Failure to Exploit Learned Spatial Value Information During Visual Search

Bo-Yeong Won[1]

Andrew B. Leber[2]

1. Center for Mind and Brain, University of California, Davis, CA

2. Department of Psychology, The Ohio State University, Columbus, OH

Author Note

Support was provided by NSF BCS-1632296 to A. B. L. We thank Rayan Magsi, Beau Snoad, Rebecca Freeman, and Eleni Christofides for help with data collection for their assistance in data collection.

Address correspondence to:

Bo-Yeong Won, Center for Mind and Brain, University of California, Davis, 267 Cousteau Pl, Davis, CA 95618. E-mail: bywon@ucdavis.edu

and

Andrew B. Leber, Department of Psychology, The Ohio State University, 225 Psychology Building, 1835 Neil Avenue, Columbus, OH, 43210. E-mail: leber.30@osu.edu

Abstract

In recent years there has been rapid proliferation of studies demonstrating how reward learning guides visual search. However, most of these studies have focused on feature-based reward, and there has been scant evidence supporting the learning of space-based reward. We raise the possibility that the visual search apparatus is impenetrable to spatial value contingencies, even when such contingencies are learned and represented online in a separate knowledge domain. In three experiments, we interleaved a visual choice task with a visual search task in which one display quadrant produced greater monetary rewards than the remaining quadrants. We found that participants consistently exploited this spatial value contingency during the choice task but not during the search task – even when these tasks were interleaved within the same trials and when rewards were contingent on response speed. These results suggest that the expression of spatial value information is task specific and that the visual search apparatus could be impenetrable to spatial reward information. Such findings are consistent with an evolutionary framework in which the search apparatus has little to gain from spatial value information in most real world situations.

Keywords: visual search, choice, reward learning, spatial value

Failure to Exploit Learned Spatial Value Information During Visual Search

The principle of utility maximization, by which individuals seek the greatest rewards and smallest losses, is among the most fundamental motivators of human behavior (von Neumann & Morgenstern, 1953). With all other things being equal, we take the bet with the largest expected payout, we buy from the seller offering the lowest price, and so on. Humans do demonstrate a variety of non-normative tendencies (Tversky & Kahneman, 1975; Hastie & Dawes, 2010), but we nevertheless are consistently sensitive to value and pursue strategies to maximize our gains.

It is thus intriguing when people demonstrate insensitivity to value, as these examples offer important insights into our cognitive architecture. Such behavior can be attributed to several causes. First, people could fail to learn the relevant value contingencies; classic examples of such acquisition failures include blocking (Kamin, 1969) and overshadowing (Pavlov, 1927), in which previous or concurrent exposure of a conditioned stimulus prevents a second conditioned stimulus from being associated with an unconditioned stimulus. Second, people could successfully learn a value contingency but fail to retrieve the memory of this contingency (Tollman, 1932; Spear, 1973; see Wasserman, 1981). Third, people could decrease their subjective valuation of the reward and thus possess lesser motivation toward it, as highlighted by the classic example of a free-feeding rat reaching satiety and then abstaining from seeking food (Richter, 1922).

In this article, we consider a more perplexing scenario, in which an individual demonstrates learning, retrieval, and motivation yet only behaviorally exploits the value contingency during some task settings and not others. Such a *task-dependent expression*, should it exist, is theoretically significant because it would show domain specificity in the exploitation of reward.

Here we examine the case of spatial value and visual search. Recently, two studies using this combination of value manipulation and task, respectively, failed to show any sensitivity to reward (Jiang, Sha & Remington, 2015; Won & Leber, 2016). In both of these studies, participants were instructed to perform a search task (target T among L distractors), in which correct target identification was followed by either a small or large monetary reward. Unbeknownst to participants, the expected value (EV) of the targets varied as a function of space. For instance, targets in "high-EV" display quadrants could yield rewards averaging approximately 6 times greater value than targets appearing in the "low-EV" quadrants (e.g., Won & Leber, 2016, Experiment 1a). Note that targets appeared with equal frequency in each of the quadrants, so only the reward magnitude was varied. What the studies showed repeatedly across numerous experiments was that participants' behavior (i.e., response time and accuracy) was totally insensitive to quadrant EV.

These findings are situated in a literature on attention and reward in which dozens of studies have reported robust effects from feature-based value manipulations (e.g., Anderson, Laurent & Yantis, 2011; Della Libera & Chelazzi, 2006; Hickey et al., 2010; Kiss et al., 2009; Navalpakkam et al., 2010), while virtually no studies have reported effects of spatial reward (but see Chelazzi et al., 2014). In line with the lopsided state of the literature, we found that participants did exploit color-based value information when we modified our task to endow colors instead of spatial locations with value, using the same EV ratios (Won & Leber, 2016, Experiment 2). Thus, there appears to be a clear dissociation between the exploitation of spatial vs. non-spatial value information during visual search.

Why would people be insensitive to spatial value? Given that similar payoff schedules elicited value-sensitive performance in other tasks, it is unlikely that participants were unmotivated to seek the monetary reward. Instead, it is possible that the act of conducting visual search interferes with either learning or retrieval of the spatial value contingency. Given the brain's capacity limitations, it is simply not possible to encode or retrieve task-relevant information about all potential relationships among variables. Such an encoding-based explanation suggests that people would demonstrate sensitivity to spatial reward during search if they were just able to properly represent the relevant information. Alternatively, expression of spatial value knowledge could be task dependent; that is, observers could know the spatial value contingencies but the visual search apparatus could be impenetrable to incorporating this information. Such a scenario represents the most intractable form of reward insensitivity, because any method to endow the observer with knowledge of the contingency will inevitably fail to circumvent the mode of expression and will consequently fail to produce behavioral change.

While it seems counterintuitive that the human mind would selectively fail to express actively represented knowledge, it may be sensible from an ecological standpoint. Consider that in the real world, objects are usually the things to which we assign value, not their locations. When we engage in a search task for a specific item, such as our car key, the target of our search is just as valuable to us if we find it on the desk or hanging from the doorknob. Because of the ecological oddity of a visual search target varying in value as a function of its location, the designer of a visual search apparatus – i.e., the process of natural selection – may not have been pressured to incorporate spatial reward sensitivity as a necessary component. It has previously been argued that the visual search apparatus does not take all possible information into account when deploying shifts of attention; for instance visual search might be more efficient when it is allowed to run in “anarchic” versus “orderly” fashion (Wolfe, Alvarez, & Horowitz, 2000), and it might sometimes proceed without concern of revisiting previously rejected locations (Horowitz & Wolfe, 1998; Wolfe, 2003).

Many forms of learning have been shown to modulate the visual search process, including phenomena such as contextual cueing and probability learning (e.g., Chun & Jiang, 1998; Miller, 1988; Geng & Berhmann, 2002). Moreover, individuals are highly sensitive to variations in the likelihood of target appearance, as shown in visual foraging tasks (e.g., Cain, Vul, Clark & Mitroff, 2012; Wolfe, 2013). Finally, individuals are plenty capable of claiming volitional control over their search when they want to; apropos to this discussion, Jiang et al. (2015) showed that a subset of participants were able to direct attention to high-value locations when asked to do so. Yet, in this paper, we entertain the notion that, by default, the visual search apparatus is impenetrable to spatial value information, even when the observer is actively representing such value within another knowledge domain.

Before proceeding in this venture, we must offer an important disclaimer: proving that the visual search apparatus is impenetrable to spatial value information is a tall order, which might require a substantial number of studies, using a variety of converging approaches. Therefore, we will refrain from making overly strong claims about the impenetrability of the search apparatus while presenting our current work. We will emphasize that the discussion of this theoretical possibility is important and our results provide tentative support for it.

Specifically, we conducted three experiments to seek support for whether insensitivity to spatial value during visual search constitutes a task-dependent expression failure. Our general approach was to mix together two types of tasks: Visual Choice and Visual Search. During the choice task (see, e.g., Jiang, Swallow, Won, Cistera & Rosenbaum, 2015), participants were able to select a single item (Exp. 1) or a subset of the display (Exp. 2 and 3) in the aim of obtaining a reward. We varied potential reward as a function of spatial location, and we anticipated that participants would rapidly learn to choose high-EV locations, replicating an earlier finding (Won & Leber, 2016, Exp. 4b). The visual search task was like that used in previous work, in which targets were assigned value depending on their spatial location. Critically, the locations assigned as high-EV were identical across the choice and search tasks within each participant. Assuming that the participants show a clear preference for the high-EV locations during choice, we reasoned that a) this would demonstrate that observers are actively representing spatial value information during choice and b) the information would be carried into the interleaved search task. Given these assumptions, if previous findings of reward insensitivity during search resulted only from learning, retrieval, or motivational deficits, then the current manipulations should overcome these limitations and facilitate the expression of the value contingency knowledge during search. If, however, the visual search apparatus is insensitive to spatial value, participants will persist in demonstrating a task-dependent expression failure during search while successfully expressing learning during choice. To preview, none of the current manipulations succeeded in eliciting the expression of spatial value knowledge during search. These results provide intriguing support for the notion that our search apparatus is, by default, impenetrable to spatial value information.

Experiment 1: Interleaved Visual Choice and Visual Search

In the first experiment, we interleaved 16-trial mini-blocks of Visual Choice and Visual Search. During Visual Choice, participants viewed a display of 16 "L" stimuli and clicked on any of them to receive a reward. For each participant, we designated a single High-EV quadrant, for which any object clicked inside of it would typically return a high reward. Because the remaining quadrants typically returned low rewards, we expected to replicate our previous finding that participants quickly learn to disproportionately choose the high-EV quadrant (Won & Leber, 2016). The Visual Search task had stimuli that were nearly identical to those in the choice task, except one of the search items was a T, and participants were told to click on it rapidly and accurately.

