Distinct selection mechanisms for when predictions and rewards guide visual selective attention

K.G. Garner & J.E. Raymond

School of Psychology, University of Birmingham

Correspondence to:

K.G. Garner

getkellygarner@gmail.com

Hills Building, 2.16

University of Birmingham

Edgbaston, Birmingham

B15 2TT

Abstract

How does the brain combine predictions and rewards to guide visual selective attention? Recent findings in humans have implied an additive influence, suggesting distinct sources of bias that act on visual selection (*additive-bias hypothesis*). Alternately, neuroeconomic evidence suggests that the brain applies an *expected-value* operation, where the likelihood that a spatial location will become task relevant (*spatial-certainty*) is multiplied by the reward-value of that location (*spatial-value*). We sought to arbitrate between these two alternatives, and furthermore, to test whether these conflicting accounts have arisen because both operations are available to the visual-selection system, and their manifestation is context dependent (*mixed-operations hypothesis*). Over two experiments we tested whether the influence of predictions and reward remained additive, even under conditions that stressed the optimality of an additive operation. Participants completed an endogenous cueing task. Spatial-certainty varied across blocks, and targets appeared in coloured circles that were associated with high or low reward values. In Experiment 1 we tested multiple conditions of spatial-certainty, and more densely sampled high certainty-levels, under the assumption that this would free more resources for the influence of spatial-value, should the two share underlying resources. In Experiment 2, potential reward value was manipulated so that an expected-value operation was objectively more profitable. Across both experiments, the additive-bias hypothesis clearly outperformed the expected-value and mixed-operations hypotheses. These findings refute theories that expected-value computations are the singular mechanism underlying the deployment of endogenous spatial attention. Instead, and even when suboptimal, spatial-certainties and spatial-values are parcellated by the selective attention system.

The brain is continually bombarded with more sensory information than it can represent at any given moment, it must therefore prioritise the most relevant sensory inputs to the detriment of others. We have long understood that this is an important aspect of brain function; as the philosopher Malebranche noted in 1674, "[the mind] applies itself infinitely more to those things that affect it, modify it, and that penetrate it (p. 412)[1]. Despite this, our understanding of what modifies prioritisation of visual signals, i.e. *visual-selection*, is still far from complete[2]. Most computational and neurobiological theories of visual-selection[2–4] assume an endogenous mechanism that flexibly selects information in accordance with the current goals. Outside the assumed scope of such an endogenous mechanism, visual-selection has also been shown to be influenced by learned associations between sensory-cues and consequent outcomes, such as the reliability with which a cue predicts an outcome[5], and the reward-value of an outcome[6,7], even when these associations are at odds with current task objectives[8,9]. Although we know that these elements of visual-selection are at least partly dissociable[10], far less is known about how the brain combines these influences to prioritise information processing. In this study we seek to understand how learned-associations pertaining to the predictiveness and value of sensory cues can combine with task goals to drive visual-selection.

Endogenous visual-selection

In order to assess endogenous visual-selection, we used a spatial-orienting task, where a symbolic cue provides advanced information regarding the likely location of an upcoming target[4,11]. These cues elicit more rapid and accurate responses to targets appearing at the cued location, at the expense of the uncued location (hereafter referred to as the *cueing-effect*); thereby providing a measure of the consequences of prioritising spatial locations. A central concept to theories accounting for cueing-effects is that attention amplifies the output of neurons that are sensitive to the highlighted feature of interest[2,12,13], thereby increasing the influence those neurons have on subsequent visual-selection behaviours. Features of interest are assumed to be defined by competition within a saliency[3,14,15] or priority[16,17] map, where objects of interest are represented by activity that is proportional to their behavioural and/or physical relevance. In support of these models, neurophysiological activity in both

monkey parietal cortex[17,18] and frontal eye fields[14] appear to encode stimulus features more strongly when they are behaviourally relevant or physically salient (although alternate, biologically plausible models of visual-selection do not assume a saliency map[19,20]).

