Learned value and object perception: Accelerated perception or biased decisions?

Jason Rajsic, Harendri Perera, Jay Pratt
Department of Psychology, University of Toronto

The final publication is available at Springer via http://rdcu.be/m43v

Please address correspondence to:
Jason Rajsic
Department of Psychology, University of Toronto
100 St. George St.
Toronto, ON
M5S3G3
416-978-6587
jason.rajsic@mail.utoronto.ca

This work was supported by the Natural Science and Engineering Resource Council of Canada under Grant 194537

Abstract

Learned value is known to bias visual search towards valued stimuli. However, some uncertainty exists regarding the stage of visual processing that is modulated by learned value. Here, we directly tested the effect of learned value on pre-attentive processing using temporal order judgments. Across four experiments, we imbued some stimuli with high value and some with low value using a non-monetary reward task. In Experiment 1, we replicated the value-driven distraction effect, validating our non-monetary reward task. Experiment 2 showed that high-value stimuli, but not low-value stimuli, exhibit a prior entry effect. Experiment 3, which reversed the temporal order judgment task (i.e., reporting which stimulus came second) showed no priority entry effect, indicating that while a response bias may be present for high-value stimuli, they are still reported as appearing earlier. However, Experiment 4, using a simultaneity judgment task, showed no shift in temporal perception. Overall, our results support the conclusion that learned value biases perceptual decisions about valued stimuli without speeding pre-attentive stimulus processing.

At any given moment, we can only attend to a small subset of the total amount of information in the visual environment. During each moment, there are a number of cognitive processes that collectively determine what information will be attended and what information will fall out of further processing. For the most part, different states of attention have been considered to be due to either bottom-up processes – driven by causes external to the individual – or top-down processes – driven by the goals of the observer. However, recent research has highlighted the contribution of sources of selection that are internal to the observer, yet not determined by his or her current goals (Awh, Belopolsky, & Theeuwes, 2012). The learned value of stimuli is one such source of attentional bias (e.g., Anderson, Laurent, & Yantis, 2011a). These value-driven attention biases can occur even when the value-laden features of stimuli are task-irrelevant (e,g., Anderson et al., 2011a; Raymond and O'Brien, 2009). Although reliably observed in laboratory experiments, the particular stage, or stages, of perceptual processing affected by learned value is not yet understood. In this paper, we assess the ability of learned value to affect visual priority in a task that does not require selective processing. First, however, we review what is known about the ways that learned value bias perceptual processing.

To study the effect of learned value on visual selection, studies have employed a two-phase structure, wherein different stimuli are repeatedly paired with different amounts of reward in a learning phase, and then attentional biases to these stimuli are compared in a test phase in which the reward contingency is removed (Raymond & O'Brien, 2009; Anderson, Laurent, & Yantis, 2011a; Anderson, Laurent, & Yantis, 2011b; Anderson, 2014; Miranda & Palmer, 2014; Sali, Anderson, & Yantis, 2014; MacLean & Giesbrecht, 2015). For example, Anderson, Laurent, and Yantis (2011b) trained participants to search for oriented bars within green or red circles amongst other colored distractor circles. For each participant, one target color had a high probability of producing a high reward, and the other target color had a high probability of producing a low reward. After practicing this task, reward contingencies were removed, and participants instead searched for an oriented bar within a unique, diamond shape among distractor circles (similar to the added singleton paradigm pioneered by Theeuwes, 1992). Critically, one of these circles on each trial would be colored in either red or green, and both of these singleton distractors led to slowed search times. Importantly, singletons in the color that had received high reward in the learning phase produced greater interference, indicating that the learned value of stimuli produces an attentional bias over and above that of perceptual salience.

As recently noted by Müller, Rothermund, and Wentura (2015), the majority of studies on reward and attention rely on search tasks to assess the prioritization of rewarded stimuli, and it is therefore unclear which stages of visual processing are affected by reward. These authors argued that reward effects in search may be due to delayed disengagement, as opposed to a preattentive boost for visual features with learned value. To support this argument, the authors reported data from a modified dot-probe task. After imbuing visual objects with value in a speeded-discrimination task, previously rewarded objects' ability to orient attention when acting as exogenous cues was compared neutral objects, as well as to objects associated with losses. While rewarded objects led to a larger cue validity effect, comparison with neutral cues showed that the rewarded objects led to slower disengagement (i.e., a larger difference between neutral and invalidly cued response times) but not to speeded orienting (i.e., no difference between neutral and validly cues response times). Müller et al. argued that delayed disengagement from rewarded stimuli could explain the attentional biases measured in search tasks, which are assessed by a slowed response time when an object with learned value appears as a distractor.

Using a different paradigm, Hickey, Chelazzi, and Theeuwes (2011) have argued instead that reward is able to affect early stages of target detection and localization, and that this target enhancement mechanism is distinct from a distractor suppression mechanism that operates on a later stage of selection. Although this finding is based on results of tasks where the effect of rewards on inter-trial priming, and not learned value, is measured, their conclusion is consistent with a recent electrophysiological and behavioral study showing that reward history influences the early stages of visual attention selection by altering P1 amplitude (MacLean & Giesbrecht, 2015, see Hickey, Chelazzi, & Theeuwes, 2010 for a similar result using immediate reward) and attentional capture, as indicated by the N2PC component (Qi, Zeng, Ding, & Li, 2013). Given that these studies involved associating learned values with stimuli, this result is inconsistent with Müller et al.'s conclusion that rewards solely affect delayed disengagement. Similarly, the suggestion that learned value solely delays disengagement is inconsistent with measures of oculomotor capture (Anderson & Yantis, 2012; Hickey & van Zoest, 2012; Theeuwes & Belopolsky, 2012). Instead, it points to an effect of learned value that is pre-attentive, in the sense that it does not require first focusing attention on a particular object to be measured. Behavioral evidence of preattentive locus of reward comes from Kiss, Driver, and Eimer (2009) who showed that pop-out was enhanced for targets that often deliver higher rewards (see also

Lee & Shomstein, 2014), however Kristjánsson, Sigurjónsdóttir, and Driver (2010) subsequently showed that this pop-out advantage rapidly reverses following a change in stimulus-reward contingencies, leaving uncertainty regarding whether learned value, as opposed to expected reward, operates at an early stage. What is missing is a direct, behavioral demonstration that stimuli with imbued learned value are prioritized for perception.