As we described in the introduction, we expected that participants would bias behavior toward the high-EV quadrant during the Visual Choice trials, and we thus sought to determine whether the participants would express this knowledge during the Visual Search task; such expression would be manifested via faster response times (RTs) on trials in which target appeared in the high-EV quadrant versus the low-EV quadrant.

Method

Participants

In choosing our sample size, our goal was to include a sufficient number of participants to obtain power = 0.9 for the visual choice manipulation. Based on the effect size obtained for this manipulation in our previous paper (Won & Leber, 2016, Experiment 4b), we estimated that 15 participants would be needed.[1] We ran 16 participants in all experiments reported in this manuscript. Of the 16 run in Experiment 1, 12 were female (mean age = 22.7 years). All participants reported normal or corrected-to-normal visual acuity and normal hearing. The Ohio State University IRB approved this protocol. Compensation for the 1.5 hr session was based on how many points were earned during the experiment (point values will be further explained in the Design). Participants' payouts were distributed in whole dollar amounts, as follows: 0-5000 points = $15; 5001-6000 points = $16; 6001-7000 points = $17; 7001-8000 points = $18; 8001-9000 points = $19; 9001 and up = $20.

Apparatus and Stimuli

Participants were tested in a dimly lit room. Stimuli were presented on a 24" LCD monitor and generated using MATLAB (www.mathworks.com), with Psychtoolbox extensions (Brainard, 1997; Pelli, 1997). In the Visual Search task, displays contained one target (a white T rotated 0°, 90°, 180°, or 270°) and 15 distractors (white Ls rotated 0°, 90°, 180°, or 270°) on a gray background. Targets and distractors subtended 1.02°x1.02° (all visual angles assume a typical viewing distance of 60 cm). Item locations were chosen randomly from a 10x10 invisible matrix (15.28°x15.28°), with four items appearing in each quadrant. Target and distractor orientations were all selected randomly with replacement on each trial. The number 20 or 1 (font size: .92°), indicating reward points for a given trial, was displayed at the target location, in green. Auditory feedback was either a three "chirp" sequence lasting 300ms for 20-point responses, a single high-pitched 100ms tone for 1-point responses. In the Visual Choice

[1] We do acknowledge, however, that estimates of effect size and power from previous data are prone to inflation, due to a "winner's curse," in which researchers are biased to primarily follow up on positive results (Halsey et al., 2015). Therefore, our true power may have been less than 0.9.

task, stimuli were identical to the search task except that choice displays contained 16 Ls and no Ts.

Design

In the Visual Search task, the target appeared with equal frequency across the four quadrants (25% of trials each), but the point-value was varied. Specifically, targets in the *high-EV* quadrant earned 20 points on 75% of trials and 1 point on 25% of trials. Targets in the other three *low-EV* quadrants earned the opposite: 1 point on 75% of trials and 20 points on 25% of trials. These contingencies yielded individual quadrant EVs of 3.813 and 1.438 for *high-* and *low-EV* quadrants, respectively. In the Visual Choice task, we matched the payoffs to the Visual Search task. Specifically, objects clicked in the *high-EV* quadrant yielded 20 points on 75% of trials and 1 point on 25% of trials; objects clicked in the *low-EV* quadrant yielded 1 point on 75% of trials and 20 points on 25% of trials (see Figure 1A). Critically, the location of the high-EV quadrant was the same across the search and choice tasks for each participant.

Procedure

Main Trials.

Participants completed 46 mini-blocks, each 16 trials, which consisted of alternating visual search and visual choice 23 times. Half of the participants started with visual search and the other half started with visual choice. In both tasks, participants initiated each trial by clicking on a small white square (.51°x.51°), which appeared near the screen center (jittered randomly on each trial by +/- .77° in both vertical and horizontal directions). After the click and a 500 ms delay, the search or choice display appeared (for the two respective tasks).

Participants were instructed to click on the target in the search task and click on any object in the choice task. Upon response, the display was removed, and the point-value earned was displayed, along with the auditory feedback. Next, the cumulative total points were displayed at the screen center for 500 ms (see Figure 1B). Participants completed 16 practice trials in each task before advancing to the main task.

Explicit Learning Assessment.

After the main trials, participants were told about the spatial reward manipulation and asked to complete a 16-trial *generation task* (similar to Chun & Jiang, 2003) for each task, which assessed their explicit knowledge of the quadrant EVs. Participants were shown a search or choice display (depending on the task) and were required to click on the target in the search task or a randomly circled L in the choice task. In both cases, this response then revealed two point values, 1 and 20. Participants had to then choose which of these they felt to best match the reward typically earned at that location. No feedback was provided.

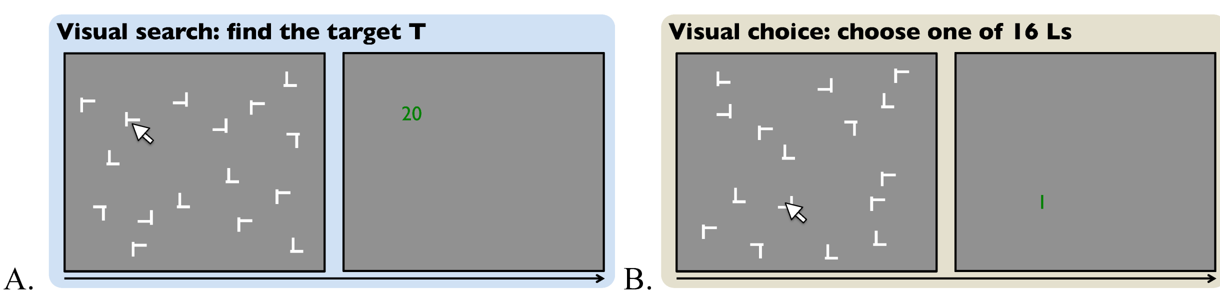

C.
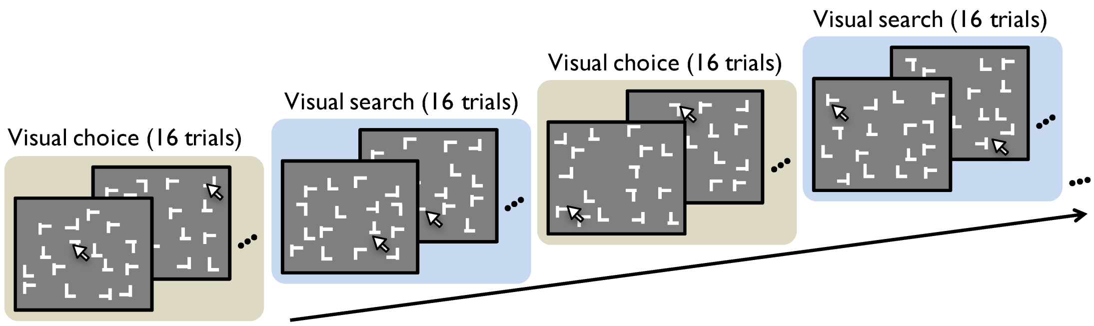

Figure 1. Design and procedure in Experiment 1. A. Choice task. Participants chose one of 16 randomly rotated Ls. The High-EV quadrant yielded 20 points on 75% of trials and 1 point on 25% of trials, while the Low-EV quadrants had the opposite payoff. B. Search task. Participants searched for a target T among 15 randomly rotated L distractors. Targets appearing in the High-EV quadrant were worth 20 points on 75% of trials and 1 point on 25% of trials, while targets in the Low-EV quadrants yielded the opposite payoff (see method for additional details). C. A schematic procedure of Experiment 1. The two tasks were presented in a sequence of 23 alternating mini-blocks (16 trials per mini-block).

Results and Discussion

Visual choice

Choice Frequency. The *choice frequency* data across quadrant type and blocks are plotted in Figure 2 (bar graph). It was calculated as the proportion of trials in which items in each of the quadrants was chosen. For the low-EV data, we collapsed across the three quadrants of this type for each participant. Because low- and high-EV data were statistically dependent, our analysis approach was to compare high-EV choices vs. chance level. A sample to hypothesized mean (of 0.25) t-test, on collapsed data from all blocks, was significant, $t(15) = 5.04$, $p < .001$, Cohen's d = 1.26. To then determine if participants' sensitivity to the high-EV quadrant changed over time, we conducted a 1-way ANOVA on choice data across 23 blocks. The results showed a significant main effect, $F(22, 330) = 7.48$, $p < .001$, $\eta_p^2 = .33$. To test whether the sensitivity to the high-EV quadrant increased over time, we entered the high-EV choices in each block into a linear regression, for each participant, and computed the slope values. These slopes were entered into a one-sample t-test (compared to a hypothesized mean of zero). Results demonstrated a significantly positive linear trend, $t(15) = 4.77$, $p < .001$, Cohen's d = 1.19, consistent with increased sensitivity to the high-EV quadrant over time.

Visual search

Because we only accepted click responses when the mouse hovered over the target, errors were not possible, so accuracy data are not reported. We thus focus solely on RT.

RT. We removed trials with RTs slower than 3 SD above the mean (1.7%), and the remaining mean RTs are plotted in Figure 2 (line graph). The quadrant type x block ANOVA revealed a main effect of block, as RT became faster as the experiment progressed, $F(22, 330) = 3.50$, $p < 0.0001$, $\eta_p^2 = .19$. However, there was neither a main effect of quadrant type, $F < 1$, $\eta_p^2 = 0.024$, nor a 2-way interaction, $F(22, 330) = 1.29$, $p = 0.174$, $\eta_p^2 = 0.079$, showing that participants were insensitive to spatial value during search.