Visual-selection and spatial cue-certainty

It is commonly assumed that we tap the voluntary deployment of attention when we measure spatial cueing-effects[4,21]. We assume that a goal is adopted (now referred to as task-set), and visual selection is deployed in accordance with that task-set. Evidence shows that cueing-effects are also influenced by factors that do not explicitly belong to the task-set. For example, cueing-effects scale with cue-reliability[22–25], even though participants are often unable to verbally report changes in cue-target contingencies[22]. Therefore, when visual-selection is applied to achieve the task-set, it is sensitive to implicit factors that may facilitate or hinder completion of that task. Neurophysiological evidence suggests that the substrates of this effect overlap with those that underpin the voluntary prioritisation of spatial locations. For example, recordings in monkey parietal cortex show that neural activity not only scales with physical salience and behavioural relevance, but also with cue-reliability[26,27]. This has motivated the idea that the likelihood that a cue predicts an upcoming spatial location (spatial-certainty) is encoded into the priority map that determines the targets of visual-selection[26].

Recently, Prinzmetal et al.[28] showed that cueing-effects follow the Hick-Hyman law of decision-time [29,30]; i.e. the size of the cueing-effect scales linearly with the spatial-certainty gained (in bits) by the cue. This is important because it provides an easily computable, parametric manipulation of *spatial-certainty*. The authors used Shannon's[31] measure of entropy (H) to calculate the spatial-certainty gained by the cue. *H* measures the amount of uncertainty in a probability distribution and is at maximum when the cue is unpredictable with regard to the target location. Therefore spatial-certainty gained by an informative cue can be calculated as:

$$Spatial\ certainty = H_{no\ information} - H_{cue} \qquad (1)$$

when H is defined in the standard manner:

$$H = -\sum_i p_i \, log_2 p_i \qquad (2)$$

and p_i is the probability that the target appears at location i, given the cue. For example, with 2 locations, and a cue that is .8/.2 valid/invalid:

$$H_{cue} = -(.8 \, log_2 .8) - (.2 \, log_2 .2) \approx .72$$

As $H_{no\ information}$ is 1, then the information gained by the cue is 1 - .72 $\approx$ .28 bits.

Although we know that cueing-effects scale linearly with spatial-certainty[28], we do not yet know how the strength of this relationship is influenced by other learned-associations. It could be that the magnitude of this relationship remains constant, or interacts with other learned-associations that complement the task-set. One other such association that is also assumed to be encoded into the visual-selection priority map is that between sensory cues and their consequent reward-values.

Visual-selection and spatial-value

Recordings from parietal cortex in monkeys show that neural activity not only scales with task-set and cue-certainty, but also with the reward value offered by cues at task-relevant locations (*spatial-value*)[26,32–34]. The influence of spatial-value is also evident in human behavioural responses. Targets appearing at locations that have been associated with high values are responded to faster and more frequently than those associated with low values[7]. Furthermore, high spatial-values can impede selection of competing, non-rewarded locations[8,9]. Taken together, these lines of evidence show that spatial-value also contributes to the computational operations that guide visual selection.

Recordings of monkey V1 neurons during a task that combines overt spatial-cues with value-cues suggest that the same neural mechanism is sensitive to both these information sources. Stanisor et al[35] recorded activity in primary visual cortex (V1) as monkeys gained rewards by making saccades to coloured discs at different spatial locations. Each colour signalled the availability of a different

reward-value, and each colour was equally likely to become a target prior to spatial-cue onset. A line drawn from fixation to one of the value-cues served as the spatial-cue. The monkeys made saccades towards the cued target either 50 ms or 400 ms after spatial-cue onset. The authors observed that when the saccade was made shortly after cue-onset (50 ms), V1 activity correlated with the relative reward values associated with the two discs. When the onset of the saccade was delayed, V1 activity evolved to correlate with the spatial location of the cued-target, regardless of its associated reward-value. Clusters of V1 neurons showing sensitivity to the relative reward-values were strongly correlated with those sensitive to the direction of the spatial cues, suggesting overlap in the neurons sensitive to both signals. The authors reasoned that as the spatial-cue was 100 % valid, this acted to change the relative expected-values - i.e. once one target was cued, it offered all the reward-value available on that trial, whereas the non-cued target offered none, regardless of its reward-value prior to spatial cue-onset. This suggests that spatial-cues act to change the relative values of potential target locations, and that spatial cues-and value-cues drive visual-selection via a single selection signal.

Combining learned associations to influence visual-selection

It is clear that both spatial-value and spatial-certainty influence visual-selection, and that there is strong overlap in the neural representation of these two information sources. However, it remains unclear how these factors are combined to facilitate visual-selection, particularly when they co-operate with, or are in direct competition to, the current task-set. For example, high spatial-certainty should combine with high spatial-value to imbue stronger prioritisation of a task-relevant location than a comparable high-certainty/low-value combination. However, it remains contested what computations underlie the mapping between these inputs and the behavioural output.