Our goal in this study was to directly test the claim that learned value can enhance preattentive processing of visual information. To do so, we employed judgments of stimulus onset (temporal order judgments [TOJs] and simultaneity judgments [SJs]), which are used to measure visual prior entry. Prior entry refers to the accelerated conscious perception of some stimuli at the expense of others, leading to earlier conscious perception of these stimuli (Sharlau, 2007; Spence & Parise, 2010). Prior entry is found to occur when attention is exogenously oriented to the location of an upcoming stimulus (Stelmach & Herdman, 1991; Hikosaka, Miyauchi, & Shimojo, 1993; Shore, Spence, & Klein, 2001; Schneider & Bavelier, 2003; Born, Kerzel, & Pratt, 2015). Although ERPs measured alongside TOJs do not always demonstrate accelerated processing (i.e., reduced peak latency of early components of the visual evoked potential), increases in the amplitude of early components (e.g., P1, N1, P2) are reliably observed, indicating that behavioral prior entry effects correspond to changes in early visual processing (McDonald, Teder-Sälejärvi, Di Russo, & Hillyard, 2005; Vibell, Klinge, Zampini, Spence, & Nobre, 2007). Importantly, these tasks can be used as a “cueless” tasks that measure the attentional biases that are intrinsic to stimuli, such as the speeded processing found for low spatial frequency patches (West, Anderson, Bedwell, & Pratt, 2010), emotional faces (West, Anderson, Pratt, 2009; West et al., 2010) and near surfaces (West, Pratt, & Peterson, 2013). Furthermore, they does not require selective processing – in fact, both stimuli must be registered to make a response – and so provides an index of visual priority when all information is equally relevant. Thus, TOJs and SJs provide an window into the perceptual biases that may exist for stimuli with learned value before focal attention is engaged, as it is difficult to envision a mechanism by which delayed disengagement alone could affect the relative perceived onset of stimuli.

In the present study we used a learned value paradigm modeled after Anderson, Laurent and Yantis’ (2011b) study, with one major exception: instead of monetary value, we assigned value using a point system. For the Experiment 1, our goal was to replicate the results of the

Anderson et al (2011b) study, especially given that our point rewards did not map onto any monetary value. To do this, we followed their modified value-learning task with an additional singleton visual search task to establish that value training was successful. We show that when the additional singleton feature was associated with learned value; it slowed down visual search proportional to the size of its associated value. In Experiment 2, participants completed the same value-learning task as Experiment 1, but were then tested using a novel TOJ paradigm to assess whether learned value would modify visual priority. Experiments 3 and 4 measured the perception of temporal onset for rewarded stimuli using a reversed TOJ and a SJ task to distinguish between three accounts of changes in perceptual judgments: true prior entry, response biases, and decision biases. To preview our results, we observed that although learned value biases temporal onset responses, such that highly valued stimuli are reported to be perceived earlier, they do not bias perception when simultaneity, and not order, is measured. This supports the proposal that learned value has effects on visual processing beyond delayed disengagement; specifically, in biasing perceptual decisions.

Experiment 1

As noted above, the main purpose of this experiment was to verify that rewarding participants with points rather than money would result in typical value-learning effects.

Participants

Twenty-two undergraduate psychology students naïve to the experiment were recruited from University of Toronto. Each participant reported normal or corrected-to-normal visual acuity and color vision. Participants gave written informed consent for the experiment and were provided with a course credit for participating in the experiment. All experimental procedures were approved by University of Toronto’s Office of Research Ethics and were in accordance with the Declaration of Helsinki.

Apparatus

The experiment was conducted using a Windows-run PC with a 19” CRT (1024 x 768 resolution with 85 Hz refresh rate) in a quiet and dimly lit room. Participants sat and viewed the monitor from a distance of 50cm with their chin rested on the chin-rest throughout the experiment. The experiment was run in MATLAB, using Psychophysics toolbox. Participants entered responses by using a standard keyboard**.**

Stimuli and Procedure

Participants were tested in a dimly lit room for a single 1-hour session. Prior to the experiment participants were presented with the instructions of the experiment using a PowerPoint presentation that included images of the visual stimuli used in the experiment alongside with the written instructions. Participants were told to place their chin on the chin-rest and to make fast and accurate responses on each trial of the experiment.

Each phase of the experiment began with a screen with instructions reiterating the instructions that had been orally provided to the participants. Stimuli for both phases were presented against a uniform grey background with a white fixation cross, 0.4° in size, centered on the screen.

Training phase

The training phase of Experiment 1 was used to imbue stimuli with learned value by repeatedly pairing them different rewards. Trials in the training phase were made up of displays composed of four Landolt-Cs, 1.5° in radius, drawn in four different colors, appearing at random positions, all centered 6.4° from fixation. Of these Landolt Cs, three, with their gaps (0.36° in size) on the top or bottom, were distractor stimuli, and one, with its gap on the right or left, was the target stimulus. The possible colors of each the distractor stimuli were orange (RGB: 192, 192, 0), blue (RGB: 0, 192, 192), yellow (RGB: 255, 128, 0) and, depending on the trial, the target stimulus could be either red (RGB: 255, 0, 0) or green (RGB: 0, 255, 0). The search display was presented until participants made their response. Participants had to identify if the gap on the colored circle was left or right by pressing the left-arrow or right-arrow keys, respectively. A feedback display followed the response to inform the participant of how many points he or she had earned for the completed trial, which was presented in the center of the screen in white Arial font varying in size depending on the reward magnitude. High-reward (200 points) were shown in large text (48 point, approximately 1.8° in height) while low-reward (20 points) were shown in smaller text (16 point, approximately 0.6° in height). The total amount of points were presented for 1 second, and added to a running tally that was continuously visible at the top of the screen.

Correct responses were followed by visual feedback indicating amount of points earned during the training phase. High-reward targets were followed by 200 points (high reward) feedback on 80% of the trials and low-reward feedback on 20% of the trials. Low-reward targets were followed by 20 points (low reward) feedback on 80% of the trials and high-reward feedback on 20% of the trials. High-reward target color and low-reward target color was randomly assigned as red or green for each participant.

The training phase consisted of a variable number of trials grouped into 12 blocks. Prior to completing the training phase, practice trials were provided. Practice trials were identical to actual trials except all visual stimuli were presented in white and points earned on each trial were equal to 0 or 10 points, for incorrect trials and correct trials, respectively. The practice phase ceased when participants had collected 100 points; in other words, once they had correctly completed 10 trials. Between each block and after completion of training phase, participants were provided with a short break. Each block was terminated after the participant had accumulated 2500 points.

Test phase

For our test phase, we used an additional singleton task (Theeuwes, 1992). During this task, eight stimuli appeared on a search display, where each search stimulus was placed, evenly spaced, the circumference of an imaginary circle, radius 6.4°, centered on fixation. Seven of these stimuli were Landolt Cs, 1.5° in radius, and the eighth stimulus was a Landolt square outline, 3.0° in width and height, Each Landolt had a 0.36° gap on either the left or right side (forward facing or reverse). The target was defined as the square outline with a 0.36° gap on either the left or right side. Depending on the trial type, either all stimuli were colored in white, or all stimuli were colored in white except for one (the additional singleton), which was either drawn in the high-reward associated color or the low-reward associated. There were no feedback or points provided following each trial.