One of potential issues with this is that failing to reject the null hypothesis could be due to limitations in statistical power. One approach to increase confidence in supporting the null hypothesis is to compute the Bayes Factor (BF), which generates a readily interpretable odds ratio of evidence for vs. against the null hypothesis (e.g., Rouder, Speckman, Sun, Morey, & Iverson, 2009). We computed BF on visual search RTs[2], which yields $BF_{01} = 3.34$. The BF indicates that the observed data were 3.34 times more likely to be observed if the null hypothesis were true than if the alternative hypothesis were true. (for guidelines in the

[2] The Bayes factors are written as BF_{10} when the evidence is in favor of H1 and as BF_{01} when the evidence is in favor of H0. We computed the BF using JASP 0.8.1.2 (JASP Team, 2017), with the default prior width of 0.707.

interpretation of Bayes factor magnitudes, see Morey, Rouder, & Jamil, 2014; Rouder, Morey, Speckman, & Province, 2012; Wetzels, Matzke, Lee, Rouder, Iverson, & Wagenmakers, 2011).

Visual search vs. visual choice

We next wanted to compare visual search and visual choice results. While it is difficult to directly compare RT and choice frequency, our approach was to convert these respective dependent measures of interest to arbitrary *learning efficiency units* for each block (see Figure 3). Specifically, for visual search, we subtracted the mean RT in *high-EV* quadrant from that in *low-EV* quadrants, and then divided by the sum of the two mean RTs from *high-EV* quadrant and *low-EV* quadrant (i.e., *low-EV* quadrant's RT – *high-EV* quadrant's RT) / (*low-EV* quadrant's RT + *high-EV* quadrant's RT)). For visual choice, we subtracted expected chance level choice frequency (0.25) from choice frequency in *high-EV* quadrant, and then divided by the sum of these two measures (i.e., (*high-EV* quadrant's choice – *0.25*) / (*high-EV* quadrant's choice + *0.25*)). We then conducted a pairwise t-test on the two learning efficiency measures, revealing a significant difference between the two, $t(15) = 7.01$, $p < .001$, Cohen's d = 1.75.

We interpret this result with some caution, as we cannot verify that the choice and search manipulations were equally strong and/or sensitive. And while converting the respective choice and RT measures learning efficiency units places the two on a common scale, it does circumvent the inherent limitations in our comparison. Keeping these points in mind, the results suggest stronger expression of learning in visual choice than visual search.

Explicit Learning Assessment

A task x quadrant type ANOVA on the generation task responses revealed that participants chose 20 points more often in *high-EV* quadrant than in any of *low-EV* quadrants, $F(1, 15) = 10.29$, $p = 0.006$, $\eta_p^2 = .41$. Also, we found a significant interaction between task and quadrant type, $F(1, 15) = 14.30$, $p = 0.002$, $\eta_p^2 = .49$, which was driven by stronger recognition in visual choice than in visual search. Specifically, for visual choice, participants chose 20 points 75.0% of the time in the *high-EV* quadrant vs. 32.8% of the time in the *low-EV* quadrants. In contrast, for visual search, participants chose 20 points 53.1% of the time in the *high-EV* quadrant vs. 44.8% of the time in the *low-EV* quadrants. When comparing recognition performance with chance (50%), visual choice was significantly more above chance (69.1%), $t(15) = 4.68$, $p = 0.0003$; visual search was not (54.7%), $t(15) = 1.31$, $p = 0.210$. Clearly, the participants developed some degree of explicit knowledge for the choice task while failing to do so for the search task.

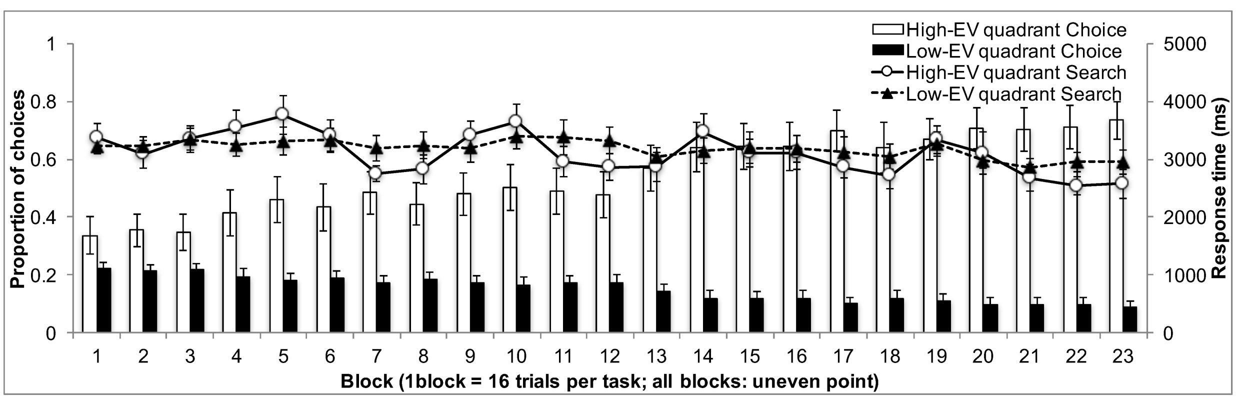

Figure 2. Results from Experiment 1, showing RT in visual search (line graph) and the choice frequency in visual choice (bar graph) as a function of quadrant type, across blocks. Error bars show ±1 S.E. of the mean.

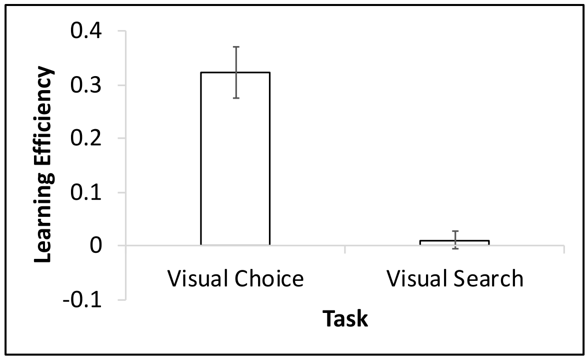

Figure 3. Learning efficiency of visual choice and visual search. Error bars show ±1 S.E. of the mean.

Experiment 2: Combined Choice and Search Within Trials

The first experiment showed a clear dissociation: even though the high-EV location was fixed across both tasks, participants prioritized it during choice but ignored it during search. This is consistent with the task-dependent expression account, in that participants demonstrated sufficient knowledge and active retrieval of the high-EV location, and they demonstrated sufficient motivation to prioritize that location, yet they failed to use the value information during search. However, there is one key limitation to this experiment: it is possible that the brief breaks in between the interleaved mini-blocks were long enough to prevent transfer between the tasks. Specifically, we presume that participants were operating with actively represented value information during the choice trials, but it could have rapidly decayed at the end of the 16-trial mini block; thus, the value information might not have been actively represented during the search trials.

To attempt to overcome this limitation, Experiment 2 combined both choice and search components within single trials. When the display appeared, observers now had to choose to search one side of the display vs. the other. After making this choice, the search items appeared – only within the chosen display side – and observers had to click on the target T. Critically, a single high-EV quadrant was assigned to each participant that could motivate both the choice and search tasks similarly. That is, to get the high reward, participants would first need to click on the correct side of the display; then to demonstrate a bias toward that same reward, participants would need to prioritize search toward that high-EV quadrant. Note that across blocks, we varied the choice task to be either upper (top) vs. lower (bottom) or left vs. right. Thus, a participant seeking to maximize reward would need to know the specific high-EV quadrant and not just a high-EV side; for instance, if the upper left were high-EV, the participant would need to choose top on some blocks and left on others.

If observers demonstrate actively retrieved knowledge of the high-EV quadrant during the choice task, we could test whether that information would be expressed moments later during the search task.

Method

Participants

Sixteen individuals participated in Experiment 2 (9 females; mean age = 19.9 years), which took approximately 1 hour. Participants' payouts were determined as follows: 0-3000 points = \$10; 3001-4000 points = \$11; 4001-5000 points = \$12; 5001-6000 points = \$13; 6001-7000 points = \$14; 7001 and up = \$15.

Material, stimuli, and procedure

The experiment was modeled off of Experiment 1, with the following changes. During the choice component, top/bottom displays contained two dots (.51°x.51°) in the center of the upper and lower sides of display, while left/right displays contained dots in the center of the left and right sides of the display (see Figure 4). Top/bottom and left/right displays were alternated across blocks, with half of the subjects doing left/right on even blocks and half doing them on odd blocks. Participants were instructed that each side has its own target, and they were asked to click which side they preferred to search. Once they clicked on a dot, that side of search array was revealed. Participants completed 24 blocks (32 trials per block). During the first half of these trials (blocks 1-12), quadrant EV was manipulated to match the same EVs used in Experiment 1. During the second half (block 13-24) all four quadrants' EVs were made to be equal, so that we could examine any persisting effects of learning.

<u>Explicit Learning Assessment</u>

We focused here on assessing explicit learning during search. The generation task was identical with that for visual search's generation task in Experiment 1, except that in Experiment 2, participants were shown two dots – white and black – before the search display, and asked to click the white dot to reveal that side of search display. This ensured that participants would provide an equal amount of data for each side. Additionally, we presented this task in two blocks of 24 trials each; one block was top/bottom and the other block was left/right. The order of blocks was counterbalanced across participants. Once the search display was revealed, participants clicked on the target T to reveal two point values, 1 and 20; they then chose which one best matched the reward typically earned at this location.

A.

