34]. Distinguishing between different underlying components of experience-driven attention, and isolating their affects, may lend new insights into the relationship between attentional bias and clinical outcomes [34–37]. For example, the associative learning component we identify here reflects a form of visual attraction that appears to be approach oriented in nature, even when approach conflicts with action-reward associations, which may reflect a fundamental aspect of incentive salience and its role in the addiction process [38, 39].

Our findings also lend insight into why attentional bias modification techniques that provide incentives not to look at particular stimuli are limited in their effectiveness [36, 37, 40]. Although we paid people to explicitly look away from a particular stimulus, this training led to after-effects opposite what was incentivized. More effective attentional-bias-modification techniques should either not use incentives, as in our second experiment, or incentivize looking toward a different stimulus that might serve as a competitor for the target of reduced orienting [41].

The specific nature of the observed dissociation in learning mechanisms remains to be clarified. One possibility is that associative reward learning biases covert attention to reward-predictive cues, whereas unrewarded selection history specifically influences the oculomotor control system, producing habitual orienting responses. In the case of reward learning, this covert attentional bias would have competed with the practiced eye movement away from the stimulus in the present study, resulting in an opposing influence on oculomotor behavior. A related possibility is that learning from repetition involves the updating of specific stimulus-response mappings, whereas associative reward learning biases the representation of predictive stimulus features independently of task set. Our findings highlight two distinct learning influences, which should be further explored.

Although the conceptual umbrella of selection history is useful for drawing distinctions between mechanisms of control that are not learning dependent, the use of this term has also promoted an assumption that now appears untenable. The role of selection

history in the control of attention reflects at least two dissociable mechanisms and possibly more. For example, learning from punishment is also known to influence the control of attention [9, 18, 42], which may also differ mechanistically from the influence of reward. Aside from a shared dependence on learning, different components of selection history may be just as dissociable from each other as they are from the two other categories of attentional control (i.e., goal-directed and stimulus-driven). Future research needs to focus on differences within the domain of selection history and reconsider the utility of grouping fundamentally different learning experiences into a common category of attentional control.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHODS DETAILS
 - Apparatus
 - Test Phase (both Experiments)
 - Measurement of Eye Position
- QUANTIFICATION AND STATISTICAL ANALYSES
- DATA AND SOFTWARE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information includes one figure and one data file and can be found with this article online at <https://doi.org/10.1016/j.cub.2019.01.030>.

ACKNOWLEDGMENTS

Special thanks to Lindsey George and Victoria Fontenot for assistance with data collection. The reported research was supported in part by a start-up package to B.A.A. from Texas A&M University and NARSAD Young Investigator Grant 26008 to B.A.A. and NIH grant R01-DA046410 to B.A.A.

AUTHOR CONTRIBUTIONS

H.K. and B.A.A. jointly contributed to all aspects of the study, including conception, design, programming, data coding, data analysis, and manuscript writing.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: October 31, 2018

Revised: December 17, 2018

Accepted: January 14, 2019

Published: February 14, 2019

REFERENCES

1. Wolfe, J.M., Cave, K.R., and Franzel, S.L. (1989). Guided search: an alternative to the feature integration model for visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 15, 419–433.
2. Folk, C.L., Remington, R.W., and Johnston, J.C. (1992). Involuntary covert orienting is contingent on attentional control settings. *J. Exp. Psychol. Hum. Percept. Perform.* 18, 1030–1044.