The search display was presented until participant made their response. Participant had to identify which side the gap, left or right, is located on the square target by pressing the z or m key, respectively. Response time was measured from the onset of the visual stimuli to the response made by each participant.

The test phase of the experiment included 320 trials that were divided into 8 blocks. Once again, practice trials were provided before the test phase was completed. In total there were four conditions in which RTs were compared for the addition singleton: no color, distractor color, high-value color and low-value color. High-value and low-value colors refer to the same colors used for the high-reward and low-reward target for the training phase of the experiment. Target and additional singletons were equally likely to appear in each of the eight positions of the search array throughout the experiment. The additional singleton always appeared as a distractor. The search display stayed on the screen until participant made their response and then the next search display will be presented.

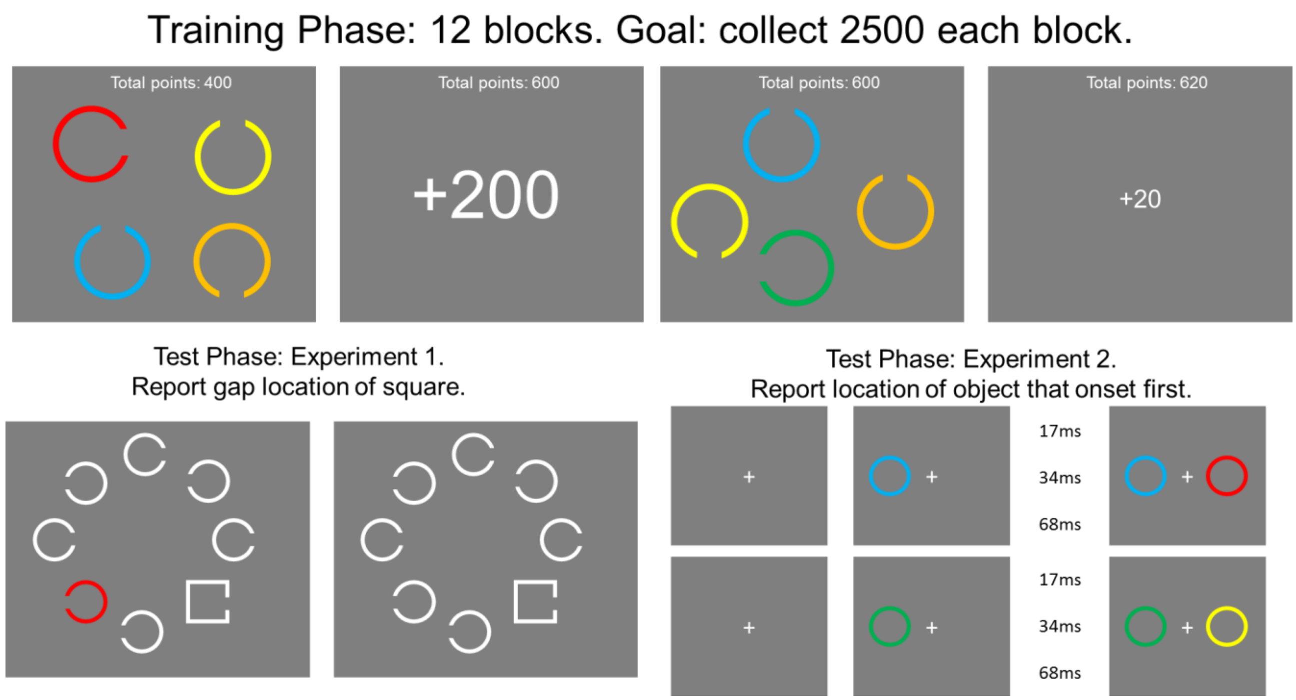

Figure 1. Upper panel: schematic of the training phase used in Experiments 1 and 2. Point-based rewards were delivered upon correct response input. Participants' task was to report the gap location of the red or green Landolt. Lower panel: Schematic of test phases for Experiment 1 and Experiment 2. Lower left panels depict a high-value singleton trial and a no-singleton trial in Experiment 1. Lower right panels depict a high-value TOJ trial (top) and a low-value TOJ trial (bottom). Stimuli are not drawn to scale.

Results and Discussion

Correct response times in the acquisition were analysed by dividing the training phase into first and last halves, each of which with high- and low-reward associated targets. Trials were trimmed within-participants by removing trials with RTs outside of 2 standard deviations of a participant's mean RT. A Block x Reward ANOVA revealed a main effect of Block, $F(1, 21)$ =

27.10, $p < .001$, $\eta_p^2 = 0.56$, but no main effect of Reward, $F(1, 21)$, = 0.84, $p = .37$, $\eta_p^2 = 0.04$, and no interaction, $F(1, 21) = 1.03$, $p = .32$, $\eta_p^2 = 0.05$, although response times were numerically faster for high-reward targets, *M* = 519, *SE* = 5ms, than low-reward targets, *M* = 531ms, *SE* = 4ms, in the last half of the training phase. Thus, we did not find reliable evidence of a difference in response time between high- and low-reward targets in our training phase.

In the test phase, correct response times and accuracy were *M* = 535 ms, *SE* = 15 ms, and *M* = 97.2%, *SE* = 0.6%, respectively. To determine whether learned value from the training phase affected the allocation of attention in the test phase, average correct response times for the additional singleton effects in the test phase were analysed using a one-way, repeated measures ANOVA with Singleton Condition (Low Value, High Value, and no Singleton) as a factor. Averaged correct response times in each condition are shown in Figure 2. A main effect of singleton type was present, $F(2, 42) = 8.80$, $p = .001$, $\eta_p^2 = 0.30$. Follow-up contrasts revealed that Low Value Singletons slowed search times relative to No Singleton trials, $F(1, 21) = 4.82$, $p = .04$, $\eta_p^2 = 0.19$, and, critically, that High Value Singletons slowed search times even further, relative to Low Value Singletons, $F(1, 21) = 6.15$, $p = .02$, $\eta_p^2 = 0.23$. No differences in accuracy were observed by Singleton Condition, $F(2, 42) = 0.99$, $p = .38$, $\eta_p^2 = 0.05$. This demonstrates that, in a task that used points in lieu of monetary reward, learned value led to stable changes in attentional priority, such that stimuli associated with more reward exhibited increased distraction in a subsequent task.