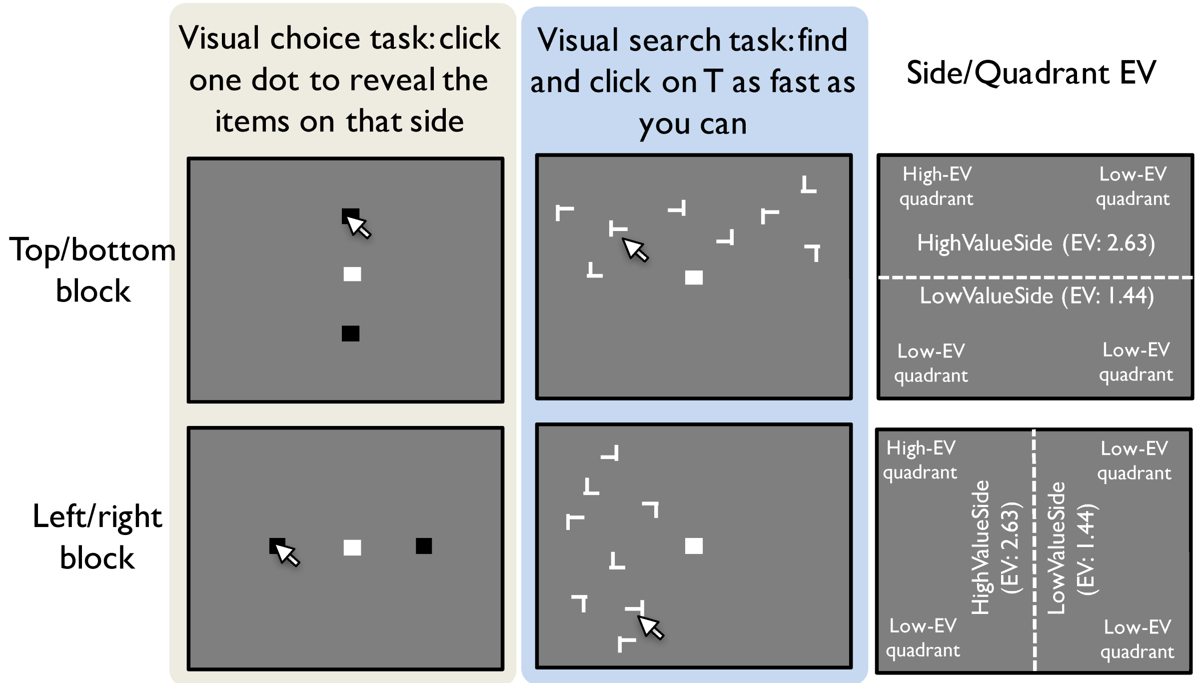

B.

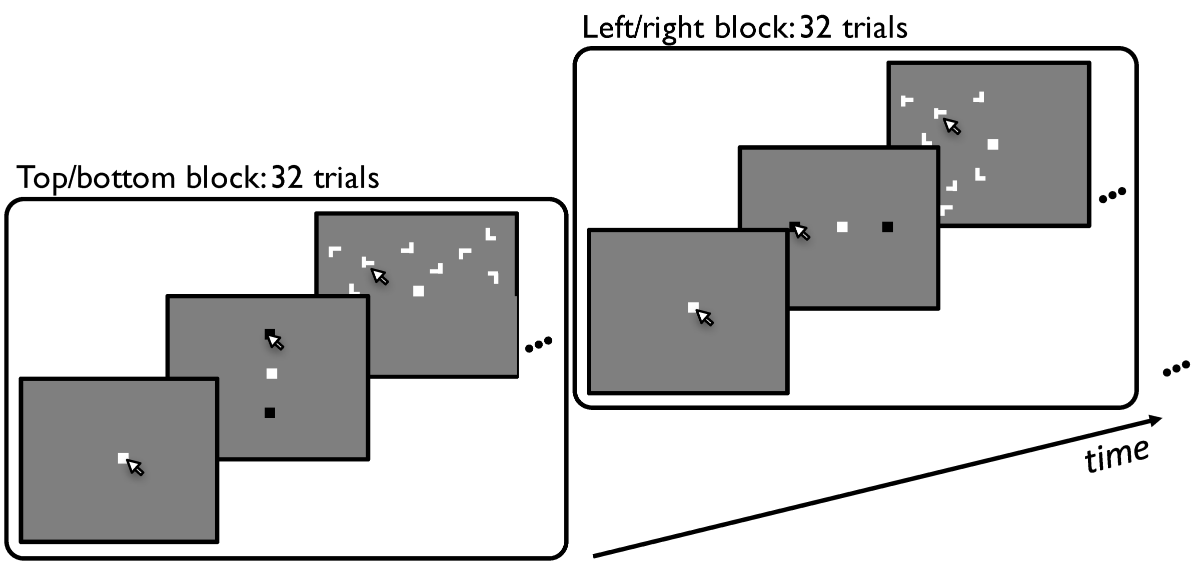

Figure 4. *Design and procedure in Experiment 2. A. In top/bottom blocks, two dots were displayed in upper side and bottom side, each. Participants clicked one dot to reveal that side of search display. In the left/right blocks, two dots were displayed in left and right side, each. Which block – odd or even – was assigned as top-down dots or left-right dots were counterbalanced across participants. Side/Quadrant EVs indicate sample reward contingencies based on location. Note that the high-EV quadrant was counterbalanced across participants. B. Sample trials of Experiment 2. Top/bottom blocks and left/right blocks were alternating 12 times.*

Results and Discussion

Visual choice

Choice frequency. Analysis focused on chosen side. We coded *low-EV side* as the one consisting of two low-EV quadrants and *high-EV side* as the one consisting of one low-EV quadrant and the high-EV quadrant. Note that we collapsed across top/bottom and left/right blocks. Again, we tested high-EV choices against chance level, because low- and high-EV data were statistically dependent. A sample to hypothesized mean (of 0.5) t-test, on collapsed data from all blocks, was significant, $t(15) = 2.79$, $p = .014$, Cohen's d = .70. To then determine if participants' sensitivity to the high-EV side changed over time, we conducted a 1-way ANOVA on choice data across 24 blocks. We did not find any significant main effect of block, $F(23, 345) = 1.07$, $p > .3$, which might be because learning emerged rapidly, in the earliest blocks.

We considered the possibility that each participant could have learned the bias from only the top/bottom (or left/right) blocks and totally ignored the left/right (or top/bottom) blocks. This is important because we inferred from the choice performance above that participants represented their high-EV quadrant rather than just a single display side. Critically, had the participants only represented a display side, the visual search performance would not be expected to vary between the high- and low-EV quadrants. To address this concern, we generated a scatterplot in which the y-axis indicates high-EV preference in the top/bottom blocks and the x-axis indicates high-EV preference in the left/right blocks (Figure 6). If participants represented the high-EV quadrant rather than only one side, the dots should all settle into the upper-right quadrant of the scatterplot; numerically speaking, this occurred for 13/16 participants. To test this preference statistically, we rank ordered the magnitude of the side preferences for each participant, (for some, the stronger numerical preference was for left-right and others it was for top-down). Then we tested the correlation of the stronger vs. weaker preference. Our logic was that, if participants only had a side bias, then the weaker preference score would not carry any meaningful information and the correlation would thus not be significant. However, if both stronger and weaker preferences scaled with participants' overall quadrant preference, then the correlation would be significant. Indeed, it was significant, $r(15) = 0.943$, $p < 0.00001$. However, this result might have been driven partially by one outlier who perfectly chose the high-EV side for both odd and even blocks and another participant who reliably chose the low-EV side. We repeated the test while excluding these individuals, and the result remained significant, $r(13) = 0.858$, $p < 0.001$. All told, this analysis provides evidence that the participants' choices were biased toward a quadrant, not a single side. It is worth noting that this result does not necessarily reflect that the participants "integrated" the representation of quadrant value from two types of trials (i.e., top/bottom and left/right trials). Whether the participants represented the quadrant value will need to be fleshed out in future experiments.

Visual search

RT. Analysis focused on target click RT, which commenced at the moment the participant completed the choice (by clicking one of the two dots). RT trimming removed 1.2% of trials. Mean RTs across quadrant types and blocks are shown in Figure 5 (line graph). A quadrant type x block ANOVA revealed a main effect of block, $F(23, 253) = 2.56$, $p < .001$, $\eta_p^2 = .19$, meaning the overall search RT became faster as the experiment progressed. However, replicating the visual search in Experiment 1, there was no RT difference between the two quadrant types, $F < 1$, $BF_{01} = 3.90$, demonstrating a failure to exploit the value information during search. Quadrant type and block did not interact, $F(23, 253) = 1.53$, $p = .062$, $\eta_p^2 = .12$.

Explicit Learning Assessment

During the generation task, on trials in which the high-EV side was presented, participants did not reliably vary in their selection of 20 points vs. 1 point across the *high-EV* (59.7% vs. 40.3%) and *low-EV* (53.5% vs. 46.5%) quadrants, $F(1, 15) = 1.47$, $p = .244$, $\eta_p^2 = .09$ The recognition performance (49.8%) was not higher than chance (50%), $t(15) = .25$, $p = .81$.