3. Theeuwes, J. (1992). Perceptual selectivity for color and form. *Percept. Psychophys.* *51*, 599–606.
4. Theeuwes, J. (2010). Top-down and bottom-up control of visual selection. *Acta Psychol. (Amst.)* *135*, 77–99.
5. Itti, L., and Koch, C. (2001). Computational modelling of visual attention. *Nat. Rev. Neurosci.* *2*, 194–203.
6. Anderson, B.A., Laurent, P.A., and Yantis, S. (2011). Value-driven attentional capture. *Proc. Natl. Acad. Sci. USA* *108*, 10367–10371.
7. Della Libera, C., and Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychol. Sci.* *20*, 778–784.
8. Hickey, C., Chelazzi, L., and Theeuwes, J. (2010). Reward changes salience in human vision via the anterior cingulate. *J. Neurosci.* *30*, 11096–11103.
9. Schmidt, L.J., Belopolsky, A.V., and Theeuwes, J. (2015). Attentional capture by signals of threat. *Cogn. Emotion* *29*, 687–694.
10. Shiffrin, R.M., and Schneider, W. (1977). Controlled and automatic human information processing II: Perceptual learning, automatic attending, and general theory. *Psychol. Rev.* *84*, 127–190.
11. Chun, M.M., and Jiang, Y. (1998). Contextual cueing: implicit learning and memory of visual context guides spatial attention. *Cognit. Psychol.* *36*, 28–71.
12. Awh, E., Belopolsky, A.V., and Theeuwes, J. (2012). Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends Cogn. Sci.* *16*, 437–443.
13. Theeuwes, J. (2018). Visual selection: usually fast and automatic; seldom slow and volitional. *J. Cogn.* *1*, 21.
14. Failing, M., and Theeuwes, J. (2018). Selection history: How reward modulates selectivity of visual attention. *Psychon. Bull. Rev.* *25*, 514–538.
15. Sha, L.Z., and Jiang, Y.V. (2016). Components of reward-driven attentional capture. *Atten. Percept. Psychophys.* *78*, 403–414.
16. Peck, C.J., and Salzman, C.D. (2014). Amygdala neural activity reflects spatial attention towards stimuli promising reward or threatening punishment. *eLife* *3*, e04478.
17. Wentura, D., Müller, P., and Rothermund, K. (2014). Attentional capture by evaluative stimuli: gain- and loss-connoting colors boost the additional-singleton effect. *Psychon. Bull. Rev.* *21*, 701–707.
18. Stankevich, B.A., and Geng, J.J. (2014). Reward associations and spatial probabilities produce additive effects on attentional selection. *Atten. Percept. Psychophys.* *76*, 2315–2325.
19. Failing, M.F., and Theeuwes, J. (2014). Exogenous visual orienting by reward. *J. Vis.* *14*, 6.
20. Anderson, B.A., and Halpern, M. (2017). On the value-dependence of value-driven attentional capture. *Atten. Percept. Psychophys.* *79*, 1001–1011.
21. Wolfe, J.M., Butcher, S.J., Lee, C., and Hyle, M. (2003). Changing your mind: on the contributions of top-down and bottom-up guidance in visual search for feature singletons. *J. Exp. Psychol. Hum. Percept. Perform.* *29*, 483–502.
22. Roelfsema, P.R., and van Ooyen, A. (2005). Attention-gated reinforcement learning of internal representations for classification. *Neural Comput.* *17*, 2176–2214.
23. Roelfsema, P.R., van Ooyen, A., and Watanabe, T. (2010). Perceptual learning rules based on reinforcers and attention. *Trends Cogn. Sci.* *14*, 64–71.
24. Sasaki, Y., Nanez, J.E., and Watanabe, T. (2010). Advances in visual perceptual learning and plasticity. *Nat. Rev. Neurosci.* *11*, 53–60.
25. Seitz, A., and Watanabe, T. (2005). A unified model for perceptual learning. *Trends Cogn. Sci.* *9*, 329–334.
26. Le Pelley, M.E., Pearson, D., Griffiths, O., and Beesley, T. (2015). When goals conflict with values: counterproductive attentional and oculomotor capture by reward-related stimuli. *J. Exp. Psychol. Gen.* *144*, 158–171.
27. Bucker, B., and Theeuwes, J. (2017). Pavlovian reward learning underlies value driven attentional capture. *Atten. Percept. Psychophys.* *79*, 415–428.
28. Mine, C., and Saiki, J. (2015). Task-irrelevant stimulus-reward association induces value-driven attentional capture. *Atten. Percept. Psychophys.* *77*, 1896–1907.
29. Mine, C., and Saiki, J. (2018). Pavlovian reward learning elicits attentional capture by reward-associated stimuli. *Atten. Percept. Psychophys.* *80*, 1083–1095.
30. Wolfe, J.M., and Horowitz, T.S. (2017). Five factors that guide attention in visual search. *Nat. Hum. Behav.* *1*, 0058.
31. Jiang, Y.V. (2018). Habitual versus goal-driven attention. *Cortex* *102*, 107–120.
32. Field, M., Munafò, M.R., and Franken, I.H.A. (2009). A meta-analytic investigation of the relationship between attentional bias and subjective craving in substance abuse. *Psychol. Bull.* *135*, 589–607.
33. Marissen, M.A.E., Franken, I.H.A., Waters, A.J., Blanken, P., van den Brink, W., and Hendriks, V.M. (2006). Attentional bias predicts heroin relapse following treatment. *Addiction* *101*, 1306–1312.
34. Field, M., and Cox, W.M. (2008). Attentional bias in addictive behaviors: a review of its development, causes, and consequences. *Drug Alcohol Depend.* *97*, 1–20.
35. Anderson, B.A. (2016). What is abnormal about addiction-related attentional biases? *Drug Alcohol Depend.* *167*, 8–14.
36. Field, M., Marhe, R., and Franken, I.H.A. (2014). The clinical relevance of attentional bias in substance use disorders. *CNS Spectr.* *19*, 225–230.
37. Christiansen, P., Schoenmakers, T.M., and Field, M. (2015). Less than meets the eye: reappraising the clinical relevance of attentional bias in addiction. *Addict. Behav.* *44*, 43–50.
38. Robinson, T.E., and Berridge, K.C. (1993). The neural basis of drug craving: an incentive-sensitization theory of addiction. *Brain Res. Brain Res. Rev.* *18*, 247–291.
39. Berridge, K.C., and Robinson, T.E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Res. Brain Res. Rev.* *28*, 309–369.
40. Field, M., Duka, T., Tyler, E., and Schoenmakers, T. (2009). Attentional bias modification in tobacco smokers. *Nicotine Tob. Res.* *11*, 812–822.
41. Sigurjónsdóttir, Ó., Björnsson, A.S., Ludvigsdóttir, S.J., and Kristjánsson, Á. (2015). Money talks in attention bias modification: Reward in a dot-probe task affects attentional biases. *Vis. Cogn.* *23*, 118–132.
42. Nissens, T., Failing, M., and Theeuwes, J. (2017). People look at the object they fear: oculomotor capture by stimuli that signal threat. *Cogn. Emotion* *31*, 1707–1714.
43. Brainard, D.H. (1997). The psychophysics toolbox. *Spat. Vis.* *10*, 433–436.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited Data		
Raw data	This paper	Data S1
Software and Algorithms		
MATLAB	Mathworks	version 2017b
Psychophysics Toolbox for MATLAB	http://psychtoolbox.org/	version 3.0.13
EyeLink 1000Plus software	SR Research	version 5.12
Custom MATLAB code (for generating experiment stimuli, online recording of eye position, and analysis of eye data)	Haena Kim (hannah.kim@tamu.edu)	N/A
Other		
EyeLink eye tracker	SR Research	EyeLink 1000Plus, desktop mounted model