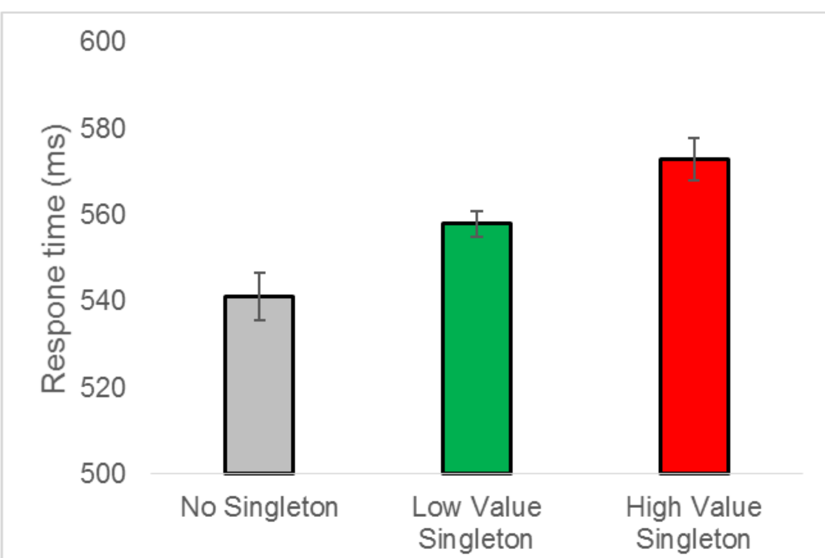

Figure 2. Correct response times in the test phase of Experiment 1. Error bars represent 1 within-subjects standard error.

Experiment 2

Given that we were able to show a learned value effect on the allocation of attention in our version of the task used by Anderson et al. (2011b), we substituted a temporal order judgment task in to the test phase to measure whether learned value affects the speed with which stimuli are processed. If learned value does increase pre-attentive visual priority, we expected that stimuli associated with higher value should be perceived earlier than stimuli with lower value.

Participants

Thirty-one undergraduate psychology students naïve to the experiment were recruited from the University of Toronto. All reported normal or corrected-to-normal visual acuity and color vision. Participants were provided with a course credit in return for the participation of the experiment. None of the participants who participated in Experiment 1 were participants in Experiment 2. All experimental procedures were approved by University of Toronto’s Office of Research Ethics and in accordance with the Declaration of Helsinki.

Apparatus

The apparatus used were identical to Experiment 1

Stimuli and Procedure

Similar to Experiment 1, instructions were presented orally using a PowerPoint presentation, which included written instructions along with all the visual stimuli included in the experiment. All procedures used were identical to Experiment 1 with the exception of the test phase. The test phase began with a screen reiterating the instructions presented prior to the experiment. Participants were asked to identify which of the two filled circles they thought appeared first by pressing the z key if the left circle appeared first, or the m key if the right circle appeared first. Similar to the training phase, they were asked to make fast and accurate responses and were given the opportunity to take a break between each block.

Similar to the test phase of Experiment 1, participants were provided with 10 practice trials where the task was identical to the actual experiment with the exception that the stimulus circles were white. In total, there were 384 trials that were divided into 8 blocks. Two circles resembling the Landolt C's from the training phase, but with no gap (with a radius of 1.5°) were presented 6.4° away from the fixation cross, on the horizontal meridian. The two circles were drawn in the colors used in the training phase. Two types of trials were used; the high-value color appearing with a distractor color, and the low-value appearing with a distractor color. The first circle appeared on the left or right side of the fixation cross and was followed by the second circle that appeared following a stimulus-onset asynchrony (SOA) of 64 ms, 32ms, or 16 ms. The two circles remained on the screen for 100 ms, after which they offset and a response was collected.

Results and Discussion

Three of the 31 participants were excluded from the TOJ analysis because their response accuracy was not significantly above chance across all reward colors and SOAs (in other words, they did not temporally discriminate the two stimuli). All analyses were performed on the remaining 28 participants. Average correct response time during the training phase were again analysed using a Block X Reward ANOVA. Unlike Experiment 1, the learning phase of Experiment 2 revealed a marginal main effect of Value, $F(1, 27) = 4.31$, $p = .05$, $\eta_p^2 = 0.14$, such that High Value targets were reported faster, *M* = 549ms, *SE* = 14ms, than Low Value targets, *M* = 560ms, *SE* = 15ms, as well as a main effect of Block, $F(1, 27) = 37.54$, $p < .001$, $\eta_p^2 = 0.58$.

For the TOJ task, trials were organized by two factors: the SOA between Valued and Non-Valued color stimuli (six levels: -64 ms, -32 ms, -17 ms, 17 ms, 32 ms, 64 ms), and which Valued Color was used (Low Value, High Value). The responses on these trials were used to fit two Psychometric functions (cumulative Gaussian distributions) for each participant, parametrizing the probability of choosing the color with learned value as having appeared first at each SOA separately for the two valued colors. Fitting was accomplished using a maximum likelihood approach, with Matlab's (by MathWorks) *fminsearch* function used to minimize the negative Log-Likelihood of the parameters of the Psychometric function. As a result, prior entry could be assessed for each Valued Color by comparing the point of subjective equality (PSE) defined by the Psychometric function (the point at which each stimulus is equally likely to be chosen, corresponding the μ, or mean, parameter of the function).

The PSS for the Low Value color, *M* = -3.78 ms, *SE* = 3 ms, was not significantly different from 0, $t(27) = 1.30$, $p = .21$, indicating no prior entry for the Low Value color, compared to a neutral color. Importantly, the PSS for the High Value color, *M* = -11ms, *SE* = 4 ms, was significantly different from 0, $t(24) = 2.88$, $p = .007$, indicating that High Value colors did receive prior entry (see Figure 3). A direct comparison of PSS values (Figure 3) yielded the same conclusion, $t(27) = 2.08$, $p = .047$, while no differences in the slope of temporal order judgments was evident, $t(27) = 1.07$, $p = .29$. These prior entry results suggest that learned value is able to affect pre-attentive visual priority.

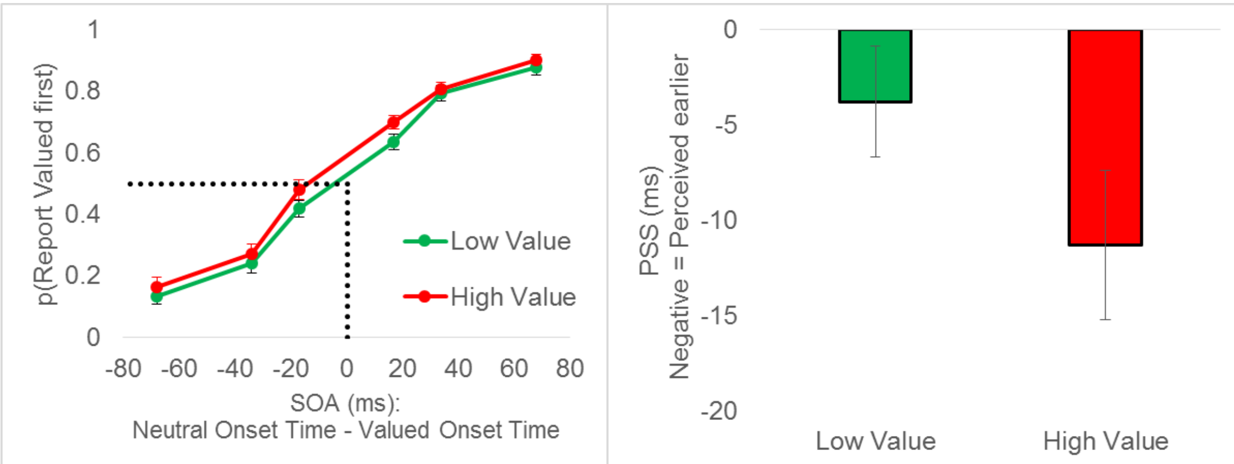

Figure 3. Results from Experiment 2's test phase. The left panel depicts across-participant average probabilities of reporting the valued stimulus as onsetting first for each stimulus onset asynchrony. The right panel depicts averaged PSS values derived from individual participant fits. Error bars reflect one standard deviation of the mean.