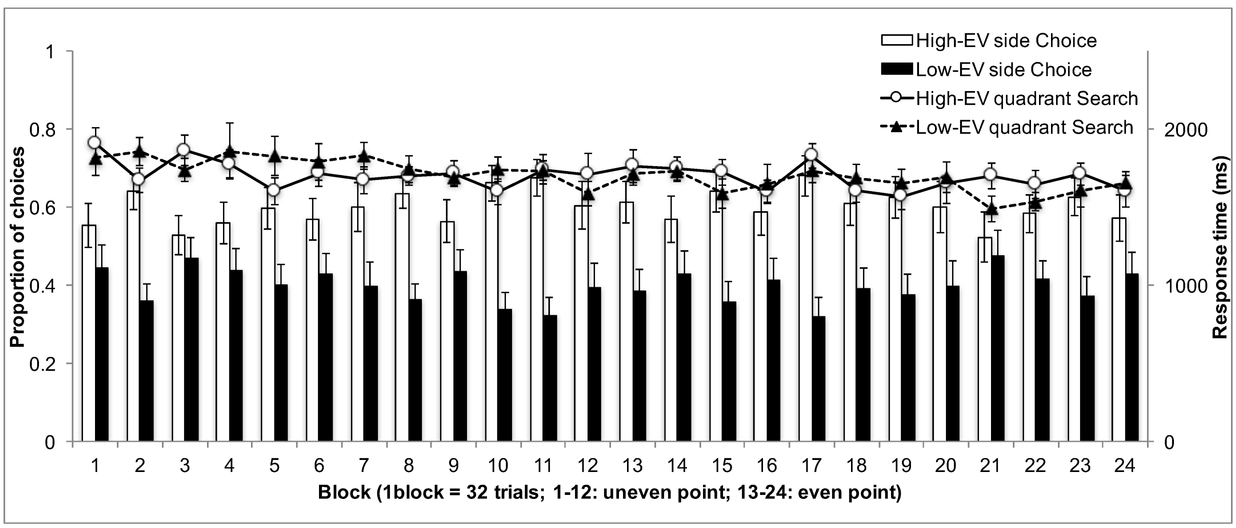

Figure 5. Results from Experiment 2 showing RT in visual search (line graph) and the choice frequency in visual choice (bar graph). Error bars show ±1 S.E. of the mean.

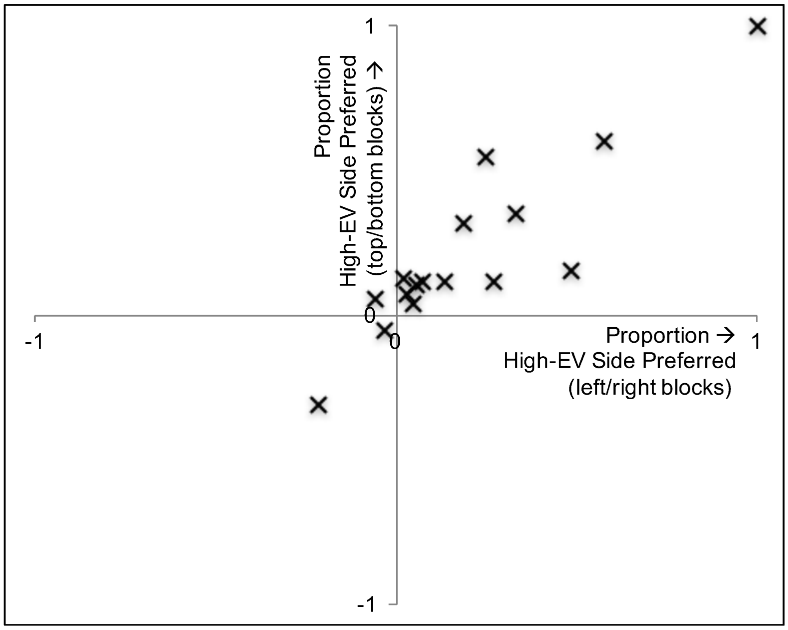

Figure 6. High-EV quadrant preference in two block types – left/right blocks and top/bottom blocks. Scores are calculated as the different in proportion of trials in which the High-EV side was chosen minus the proportion of trials in which the Low-EV side was chosen. Each "X" represents one participant.

<u>Visual search vs. visual choice</u>

To compare results from visual search and visual choice, as done in Experiment 1, we again calculated *learning efficiency units* (see Figure 7). A pairwise t-test revealed a significant difference between the two tasks, $t(15) = 2.77$, $p = .014$, Cohen's d = .69, showing a reliably stronger learning effect during choice than during search.

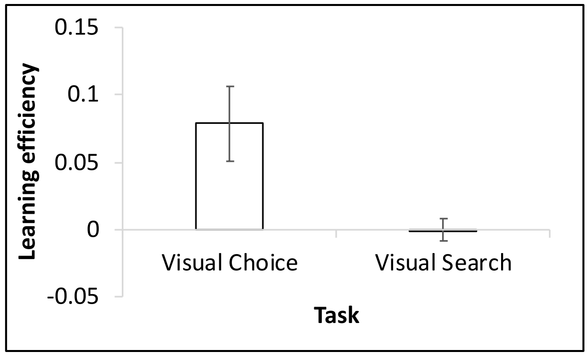

Figure 7. Learning efficiency of visual choice and visual search. Error bars show ±1 S.E. of the mean.

Experiment 3: Combined Choice and Search for Single Target

Experiment 2 showed that high-value locations could be represented and acted upon via Visual Choice while neglected during Visual Search *within the same trial.* Along with the results of Experiment 1, these results are consistent with the task-dependent expression account.

Nevertheless, there are a few remaining potential limitations we must address. One such limitation is that, even though we fixed the same high-EV quadrant across choice and search, it is possible that participants did not view the two tasks to be related to one another. This is because the choice task produced a target on either display side that was clicked, which stands in contrast to the visual search task, whose target is pre-determined to appear in only one location. There is some evidence that learning could fail to transfer across a choice and search task if they are viewed as contextually dissimilar to one another. For instance, Jiang, Swallow, et al. (2015) used a training/transfer procedure to see if a location preference learned during a spatial probability cueing manipulation – in which targets appeared more frequently in one display quadrant than others – transferred to a visual choice task. They tested whether learning from visual choice task would transfer to visual search. Learning did occur within each task type, but the critical result showed no transfer in either direction. Thus, the two tasks were represented as distinct from one another. That study differed from the present work in a few important ways. First, it used probability cueing instead of a reward manipulation; second, its design included hundreds of training trials followed by hundreds of test trials, whereas we interleaved or combined our tasks. Nevertheless, the possibility remains that participants simply viewed our choice and search tasks as unrelated to one another and thus did not transfer learning from choice to search. To address this concern, we now told participants that they had to find a single target on each trial, which would be completed in two steps. First, the participants would only be able to reveal half of the display, by clicking on any two quadrants. If neither of the quadrants they chose on a given trial contained the target, the trial would be terminated with no reward, and they would advance to the next trial. If, however, their chosen quadrants did include the target location, the participants would then proceed to the search portion of the trial. For the search portion, we revealed items in both of the chosen quadrants, and participants had to click on the target. This procedure allowed us to determine if a) participants developed a bias toward the high-EV quadrant during choice and b) if they prioritized the same high-EV quadrant during search.

Additionally, we addressed one further potential drawback of Experiments 1 and 2, which is that the outcome of selecting high-EV locations during choice is much more consequential than for prioritizing search toward the high-EV location. This is because, in the previous experiment, the choice determined the payoff on each trial; that is, choosing high-EV locations more frequently resulted in increased overall earnings for the participants. In contrast, directing visual search toward the high-EV location before inspecting the other locations may not have had as great of an impact on earnings. For instance, in Experiment 1, participants searched the full display exhaustively until they found the target, which was the only way for them to advance to the next trial. Because the target was equally likely to appear in all four quadrants, directing the search initially to the high-EV quadrant would not, on average, speed target identification compared to searching a low-EV quadrant first. In Experiment 2, the same logic can be applied, albeit within the chosen side of the display. As a result, even if participants knew where the high-EV location was, it would not necessarily have benefitted them to prioritize their search one way or another.

Note that the same criticism can be leveled at experiments showing color-based prioritization, yet those experiments generate robust expression of learned value. For instance, as mentioned in the introduction, we previously ran an experiment using the same 16-item T among L displays in which target color, not location, was associated with reward (Won & Leber, 2016, Exp. 2). We matched the reward ratios in that experiment (i.e., high-EV color vs. low-EV colors) to a spatial reward manipulation much like Exp. 1 of this article. While participants could gain no monetary advantage for initially prioritizing search within the high-EV color, they did so anyway, showing faster RTs to high-EV targets. This is reminiscent of a study by Goldstein & Spence (1963), in which rats were placed in one of two 60-inch alleys that contained a food reward at the other end. One alley consistently had a greater reward than the other, which rats learned, and they expressed this learning by running faster when placed in the high-reward alley than the low-reward one. Thus, even though the rats could not alter their behavior to influence the reward outcome, they demonstrated greater motivation in one condition than another. Given these results, we should have reasonably expected our observers to search the high-EV locations more rapidly than low-EV locations in the previous experiments.

Nevertheless, there is a clear way to experimentally address this concern. Indeed, the two papers that previously failed to find spatial reward effects each included a manipulation to incentivize participants to prioritize search to the high-EV location (Jiang, Sha & Remington, 2015, Exp. 3; Won & Leber, 2016, Exp. 3). Jiang and colleagues only rewarded trials in which RTs were faster than each participant's median RT; in this case, strategically searching the high-EV quadrant first would yield greater overall rewards than searching low-EV quadrants first. Won and Leber limited display exposure, which reduced overall accuracy and provided a similar incentive in which prioritizing the high-EV quadrant would bring greater overall earnings (since rewards were only provided on correct trials). In the current experiment, we adopted Jiang et al.'s approach by only providing a reward on faster-than-median trials during search.