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Haena Kim (hannah.kim@tamu.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Thirty human participants were recruited from the Texas A&M University community for Experiment 1 (14/15/1 male/female/no response, ages 18-26y [mean = 21.6y]), and thirty new participants were recruited for Experiment 2 (15/15 male/female, ages 18-33y [mean = 23.6y]). All reported normal or corrected-to-normal visual acuity and normal color vision. All procedures were approved by the Texas A&M University Institutional Review Board and conformed with the principles outlined in the Declaration of Helsinki. Written informed consent was obtained from each participant. Participants were compensated with their task earnings in Experiment 1 and at a rate of \$10/hr in Experiment 2.

METHODS DETAILS

Apparatus

A Dell OptiPlex equipped with MATLAB software and Psychophysics Toolbox extensions [43] was used to present the stimuli on a Dell P2717H monitor. The participants viewed the monitor from a distance of approximately 70 cm in a dimly lit room. Eye position was monitored using an Eye Link 1000-plus desktop mount eye tracker (SR Research). Head position was maintained using an adjustable chin rest (SR Research).

Training Phase of Experiment 1

Each trial consisted of a fixation display, a stimulus array, and a reward feedback display (see [Figure 1](#)). The fixation display remained on screen until eye position was registered within 1.1° of the center of the fixation cross for a continuous period of 500 ms, which triggered the beginning of a trial, and then for an additional 400-600 ms. The stimulus array, which consisted of a $4.7^\circ \times 3.4^\circ$ color square appearing 12.2° center-to-center from fixation, was then presented for 800 ms or until an eye movement exceeding 8.2° in amplitude to the left or right was registered. The reward feedback display was presented for 1500 ms, and consisted of the money earned on the current trial along with the updated total earnings (if an eye movement exceeding 8.2° in amplitude in the direction of the color square was registered, the word “Incorrect” was presented in place of the money earned; if no eye movement exceeding 8.2° in amplitude in either direction was detected, the words “Too slow” were presented). A 1000 ms blank screen was inserted between the stimulus and feedback displays, and each trial concluded with a 1000 ms blank interval.