Experiment 3

Although Experiment 2 provided evidence that learned value leads to prior entry, the results are equally consistent with the possibility that learned value increases the choice salience of an object. A number of investigators have remarked that an increase in the probability of an object being chosen first in a temporal order judgment can be observed because of a true change in perceived temporal order, or simply a bias to choose a particular object for report (Shore, Spence, & Klein, 2001; Schneider & Bavelier, 2003). As such we ran a new group of participants through a task identical to the one we used in Experiment 2, save for the fact that participants were instructed to report the object that onset last. If the results of Experiment 2 were due to a bias towards choosing the rewarded stimulus, then we should observe a reversed effect on the PSS of stimuli with learned value. However, if the results of Experiment 2 were due to perceptual prior entry, then no such reversal should occur.

Participants

Thirty one adult volunteers were recruited for Experiment 3. Each participant was compensated with either course credit or $10 for participation. All participants provided informed consent, and no participants had participated in either Experiment 1 or 2. All experimental procedures were approved by University of Toronto's Office of Research Ethics and in accordance with the Declaration of Helsinki.

Apparatus, Stimuli, and Procedure

All apparatus, stimuli, and procedure were identical to those used in Experiment 2. Participants were simply instructed that, during the temporal order judgment task, they should report which of the two stimuli onset last.

Results and Discussion

Four participants were excluded, as in Experiment 2, on the basis of poor TOJ performance. Correct mean response times for the training phase were analysed in a Block X Value ANOVA. No main effect of Value was observed, $F(1, 26) = 1.54$, $p = .23$, $\eta_p^2 = 0.06$, but RT was affected by Block, $F(1, 26) = 14.82$, $p = .001$, $\eta_p^2 = 0.36$, such that RT was lower in the second half. A marginal Value X Block interaction was present, $F(1, 26) = 3.57$, $p = .07$, $\eta_p^2 = 0.12$, and so we analysed the effect of Value for the first and last halves of the training phase separately. In the first half, RT did not differ for High- and Low-reward targets, $t(26) = 0.23$, $p = .82$, but did differ for the second half, $t(26) = 2.49$, $p = .02$, suggesting that an RT benefit for High Value targets emerged later into the training phase.

PSS values for High- and Low-value stimuli were estimated again by fitting a cumulative Gaussian distribution, except that now the fitted distribution was inverted (i.e., $1 – \varphi$). Unlike Experiment 1, neither Low-value stimuli, $t(26) = 1.35$, $p = .19$, nor High-value stimuli, $t(26) = 0.03$, $p = .97$, showed a PSS shift from 0 (see Figure 4). Following the analysis of TOJ effects by Shore et al., (2001), this indicates that the responses in Experiment 2 were likely due to a mixture of prior entry and decision biases. In the present experiment, the which-came-second task pitted these two effect against each other, and they cancelled each other out. Thus, these results support the conclusion that learned-value affects both perceptual and response biases in temporal order judgments.

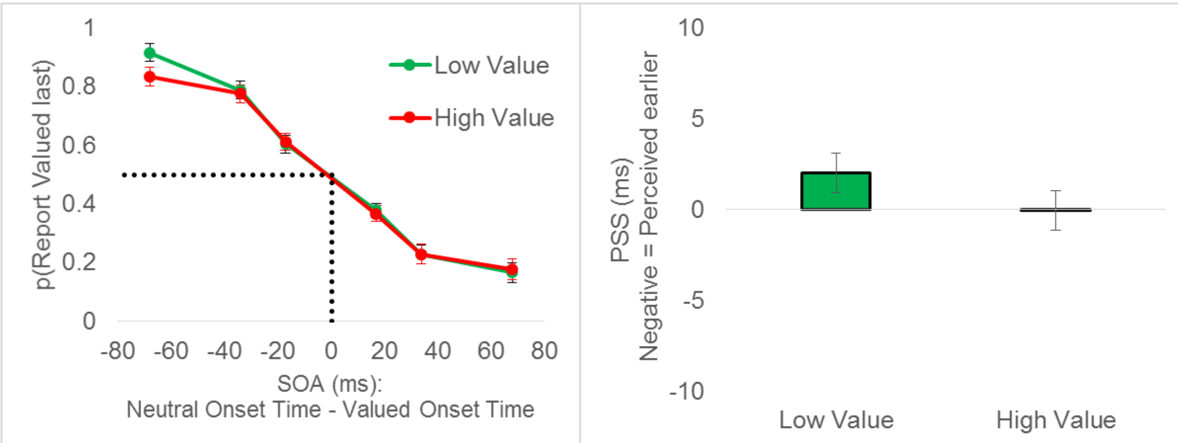

Figure 4. Results from Experiment 3's test phase. The left panel depicts across-participant average probabilities of reporting the valued stimulus as onsetting last for each stimulus onset asynchrony. The right panel depicts averaged PSS values derived from individual participant fits. Error bars reflect on standard deviation of the mean.

Experiment 4

Experiments 2 and 3 demonstrate that learned value can affect the perception of temporal order. However, whether this reflects true prior entry or not is still unclear. As first argued by Schneider and Bavelier (2003), TOJ tasks may be contaminated by a third type of bias – a decision bias. The TOJ task requires the detection two signals and comparing their onsets. Given the presence of sensory noise, some evidence threshold is necessary for the detection of onsets. A bias to report a valued stimulus may therefore reflect either an increase in the signal strength (i.e., a true change in the stimulus onset signal) or a change in its decision threshold. In Experiment 4, we measured the perception of onset for stimuli with learned value using a SJ task, where participants report whether two stimuli appear at the same time or different times. If stimuli with learned value indeed receive accelerated visual processing, we should observe a shifted PSS using this SJ task.

Participants

Thirty-one participants were again recruited to participate in Experiment 4. All participants were compensated for their participation with $10. All experimental procedures were approved by University of Toronto's Office of Research Ethics and in accordance with the Declaration of Helsinki.