Altogether, this experiment maintained the same approach as Exp. 2, in that it provided for the active representation of spatial value information during choice, which we assumed to be available a moment later during search. Moreover, it merged the task context across choice and search, as participants now sought the same target on each trial. And finally, the incentive to exploit the value information to maximize reward was roughly equivalent in both the choice and search components of the task.

Method

<u>Participants</u>

We ran sixteen individuals (9 females; mean age = 23.2 years), which took approximately 1.5 hours. Participants' payouts were determined as follows: 0-3000 points = $15; 3001-4000 points = $16; 4001-5000 points = $17; 5001-6000 points = $18; 6001-7000 points = $19; 7001 and up = $20. This payoff schedule was slightly modified from the previous experiments. This was because we made rewards dependent on search RT, which reduced the total number of points earned; the modified payoff schedule ensured comparable total dollars earned per hour across the experiments.

Material, stimuli, and procedure

The first half of the experiment began with the same Visual Search used in Exp. 1. In particular, we did not condition reward based on search RT; our goal was to have the greatest possible chance of forming an association between quadrant and reward magnitude, and providing rewards on every correct trial maximized this. We also wanted to familiarize participants with this task before layering on the choice component; this served to reinforce the participants' interpretation that choice and search were part of the same task.

In the second half of the experiment, each trial began with the presentation of four black dots, each centered in one of the four quadrants. Participants were told to click any two of the dots to reveal the search items in those quadrants, then click on a small white square which appeared near the screen center to initiate search. If neither quadrant contained the target (which was equally likely to appear in each of the four quadrants), the message “no T” (font size: 0.92°) appeared for 1 sec, before proceeding to the next trial. If one of the two chosen quadrants did contain the target, the search items in those chosen quadrants were revealed. Participants then had to click on the target, at which point their RT was calculated. For trials in which the RT was slower than the previous block's median RT, the message "too slow" appeared in white, along with the point-value they would have earned, which appeared in red. For trials in which the RT was faster than previous block's median, a reward was provided, with the point value displayed in green. Note that we needed a full block of this combined choice/search task to calculate a stable median RT, so the median cutoff was not used until the 2nd block of this phase (i.e., the Block 13 of the experiment).

There were 24 total blocks, with 32 trials per block. The first 12 blocks consisted of the search task only, and the remaining 12 blocks included the combined choice and search task (see Figure 8 for design and procedure details). The explicit learning assessment was identical with that of the search task in Experiment 1.

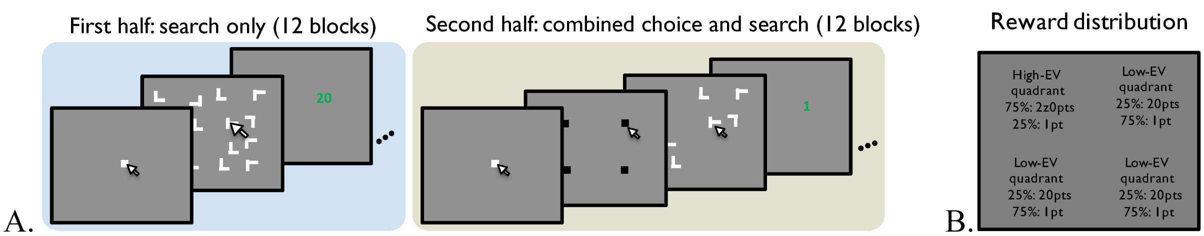

Figure 8. Design and procedure in Experiment 3. A. The first half of the experiment (blocks 1-12), consisted of visual search. In the second half of experiment (blocks 13-24), participants clicked any two of four presented dots (one in each quadrant) to be able to search for the target within those chosen quadrants. When targets were clicked on, rewards were determined by 1) the target quadrant's EV and 2) whether the response was faster than the previous block's median RT. B. Sample reward contingencies in each quadrant. The actual location of the high-EV quadrant was counterbalanced across participants.

Results and Discussion

First half: visual search only

RT. Analysis focused on target RT. Trimming removed 1.8% of trials, and mean RTs across quadrant types and blocks are shown in Figure 9A. A quadrant type x block ANOVA revealed a main effect of block, $F(11, 165) = 2.07$, $p = 0.014$, $\eta_p^2 = .12$, meaning the overall search RT became faster as the experiment progressed. The main effect of quadrant type was significant, $F(1, 15) = 8.84$, $p = 0.0095$, $\eta_p^2 = .37$. This result was very surprising, given the multiple previous failures of us and others to produce any evidence of spatial value exploitation during search under highly similar conditions. The effect of quadrant did not increase across blocks, as shown by a non-significant quadrant x block interaction, $F(15, 165) = 1.06$, $p = 0.397$. Moreover, an inspection of the means shows the largest numerical effects of quadrant to appear in the first phase, when the manifestation of learning should be the weakest. When excluding these two blocks from the ANOVA, the main effect is no longer significant, $F(1, 15) = 2.191$, $p = 0.160$, $\eta_p^2 = 0.13$.

Overall, the results from the first half of this experiment are not easily interpreted as support for the expression of spatial reward learning. However, to ensure that we were not too quick to dismiss this result, we ran a more powerful analysis across pooled results from five total experiments run under highly similar conditions. These experiments all contained a T among L search task in which quadrant value was manipulated, the displays had unlimited exposure, and rewards were provided on all correct trials. The analysis included: Experiment 2 from Jiang et al. (2015), who kindly shared their data (blocks 1-8, n=16); Won and Leber (2016), Experiment 1a (all 12 blocks, n=12); Won and Leber (2016), Experiment 3 (blocks 1-4, n=12); Won and Leber (2016), Experiment 4a (blocks 1-12, n=12); and the current Experiment 3 (blocks 1-12, n=16). Given the highly similar experimental conditions across experiments, we simply ran a paired-samples t-test on the pooled group of 68 participants, comparing mean RT on the lowest EV vs. highest EV quadrants. Results showed a total mean quadrant effect of M = 20.7 ms, SD = 46.3, which was not significant, $t(68) = 0.351$, $p = 0.727$, $d = 0.043$. This analysis increases our confidence that the surprising main effect of quadrant in the first half of Experiment 3 was likely a Type I error.

Second half: Visual choice

Choice frequency.

Analysis focused on the frequency of trials in which participants chose the high-EV. Note that while we requested participants choose two quadrants on each trial, our experimental code allowed for them to click on the same quadrant twice; we excluded such trials from the analysis (2.3%). A sample to hypothesized mean (of 0.5) t-test, on collapsed data from all blocks, was significant, $t(15) = 2.40$, $p = .03$, Cohen's d = .60, reflecting that participants more often chose the *high-EV* quadrant in one of their two clicks than chance. We performed a 1-way ANOVA on the choice data across the 12 blocks to determine if the learning effect changed over time. Mauchley's test showed a violation of the sphericity assumption, so we applied the Greenhouse-Geisser correction; results were significant, $F(4.12, 143) = 2.65$, $p = .042$, $\eta_p^2 = .17$. As we did previously, we used linear regression to calculate slopes in high-EV choices across the blocks for each participant, and the one-sample t-test was significant, $t(15) = 2.60$, $p = .020$, Cohen's d = .65, consistent with an increase in learning over time. Mean choice frequency across quadrant types and blocks are shown in Figure 9B (bar graph).

Second half: Visual search

RT. Trimming removed 1.1% of trials. Because we could not measure RTs from the trials in which participants did not find a target, we only analyzed the trials in which participants found a target, which was 49.1% (among those trials, “too slow” trials were 43.4%). This relatively small number of trials produced several missing data points and noisy data when we separated RTs separately by quadrant type and block (Figure 9B line graph). Therefore, we collapsed across blocks, and conducted a paired samples t-test, which showed no effect of quadrant type, $t(15) = 0.912$, $p = 0.376$, $d = 0.228$, $BF_{01} = 2.85$ (Figure 9C). Therefore, any evidence of participants prioritizing the high-EV quadrant during search in the first half of the experiment had vanished during the second half. This is notable, given that the second half was specifically designed to maximize any chance of such prioritization.

One potential reason for our failure to observe prioritization of the EV quadrant during search could have stemmed from the order of clicks in the choice task. That is, if participants routinely clicked the high-EV quadrant first, followed by a low-EV quadrant, then an attentional bias could remain in the last clicked quadrant (low-EV) at the start of the search. While our requirement that participants click the center point prior to the onset of the search was designed to prevent the persistence of such a bias, we further checked to see if any click order preference was present in the choice data. We found that, on trials in which the high-EV quadrant was clicked, it was no more likely to be clicked first or second, $t(15) = 1.35$, $p = .20$, $BF_{01} = 1.83$.

<u>Explicit Learning Assessment</u>

During the generation task, participants did not reliably vary in their selection of 20 points vs. 1 point across the *high-EV* (57.7% vs. 42.3%)and *low-EV* quadrants (53.4% vs. 46.6%), $F(1, 15) = 1.88$, $p = 0.199$, $\eta_p^2 = .11$.