Participants were instructed to move their eyes opposite the color square on each trial as quickly as possible, and were informed that correct responses would sometimes be rewarded with a small amount of money. Each of two colors appeared equally-often across trials within a block, with each color appearing equally-often on each side of the display. Trials were presented in a random order. Correctly generating an antisaccade to one color target was associated with a reward of 15¢ (high-value color), while for the other target color a correct response always yielded 0¢ (unrewarded color). Red, green, and blue were used for the color of the squares, with each color serving as the high-value, unrewarded, and left-out color equally-often across participants. Participants

were not explicitly informed of the reward contingencies, which had to be learned from experience in the task. Participants completed 4 runs of 60 trials.

Training Phase of Experiment 2

The training phase was identical to that used in Experiment 1, with the exception that no monetary rewards were delivered and one color target appeared more frequently than the other. The frequent target color appeared on 90% of trials within a run, with the infrequent target color appearing on the remaining 10% of trials. Correct responses resulted in no feedback, with an errant saccade resulting in “Incorrect” feedback and a timeout resulting in “Too slow” feedback. Participants completed 7 runs of 80 trials on the first day of training, and 8 runs on each of days 2–4. Day 5 began with 2 runs of training before transitioning into the test phase. All 5 experiment sessions were completed on consecutive days.

Test Phase (both Experiments)

In the test phase, the stimulus array now consisted of a square and a circle, presented equidistant from fixation on the left and right. The size and spacing of these shapes matched that used in the training phase. Participants were instructed to fixate the circle on each trial, regardless of the color of the shapes. Red, green, and blue were used for the shapes, with no color ever appearing twice in the same display. Each color served as the target and distractor equally-often in a run, with the color by position pairings fully crossed and counterbalanced. Trials were presented in a random order. A correct response was registered if eye position was measured at more than 8.2° from fixation in the direction of the target within the 1000 ms timeout limit. If the participant generated a saccade landing more than 8.2° from fixation in the direction of the distractor, the trial was scored as containing an errant eye movement. No rewards were provided, and feedback (the word, “Miss”) was provided for 1000 ms only if participants failed to generate a saccade toward the target within the timeout limit. Each trial ended with a 1000 ms blank interval. Participants completed 3 runs of 96 trials.

Measurement of Eye Position

Head position was maintained throughout the experiment using an adjustable chin rest that included a bar upon which to rest the forehead (SR Research). Participants were provided a short break between each run of the task in which they were allowed to reposition their head to maintain comfort. Eye position was calibrated prior to each block of trials using 5-point calibration, and was manually drift corrected by the experimenter as necessary (the next trial could not begin until eye position was registered within 1.1° of the center of the fixation cross for 500 ms; see, e.g., [42]). During the presentation of the stimulus array, the X and Y position of the eyes was continuously monitored in real time, such that eye position was coded on line [26]. In the training phase, to ensure that only antisaccades were rewarded, an eye movement toward the target resulted in the termination of the trial and no reward. In the test phase, to maximize our sensitivity to measure selection bias, errant eye movements were recorded but did not result in the termination of the trial, such that participants could still fixate the target within the timeout limit and register a correct response even if gaze was initially directed toward the distractor.

QUANTIFICATION AND STATISTICAL ANALYSES

The coding of response time and errant saccades was performed on line during each trial of the experiment as described above. Response times exceeding 3 SD of the mean of a given condition were trimmed. For Experiment 1, data from three participants were not analyzed due to an inability to remain alert for the duration of the task, resulting in poor tracking quality. For Experiment 2, 11 participants withdrew from the study before completing all five sessions and were replaced; all participants who completed the entire experiment produced usable data. The training phase data were subjected to 4 × 2 analyses of variance (ANOVAs) with epoch (block in Experiment 1 and day of training in Experiment 2) and target condition as factors, separately for response time and error rate. For the test phase, paired samples t tests were performed separately for saccades with respect to target color and distractor color (to preserve independence of conditions), separately for response time and error rate. For each condition (bar depicted in Figure 3, the conditional mean collapses across the color of the non-referenced stimulus (e.g., high-value target trials collapse across trials on which the distractor was the unrewarded color and the color not used during training). For assessment of interactions between experiments, experiment was added as a between-subjects factor for the aforementioned pairwise comparisons and subjected to 2 × 2 ANOVAs.

DATA AND SOFTWARE AVAILABILITY

Raw data for Experiments 1 and 2 are provided as Supplemental Material ([Data S1](#)).