Apparatus, Stimuli, and Procedure

Identical apparatus, stimuli, and procedure from the previous experiments were used in Experiment 4 with two exceptions. First, because simultaneity was to-be-reported in this task, we introduced trials in the test phase wherein stimuli onset simultaneously, randomly intermixed. To accommodate these extra trials, we increased the number of trials in the test phase from 384 to 448. In total, there were 28 trials of the six asynchronous onsets used in Experiments 2 and 3, and 56 trials with simultaneous onsets, per stimulus type (low value vs. neutral, high value vs. neutral). Second, instead of being instructed to report the stimulus that onset first, participants were instructed to report whether the stimuli appeared at the same or different times. The "z" key was used to indicate perceived simultaneous onset and the "/" key was used to indicate perceived asynchronous onset.

Results and Discussion

Three participants were excluded from analysis due to poor performance in the SJ task. Data from the remaining 28 participants was analysed for both the training and test phases. In the training phase, RTs were faster in the second half than the first, $F(1, 27) = 7.99$, $p = .009$, $\eta_p^2 = 0.23$. A main effect of Value was present, $F(1, 27) = 25.18$, $p < .001$, $\eta_p^2 = 0.48$, but Value and Block interacted, $F(1, 27) = 4.47$, $p = .044$, $\eta_p^2 = 0.14$. Paired-samples *t*-tests indicated that, as in Experiment 3, no difference in RT was present between High- and Low-value trials in the first half of the training phase, $t(27) = 1.01$, $p = .32$, but RT was faster for High- than Low-value trials in the second half, $t(27) = 3.54$, $p = .001$.

Simultaneity judgments were analysed by fitting responses to the difference between a cumulative Gaussian distribution and an inverse Gaussian distribution, as in Schneider and Bavelier (2003). Paired-samples *t*-tests showed no differences between the parameters fitted for High- and Low-value SJs, *t*s(27) < 0.44, *p*s > .66, and, critically, no difference between the PSS between High- and Low-value stimuli (see Figure 5), *t*(27) = 0.25, *p* = .80. These results challenge the conclusion that learned value leads to accelerated visual processing per se, and instead favor an account wherein perceptual decision thresholds are lowered for valued stimuli, leaving the speed of sensory processing unchanged. One potential explanation for why we found a PSS shift in our TOJ tasks, but not the SJ task, comes from van Eijk, Kohlrausch, Juolaand van de Par (2010), who showed that PSS estimates in TOJ tasks can reflect a shift towards the stimulus with greater temporal sensitivity. Fitting our data with asymmetric slopes (i.e., different mean and standard deviations for the Value-stimulus leading and Value-stimulus trailing components of the response distribution), however, did not yield slope differences, *t*(27)s < 1.57, *p*s > .13. As such, our data cannot speak to this possibility.

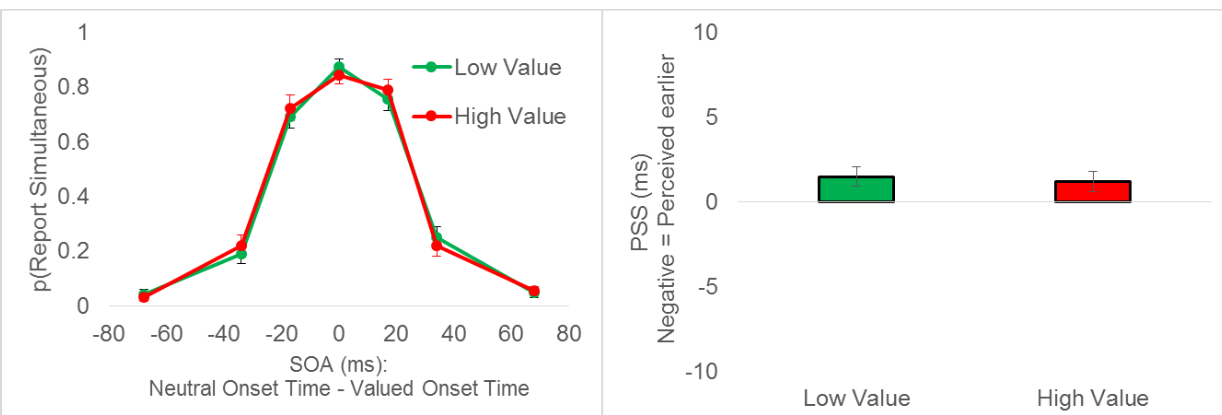

Figure 5. Responses to High- and Low-value stimuli in the SJ task. The left panel depicts aggregate mean simultaneity reports for each SOA. The right panel depicts average, estimated PSS values. Error bars depict 1 within-subjects standard error.

General Discussion

The present study sought to establish whether learned value can affect pre-attentive processing of previously rewarded visual information using a behavioral measurement. In Experiment 1, we replicated the findings of Anderson et al. (2011b), confirming that our training procedure was able to produce a value-driven attentional bias. In Experiment 2, we used an identical training phase to imbue stimuli with differential learned value. Using a temporal order judgment task, we observed that stimuli with greater learned value were perceived to onset earlier than stimuli with lower learned value, but equivalent exposure and task-relevance history. Experiment 3 showed that these effects where not entirely due to simple response biases. Critically, however, Experiment 4 showed no such difference in perceived simultaneity between High- and Low-value stimuli. Schneider and Bavelier (2003) argued that such a pattern of results – a shifted PSS for attended stimuli in TOJ tasks, but not in SJ tasks - indicates that no sensory acceleration occurs due to attention, but rather the decision criteria used to estimate onset time are affected. Indeed, research has shown that PSS estimates in TOJ and SJ tasks do not necessarily correlate (van Eijk, Kohlrausch, Juola, & van de Par, 2008). One explanation for why this occurs is that biased PSS values in TOJ tasks occur due to a bias to report the stimulus that has better temporal resolution (van Eijk, Kohlrausch, Juola, & van de Par, 2010). However, we did not observe differences in sensitivity when value-laden stimuli onset first, compared to last, in our SJ task. It is important to note that SJ and TOJ tasks may reflect decisions based on different sensory information; specifically, SJ judgments may often be based on the total duration of both stimuli, if stimulus durations are fixed (see Love, Petrini, Cheng, & Pollick, 2013). Therefore, the inference that no prior entry occurs for stimuli with learned value from our data requires the supposition that a lack of a PSS shift in SJ tasks accompanied by shifted PSS values in TOJ tasks should be interpreted as a post-perceptual decision bias, consistent with the dominant view in the prior entry literature (see Schneider & Bavelier, 2003; García-Pérez & Alcalá-Quintana, 2015). As such, we conclude that learned value acquired in our task did not produce prior entry.