A.
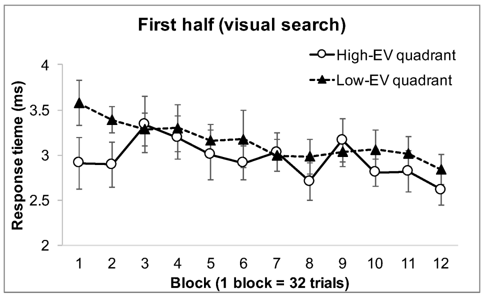

B.
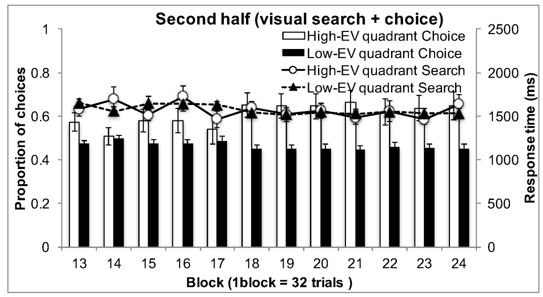

C.
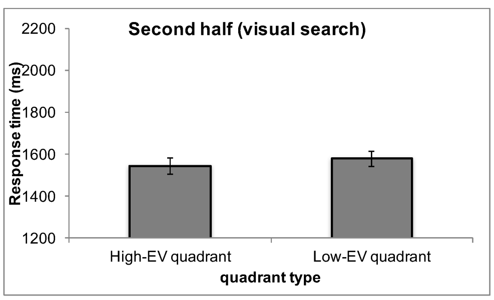

Figure 9. Results in Experiment 3. A. Search RT in the first half of experiment as a function of quadrant type, across blocks. B. The choice frequency (bar graph) in the second half of experiment as a function of quadrant type, across block and search RT (line graph) in the second half of experiment as a function of quadrant type, across blocks. C. Search RTs for the two quadrant types in the second half of the experiment, collapsed across block. Error bars show ±1 S.E. of the mean

Visual search vs. visual choice

We compare results from visual search and visual choice during the combined search and choice trials in the second half of the experiment, we calculated *learning efficiency units* for each task (see Experiment 1 for detail). A pairwise t-test revealed numerical but non-significant difference between two tasks, $t(15) = 1.513$, $p = .151$, $BF_{01} = 1.51$. Learning efficiency of visual choice and visual search is shown in Figure 10.

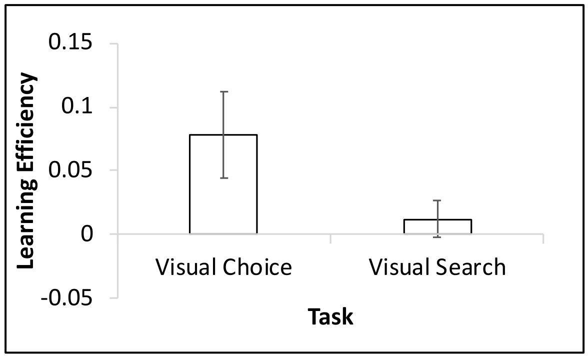

Figure 10. *Learning efficiency of visual choice and visual search in the second half of experiment. Error bars show ±1 S.E. of the mean.*

General Discussion

Across three experiments, we found a clear qualitative distinction between visual choice and visual search. Participants learned to prioritize high-value locations when their task was one of choice, but the same participants failed to prioritize high-value locations when their task was visual search (we found just one exception to this, in the first half of Experiment 3, although a pooled analysis across 5 experiments suggested the result was a Type I error). Even when we incorporated choice and search within the same trial, such that the two were jointly involved in representing the same target information – and both were incentivized such that behavior determined the reward outcome – we continued to observe significant prioritization during choice but not during search. That is, we demonstrated that participants learned the spatial value contingencies, actively represented this information, and displayed sufficient motivation to seek reward, but they nonetheless failed to express this knowledge during search. We thus find support for a *task-dependent expression failure* in the domain of visual search.

Broadly, these findings represent an example of the notion that some cognitive mechanisms are impenetrable to knowledge represented in different domains (e.g., Fodor, 1983). For instance, it has been argued that our low-level perceptual processing is uninfluenced by higher-level object knowledge, as highlighted by our inability to willfully alter how we perceive some visual illusions (e.g., the Müller-Lyer; see Firestone & Scholl, 2015; Pylyshyn, 1999). Other examples come from the domain of motor control, in which, knowledge of mirror tracing is clearly dissociated from the time-consuming process of skill acquisition (Milner, 1962).

While we present evidence favoring a default impenetrability of the visual search apparatus, we must acknowledge several caveats. First, from an experimental standpoint, we rely on producing learning in one task (visual choice) and measuring the transference to another (visual search). However, we cannot know if the choice task carried a stronger manipulation than the search task; moreover, equating the two tasks would be virtually impossible. That said, we do believe that the chosen search task is amenable to robust learning, based on a broad array of studies using this task and demonstrating strong learning effects (including some of our own, such as Won & Leber, 2016).

Second, some might question how much the spatial value learning is useful for the visual search task because, unlike in the choice task, learning that a location is more valuable than other locations does not help one to find the target any sooner. That is, because the target must be found, regardless of whether it appears in a high-value or low-value location, participants in our basic paradigm could not alter their earnings by first searching within the high-value quadrants. Although this is a valid point, consider classic animal learning studies in which the rats do prioritize high-value locations even when such behavior need not yield a better outcome (Crespi, 1942; Herrnstein, 1961; Spear & Pavlik, 1966). Furthermore, if spatial value learning would occur only when the information could facilitate search behavior (and earnings), then non-spatial features (e.g., color) that are associated with high reward should not prioritize attention in visual search. Yet, many studies including our previous study (Won & Leber, 2016) showed that high-reward associated colors prioritize attention in visual search even when the color is no longer rewarded (or even when it becomes penalized, Failing & Theeuwes, 2017; Le Pelley, Pearson, Griffiths, & Beesley, 2015). Regardless, to address these concerns, we conducted Experiment 3, where learning spatial value was actually useful and incentivized, but again, we found no significant search bias toward the high value location.

Third – and also related to the first point – our argument rests in large part on supporting the null hypothesis. Failing to reject the null hypothesis with classical statistics could occur due to limitations in statistical power, even when a true difference is present (type II error). One approach to increase confidence in the null hypothesis is to compute the Bayes Factor (BF), which generates a readily interpretable odds ratio of evidence for vs. against the null hypothesis (e.g., Rouder, Speckman, Sun, Morey, & Iverson, 2009). We computed these statistics on our test-phase visual search RTs for Experiments 1-3, using JASP 0.8.1.2 (JASP Team, 2017), with the default prior width of 0.707, and found that our data were more likely to be observed given that the null hypothesis were true vs. if the null hypothesis were false (Wetzels et al., 2011). Taken together, these results strengthen the interpretation that visual search performance in our experiments was insensitive to spatial value.

Previous concerns notwithstanding, we do not intend to argue that the visual search apparatus is impervious to all learning, an extreme argument that is contradicted by several robust findings. For one, the phenomenon of contextual cueing shows that individuals performing a visual search task learn when the search array is presented in spatial configurations that repeat multiple times. When these configurations are associated with specific target positions, participants exploit this information and begin to rapidly prioritize search to these expected target locations, facilitating both RT (Chun & Jiang, 1998; Chun, 2000; Jiang & Wagner, 2004; Gibson, Leber & Mehlman, 2015; Hout & Goldinger, 2010) and eye movements (Peterson & Kramer, 2001; Hout & Goldinger, 2012). Another well-documented phenomenon, described in this paper, is probability cueing, which shows that individuals rapidly learn to prioritize locations in the search display that contain more frequent targets (Geng & Berhmann, 2002; Jiang, Swallow, Rosenbaum & Herzig, 2013; Miller, 1988; Won & Leber, 2015).

Additionally, visual search is penetrable by volition; individuals can voluntarily choose to search some portions of the display over others, as in the classic case of endogenous cueing (Posner, Snyder & Davidson, 1984). Had we explicitly told our participants which quadrant to search first, we undoubtedly would have found that they could comply; indeed, when Jiang et al. (2015, Exp. 4) informed their participants about the spatial reward contingencies and asked the participants to prioritize the high-value locations, several of the participants were able to do exactly that. In the current experiments, we employed an incidental learning approach, in which we did not explicitly inform participants of the value manipulation. Clearly the visual search apparatus is not universally impenetrable to outside cognitive inputs (i.e., volition and certain forms of spatial learning). However, we have observed here that the visual search apparatus does not spontaneously act upon currently represented spatial value information.

This finding is parsimonious with our speculation that the search apparatus evolved without any environmental pressures to incorporate spatial value. That is, it is of great ecological significance for us to learn the value of individual object properties for which we might search. For instance, when searching for berries in a patch where red ones are desirable while yellow ones are poisonous, we could benefit greatly from a visual search apparatus that can boost the priority of red information, regardless of its location. Consistent with this notion, as mentioned above, there have been dozens of recent reports in which color-based or object-based value information drives visual search (e.g., Anderson et al., 2011; Della Libera & Chelazzi, 2006; Hickey et al., 2010; Kiss et al., 2009; Navalpakkam et al., 2010; Shomstein & Johnson, 2013), including in our own paradigm (Won & Leber, 2016). However, as we stated in the introduction, prioritizing search based only on spatial value is an ecological oddity. That is, we devised a task in which only one target of search was to be found. In the real world, when an individual has committed to finding a specific target, its objective value becomes irrelevant. Moreover, its value should rarely vary as a function of its ultimate location, as we argued when discussing the relative value of car keys that are hanging from the door vs. on a desk. Thus, perhaps it should not be surprising that the visual search apparatus does not adapt to an artificial scenario in which a target's value does vary depending on its spatial location. Of course, it is not impossible that an object's value changes as a function of location. For instance, ice cream on the table is more valuable than ice cream on the ground.[3] However, we believe that examples like this occur infrequently, compared to everyday searches for items such as keys, people, articles of clothing, etc. Such examples may not carry sufficient behavioral significance to drive evolutionary changes in the design of our search apparatus.