Several studies measuring the effects of recently delivered rewards on selective attention show what may be considered to be early effects of reward on selection (Hickey, Chelazzi, & Theeuwes, 2010; Hickey, Chelazzi, & Theeuwes, 2011). In particular, priming of pop-out is enhanced after reward delivery (Kiss, Driver, & Eimer, 2009), and visual priming similarly leads to shifts in the PSS for as measured by both TOJ and SJ tasks (Theeuwes & Van der Burg, 2013). Given the lack of a clear PSS shift across tasks, despite consistent stimulus value learning, we suggest that the consequences of recent reward and learned value for visual processing may in fact differ. As such, future research should compare the effects of recently delivered reward and learned value with caution; although the distracting effect of stimuli associated with reward over the short- and long-term in search may be similar, the broader visual effects of reward and value may not be identical. While detailed descriptions of how moment-to-moment rewards may result in lasting attentional biases have been advanced (Rombouts, Bohte, Martinez-Trujillo, & Roelfsema, 2015; Failing & Theeuwes, 2016), our results suggest that this process is worth investigating in detail. As noted earlier, Kristjánsson et al. (2010) found that the effect of reward on priming of pop-out rapidly changes when reward contingencies change, leaving open the possibility that the removing reward contingencies (as is necessarily done in studies of learned value) may affect early sensory consequences of reward more than later consequences (i.e., response and decisions biases). This is, however, inconsistent with the ERP findings of Maclean and Giesbrecht (2015).

The issue of how stimulus-reward pairings do or do not accumulate in to lasting value-driven biases may benefit from an integration with the rich literature on the mechanisms of intertrial priming of attention (Becker, 2008; Olivers & Meeters, 2006; Kristjánsson & Campana, 2010; Kruijne & Meeter, 2016). Indeed, Sha & Jiang (2106) have recently argued that value-driven attention may rely on target history-related priming. While our results showed value-dependent differences in capture -- that is, high-value stimuli were found and reported as onsetting faster than low-value stimuli -- both stimuli had a history of task-relevance, raising the possibility that stimulus value is learned for attended stimuli only (but see Le Pelley, Pearson, Griffiths, & Beesley, 2015). If the developing literature on learned value and attention seeks to account for real-world attentional biases (e.g., Field & Cox, 2008), then characterizing the mechanisms underlying the learning of stimulus value will be of critical importance.

One potentially significant difference between our experiments and those experiments that show an early effect of learned value is that we used a money-less reward-learning task. The majority of value-driven attention studies rely on monetary incentive in order to create stimuli with value associations. Three exceptions are Shomstein and Johnson (2013), who showed a reversal of object-based attention when more points were awarded for the correct detection of targets in non-cued objects, regardless of whether the points led to monetary reward or simply were accumulated, Miranda and Palmer (2014), who showed that combining points and sound effects (as well as a “high-score” counter) could produce similar learned-value effects for stimuli paired with higher reward, and Roper and Vecera (2016), who paired correct responses to different target stimuli with the presentation of different denominations of currency, finding that those stimuli paired with the appearance of larger denominations led to greater attentional capture. While it appears that monetary reward is not necessary to entrain learned value when measured using search times, it is possible that not all rewards affect perception equally, or that the longevity of different rewards' effects on attention differ. Indeed, Miranda and Palmer did not find that points alone could create value-driven attention effects, whereas points alone were sufficient to affect attention in both Shomstein and Johnson's experiments and our experiments. While differences exist in each case between the specific tasks and point values used, we note that, in our task, higher points reduced the number of trials that participants needed to complete, as each block of trials simply required a criterion value of accumulated points in order to be completed. In our paradigm, then, learned value may have been predicated on the reduction in time or effort that accompanied higher point-values. As the old adage goes, time is money, and the subjective impact of a high-reward in our task translates to a reduction in the potential number of trials to be completed. While participants were clearly sensitive to this reward, it may not bias attention in quite the same way as the receipt of money. One reason for this could be that the delivery of money (even symbolic) would be considered positive reinforcement, whereas earning points that reduce the number of trials to be completed could arguably be considered negative reinforcement (the removal of impending effort). As such, these types of rewards may produce different effects on selection.

Another salient difference between Miranda and Palmer's experiments, which showed no effect of points alone on attention, and experiments where non-monetary reward led to value-driven attention (Shomstein & Johnson, 2013; Roper & Vecera, 2016; the present experiments) is the difference in feedback complexity. Experiments where non-monetary reward has led to value-driven attention have used consistent mappings between a particular feedback stimulus and high- or low-rewards. In our experiment, high rewards were always "200 points" and low rewards were always "20 points"; in Shomstein and Johnson's experiments, high rewards were always "6 points" and low rewards were always "1 point"; in Roper and Vecera's experiments', high rewards were always depictions of $20 and low rewards were always depictions of $5. Indeed, Miranda and Palmer's successful demonstrations of non-monetary, value-driven attention (Experiments 1 and 3) seem to have occurred when a particular sound (an "electric whip") accompanied positive feedback; points, when awarded, varied by participants' response time, meaning that the high reward values, while 5 times larger than the low reward values on average, could change across trials, perhaps making reward-stimulus associations more difficult to learn. This is not to say that consistent mapping is sufficient for reward learning, as Roper and Vecera found no value-driven attentional biases when monetary amounts did not appear as monetary images, but instead as simple numeric amounts (even when preceded by a dollar-sign). However, a consistent mapping between high-value and particular stimulus that conveys high value may be important for the rapid formation of value-driven attentional biases. If this is the case, it suggests that associations between value and attention may be underlain by associations between stimuli (i.e., between target stimuli and the stimuli that signal reward, but not between target stimuli and abstract reward). While the use of a different types of reward may be responsible for conflicting results, insofar as our lack of prior entry conflicts with ERP data, ultimately, we see this as an advantage for the literature on reward, value, and attention. One goal of research into the effects of reward and value on attention must be generalizable theories, and so testing different types of rewards (e.g., positive emotional expressions; Anderson, 2015) is essential to understanding the nuances of motivated attention.

References

Anderson, B. A. (2014). Value-driven attentional priority is context specific. *Psychonomic Bulletin & Review, 22*(3), 750-756.

Anderson, B. A. (2015). Social reward shapes attentional biases. *Cognitive Neuroscience*, doi: 10.1080/17588928.2015.1047823.

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences, 108*(25), 10367-10371.

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Learned value magnifies salience-based attentional capture. *PLoS One, 6*(11), 1-6.

Anderson, B. A.& Yantis, S. (2012). Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. *Attention, Perception & Psychophysics, 74*(8), 1644-1653.

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends in cognitive sciences, 16*(8), 437-443.

Becker, S. I. (2008). The mechanism of priming: Episodic retrieval or priming of pop-out? *Acta Psychologica*, *127*(2), 324-339.