To conclude, this study takes a step toward resolving a puzzling pattern in the literature, in which dozens of papers have reported feature or object-based reward learning at the same time as a striking scarcity in reports of space-based reward learning. We have systematically

[3] We thank the anonymous reviewer for this example.

shown that individuals fail to express such learning even while simultaneously demonstrating it in other task domains. We propose that the visual search apparatus is not be designed to make use of this information source.

References

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences, USA.,108,* 10367–10371. doi: 10.1073/pnas.1104047108

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

Cain, M. S., Vul, E., Clark, K., & Mitroff, S. R. (2012). A Bayesian optimal foraging model of human visual search. Psychological Science, 23(9), 1047-1054. doi: 10.1177/0956797612440460

Chelazzi, L., Eštočinová, J., Calletti, R., Gerfo, E. L., Sani, I., Della Libera, C., & Santandrea, E. (2014). Altering Spatial Priority Maps via Reward-Based Learning. *The Journal of Neuroscience, 34*(25), 8594–8604. doi: 10.1523/JNEUROSCI.0277-14.2014

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences, 4,* 170–177. doi: 10.1016/S1364-6613(00)01476-5

Chun, M. M. & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology, 36*, 28–71. doi: 10.1006/cogp.1998.0681

Chun, M. M. & Jiang, Y. (2003). Implicit, long-term spatial contextual memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(2), 224–234. doi: 10.1037/0278-7393.29.2.224

Crespi, L.P. (1942). Quantitative variation in incentive and performance in the white rat. *The American Journal of Psychology, 55*, 467-517. doi: 10.2307/1417120

Della Libera, C. & Chelazzi, L. (2006). Visual selective attention and the effects of monetary reward. *Psychological Science, 17*, 222–227. doi: 10.1111/j.1467-9280.2006.01689.x

Failing, M. & Theeuwes, J. (2017). Don't let it distract you: how information about the availability of reward affects attentional selection. *Attention, Perception & Psychophysics, 79*(8), 2275-2298. doi: 10.3758/s13414-017-1376-8

Firestone, C., & Scholl, B. J. (2015). Can you experience 'top-down' effects on perception?: The case of race categories and perceived lightness. *Psychonomic Bulletin & Review*, *22*(3), 694-700. doi: 10.3758/s13423-014-0711-5

Fodor, J. A. (1983). *Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, Mass.: MIT Press

Geng, J. J. & Behrmann, M. (2002). Probability cuing of target location facilitates visual search implicitly in normal participants and patients with hemispatial neglect, *Psychological Science, 13*, 520–525. doi: 10.1111/1467-9280.00491

Gibson, B. M., Leber, A. B., & Mehlman, M. L. (2015). Spatial Context Learning in Pigeons (Columba livia). *Journal of Experimental Psychology: Animal Learning and Cognition*, *41*(4), 336-342. doi: 10.1037/xan0000068

Goldstein, H. & Spence, K. W. (1963). Performance in differential conditioning as a function of variation in magnitude of reward, *Journal of Experimental Psychology, 65* (1), 86-93. doi: 10.1037/h0043767

Halsey, L. G., Curran-Everett, D., Vowler, S. L., & Drummond, G. B. (2015). The fickle P value generates irreproducible results. *Nature methods, 12*(3), 179-185. doi:10.1038/nmeth.3288

Hastie, R., & Dawes, R. M. (Eds.). (2010). *Rational choice in an uncertain world: The psychology of judgment and decision making*. Sage.

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of The Experimental Analysis of Behavior, 4*(3), 267–272. doi: 10.1901/jeab.1961.4-267

Hickey, C., Chelazzi, L., & Theeuwes, J. (2010). Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience, 30,* 11096–11103. doi: 10.1523/JNEUROSCI.1026-10.2010

Horowitz, T. S. & Wolfe, J. M. (1998). Visual search has no memory. *Nature, 357,* 575-577. doi:10.1038/29068

Hout, M.C. & Goldinger, S.D. (2010). Learning in repeated visual search. *Attention, Perception & Psychophysics, 72*(5), 1267-1282. doi:10.3758/APP.72.5.1267

Hout, M.C. & Goldinger, S.D. (2012). Incidental learning speeds visual search by lowering response thresholds, not by improving efficiency: Evidence from eye movements. *Journal of Experimental Psychology: Human Perception and Performance, 38*(1): 90–112. doi: 10.1037/a0023894

JASP Team (2017). JASP (Version 0.8.1.2) [Computer Software].

Jiang, Y. V., Sha, L. Z., & Remington, R. W. (2015). Modulation of spatial attention by goals, statistical learning, and monetary reward. *Attention, Perception, & Psychophysics, 77*(7), 2189-2206. doi: 10.3758/s13414-015-0952-z

Jiang, Y. V., Swallow, K. M., Rosenbaum, G. M., & Herzig, C. (2013). Rapid acquisition but slow extinction of an attentional bias in space. *Journal of Experimental Psychology: Human Perception and Performance, 39*, 87–99. doi:10.1037/a0027611

Jiang, Y. V., Swallow, K. M., Won, B. -Y., Cistera, J. D., & Rosenbaum, G. M. (2015). Task specificity of attention training: The case of probability cuing. *Attention, Perception, & Psychophysics, 77*(1), 50-66. doi: 10.3758/s13414-014-0747-7

Jiang, Y. & Wagner, L. C. (2004). What is learned in spatial contextual cueing: Configuration or individual locations? *Perception & Psychophysics,66*(3), 454-463. doi: 10.3758/BF03194893

Kamin, L. J. (1969). *Selective association and conditioning*, N.J. Mackintosh, W.K. Honig (Eds.), Fundamental issues in instrumental learning, Dalhousie University Press, Halifax (1969), pp. 42–64

Kiss, M., Driver, J., & Eimer, M. (2009). Reward priority of visual target singletons modulates event-related potential signatures of attentional selection, *Psychological Science, 20,* 245–251. doi: 10.1111/j.1467-9280.2009.02281.x

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. http://dx.doi.org/10.1037/xge0000037

Miller, J. (1988). Components of the location probability effect in visual search tasks. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 453–471. doi: 10.1037/0096-1523.14.3.453

Milner, B. (1962). *Les troubles de la me´ moire accompagnant les lésions hippocampiques bilatérales*. In Physiologie de l'Hippocampe, Colloques Internationaux No. 107 (Paris, C.N.R.S.), pp. 257–272. [English translation (1965). In Cognitive Processes and the Brain, P.M. Milner and S. Glickman, eds. (Princeton, NJ: Van Nostrand), pp. 97–111.]

Morey, R. D., Rouder, J. N., & Jamil, T. (2014). Bayes factor: Computation of Bayes factors for common designs. R package version 0.9.

Navalpakkam, V., Koch, C., Rangel, A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences, USA, 107*, 5232–5237. doi: 10.1073/pnas.0911972107

Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex* (G. V. Anrep, Trans.). London: Oxford Univ. Press.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437–442. doi:10.1163/156856897X00366

Peterson, M. S., Kramer, A. F., Wang, R. F., Irwin, D. E., & McCarley, J. S. (2001). Visual search has memory. *Psychological Science, 12*(4), 287-292. doi: 10.1111/1467-9280.00353

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology*, *109*(2):160-174. doi: 10.1037/0096-3445.109.2.160

Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences, 22*(3), 341-65; discussion 366-423. doi: 10.1017/S0140525X99002022

Richter, C. P. (1922). A behavioristic study of the activity of the rat. *Comparative Psychology Monographs*, *1*. 1-55.

Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology, 56*, 356–374. http://dx.doi.org/10.1016/j.jmp.2012.08.001

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review, 16*, 225–237. http://dx.doi.org/10 .3758/PBR.16.2.225

Shomstein, S., & Johnson, J. (2013). Shaping attention with reward effects of reward on space- and object-based selection. *Psychological Science*, *24*(12), 2369-2378. doi: 10.1177/0956797613490743

Spear, N. E. (1973). Retrieval of memory in animals. *Psychological Review, 80,* 163-194. doi: DOI: 10.1037/h0034326

Spear, N. E., & Pavlik, W. B. (1966). Percentage of reinforcement and reward magnitude effects in a T maze: Between and within subjects. *Journal of Experimental Psychology, 71*(4), 521-528. doi: 10.1037/h0023014

Tolman, E. C. (1932). *Purposive behavior in animals and men*. The Century Co, New York

Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185(4157), 1124-1131. doi: 10.1126/science.185.4157.1124

von Neumann, J. & Morgenstern, O. (1953). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

Wasserman, E. A. (1981). Comparative psychology returns: A review of Hulse, Fowler, and Honig's "Cognitive processes in animal behavior." *Journal of the Experimental Analysis of Behavior, 35*, 243-257. doi: 10.1901/jeab.1981.35-243

Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., & Wagenmakers, E-J. (2011). Statistical evidence in experimental psychology: An empirical comparison using 855 t tests. *Perspectives on Psychological Science, 6*, 291-298. doi: 10.1177/1745691611406923

Wolfe, J. M. (2003). Moving towards solutions to some enduring controversies in visual search. *Trends in Cognitive Sciences, 7*(2), 70-76. doi: 10.1016/S1364-6613(02)00024-4

Wolfe, J. M. (2013). When is it time to move to the next rasberry bush? Foraging rules in human visual search. Journal of Vision, 13(3), 1-17. doi: 10.1167/13.3.10

Wolfe, J. M., Alvarez, G. A., & Horowitz, T. S. (2000) Attention is fast but volition is slow. *Nature, 406*, 691. doi:10.1038/35021132

Won, B. -Y. & Leber, A. B. (2016). How do magnitude and frequency of monetary reward guide visual search? *Attention, Perception & Psychophysics, 78*(5), 1221-1231. doi: 10.3758/s13414-016-1154-z