Born, S., Kerzel, D., & Pratt, J. (2015). Contingent Capture Effects in Temporal Order Judgments. *Journal of Experimental Psychology: Human Perception and Performance, 41*(4), 995-1006.

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

García-Pérez, M. A. & Alcalá-Quintana, R. (2015). Converging evidence that common timing processes underlie temporal-order and simultaneity judgments: a model-based analysis. *Attention, Perception, & Psychophysics, 77*(5), 1750-1766.

Hickey, C., Chelazzi, L., & Theeuwes, J. (2010). Reward changes salience in human vision via the anterior cingulate. *The Journal of Neuroscience, 30*(33), 11096-11103.

Hickey, C., Chelazzi, L., & Theeuwes, J. (2011). Reward has a residual impact on target selection in visual search, but not on the suppression of distractors. *Visual Cognition, 19*(1), 117-128.

Hikosaka, O., Miyauchi, S., & Shimojo, S. (1993). Focal visual attention produces illusory temporal order and motion sensation. *Vision research, 33*(9), 1219-1240.

Kiss, M., Driver, J., & Eimer, M. (2009). Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychological Science, 20*(2), 245-251.

Kristjánsson, À. & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, & Psychophysics*, *72*(1), 5-18.

Kristjánsson, À., Sigurjónsdóttir, Ó., & Driver, J. (2010). Fortune and reversals of fortune in visual search: Reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Attention, Perception, & Psychophysics, 72*(5), 1229-1236.

Kruijne, W. & Meeter, M. (2016). Implicit short- and long-term memory direct our gaze in visual search. *Attention, Perception, & Psychophysics, 78*(3), 761-773.

Le Pelley, M. E., Pearson, D., Griffiths, O., & Beesley, T. (2015). When goals conflict with values: counterproductive attentional and oculomotor capture by reward-related stimuli. *Journal of Experimental Psychology: General, 144*(1), 158-171.

Lee, J. & Shomstein, S. (2014). Reward-based transfer from bottom-up to top-down search tasks. *Psychological Science*, *25*(2), 466-475.

Love, S. A., Petrini, K., Cheng, A., & Pollick, F. E. (2013). A psychophysical investigation of differences between synchrony and temporal order judgments. *PLOS One, 8*(1), e54798.

MacLean, M. H., & Giesbrecht, B. (2015). Neural evidence reveals the rapid effects of reward history on selective attention. *Brain research, 1606*, 86-94.

McDonald, J. J., Teder-Sälejärvi, W. A., Di Russo, F., & Hillyard, S. A. (2005). Neural basis of auditory-induced shifts in visual time-order perception. *Nature Neuroscience*, *8*, 1197-1202.

Miranda, A. T., & Palmer, E. M. (2014). Intrinsic motivation and attentional capture from gamelike features in a visual search task. *Behavior research methods, 46*(1), 159-172.

Müller, S., Rothermund, K., & Wentura, D. (2015). Relevance drives attention: Attentional bias for gain- and loss-related stimuli is driven by delayed disengagement. *The Quarterly Journal of Experimental Psychology.*

Olivers, C. N. L. & Meeter, M. (2006). On the dissociation between compound and present/absent tasks in visual search: Intertrial priming is ambiguity driven. *Visual Cognition*, *13*(1), 1-28.

Qi, S., Zeng, Q., Ding, C., & Li, H. (2013). Neural correlates of reward-driven attentional capture in visual search. *Brain Research, 1532*, 32-43.

Raymond, J. E., & O'Brien, J. L. (2009). Selective visual attention and motivation the consequences of value learning in an attentional blink task. *Psychological Science*, *20*(8), 981-988.

Rombouts, J. O., Bohte, S. M., Martinez-Trujillo, & Roelfsema, P. R. (2015). A learning rule that explains how rewards teach attention. *Visual Cognition, 23*(1-2), 179-205.

Roper, Z. J. J. & Vecera, S. (2016). Funny money: The attentional role of monetary feedback detached from expected value. *Attention, Perception, & Psychophysics.* doi: 10.3758/s13414-016-1147-ys

Sali, A. W., Anderson, B. A., & Yantis, S. (2014). The role of reward prediction in the control of attention. *Journal of Experimental Psychology: Human Perception and Performance, 40*(4), 1654.

Scharlau, I. (2007). Perceptual latency priming: A measure of attentional facilitation. *Psychological Research, 71*(6), 678-686.

Schneider, K. A., & Bavelier, D. (2003). Components of visual prior entry. *Cognitive Psychology, 47*(4), 333-366.

Shomstein, S. & Johnson, J. (2013). Shaping attention with reward: Effects of reward on space- and object-based selection. *Psychological Science, 24*(12), 2369-2378.

Shore, D. I., Spence, C., & Klein, R. M. (2001). Visual prior entry. *Psychological Science, 12*(3), 205-212.

Spence, C., & Parise, C. (2010). Prior-entry: A review. *Consciousness and Cognition, 19*(1), 364-379.

Stelmach, L. B., & Herdman, C. M. (1991). Directed attention and perception of temporal order. *Journal of Experimental Psychology: Human Perception and Performance, 17*(2), 539.

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics, 51*(6), 599-606.

Theeuwes, J. & Belopolsky, A. V. (2012). Reward grabs the eye: Oculomotor capture by rewarding stimuli. *Vision Research, 74*, 80-85.

Theeuwes, J. & Van der Burg, E. (2013). Priming makes a stimulus more salient. *Journal of Vision*, *13*(21), doi:10.1167/13.3.21

van Eijk, R. L. J., Kohlrausch, A., Juola, J. F., & van de Par, S. (2008). Audiovisual synchrony and temporal order judgments: Effects of experimental method and stimulus type. *Perception & Psychophysics, 70*(6), 955-968.

van Eijk, R. L. J., Kohlrausch, A., Juola, J. F., & van de Par, S. (2008). Temporal order judgment criteria are affected by synchrony judgment sensitivity. *Attention, Perception & Psychophysics, 72*(8), 2227-2235.

Vibell, J., Klinge, C., Zampini, M., Spence, C., & Nobre, A. C. (2007). Temporal order is coded temporally in the brain: Early event-related potential latency shifts underlying prior entry in a cross-model temporal order judgment task. *Journal of Cognitive Neuroscience, 19*(1), 109-120.

West, G. L., Anderson, A. K., Bedwell, J. S., & Pratt, J. (2010). Red diffuse light suppresses the accelerated perception of fear. *Psychological Science, 21*(7), 992-999.

West, G. L., Anderson, A. K., & Pratt, J. (2009). Motivationally significant stimuli show visual prior entry: evidence for attentional capture. *Journal of Experimental Psychology: Human Perception and Performance, 35*(4), 1032-1042.

West, G. L., Pratt, J., & Peterson, M. A. (2013). Attention is biased to near surfaces. *Psychonomic Bulletin & Review, 20*, 1213-1220.