Stankevich and Geng[36] showed evidence for an additive relationship between spatial-value and certainty that a target will appear in a specific spatial-location. However, spatial certainty was manipulated in the absence of overt spatial cues. Participants associated two circles of differing colours with differing reward values. A dot then appeared in each circle; one was positioned above or below the meridian (the target), the second appeared at the meridian. Participants were required to detect whether

the target was above or below the meridian. Over time and without explicit instruction, the likelihood that the target would appear in one location (e.g. left) relative to the other location increased. Although participants showed performance benefits from both value and location-likelihood associations, these benefits were additive. According to additive-factors logic[37], this suggests that spatial-value and certainty act distinctly to influence visual-selection mechanisms. The authors thus concluded that spatial-value and spatial-certainty associations represent independent sources of influence. However, this manipulation of spatial-certainty acts distinctly to the influence of symbolic spatial cues[38], therefore it remains to be identified whether a comparable additive influence between spatial-certainty and spatial-value is observed when spatial-certainty is conveyed via a symbolic cue.

In contrast, because signals corresponding to spatial-certainty and spatial-value are both observed in monkey parietal[26,32,33] and visual[35] cortices, it has been claimed that the brain makes a priority computation based on relative expected-value; where spatial-value (r) is multiplied by the likelihood that it will be obtained (p), and is then normalised by all expected-values in the visual-field:

$$Relative\ expected\ value = \frac{p_i * r_i}{\sum_{i=1}^{n} p_i * r_i} \qquad (3)$$

In support of this idea, saccadic response times in a visual detection task have been shown to correlate with the relative expected-value between potential visual-selection targets[39]. In this study, participants were required to saccade to single red-dots that appeared within the left or right hemifield of a blank display. Over blocks, the authors varied the likelihood of the target-location, and the reward-value associated with the target, given its appearance at that location. The time taken to initiate a saccade to the target correlated more highly with the relative expected-value of the potential target locations, than with either the relative likelihood of the locations, or the relative reward values. However, as both factors have been shown to influence performance[6,28], it could be that any model combining both factors would outperform a model based on either factor alone. Therefore, we cannot conclude that an additive-bias model would not account better for this data. It is therefore important to directly compare predictions

made by the additive-bias and expected-value models. Providing this more appropriate test is a central aim of the current work.

We also considered a third alternative. It is possible that it is optimal for the visual system to maintain independence between spatial-value and spatial-certainty when these values are orthogonal, as is the case when good experimental design ensures that variables are unconfounded. However, there must be limits to the additive nature of the influence, after which the independent influence may break down to reveal a non-additive relationship, given that the same mechanism is influenced by both information sources. In this case, the influence should be additive until the system is sufficiently stressed, after which an expected-value operation may be revealed (*mixed-operations hypothesis*).

Aims and Hypothesis

Over two experiments, we sought to understand whether the additive-bias, expected-value, or mixed-operations hypothesis provided a better account of visual-selection, under conditions that could be expected to stress the optimality of additive operations.

In Experiment 1, we tested whether an additive-bias would hold under conditions where spatial-certainty approaches maximum. If the influence of spatial-certainty and spatial-value is always serial, then additivity should be maintained across all levels of spatial-certainty. In contrast, if the two influences can be integrated, we expect that this relationship should either be non-additive across all levels of spatial-certainty (*expected-value hypothesis*) or become non-additive as spatial-certainty nears maximum (i.e. as the cue becomes close to 100 % valid, *mixed-operations hypothesis*). This is because as the computation of spatial-certainty becomes trivial, the mechanism using both information sources should have more availability to implement the influence of spatial-value. Therefore in Experiment 1 we pitted spatial-values against cues with varying spatial-certainties, more densely sampling probabilities between .8 and 1. Against the expected-value and mixed-operations hypotheses, we find that the the influence of spatial-certainty and spatial-value remains additive across all tested levels of spatial-certainty.

In Experiment 2 we tested whether the reward structure of the task can influence how spatial-certainty and spatial-value interact. It has been shown that the visual system learns to sample sensory information to optimise reward accrual[40–42]. We reasoned that if an expected-value operation is a possible computation underlying visual-selection, then reward conditions that favour an expected-value operation should yield a non-additive influence of spatial-value and spatial-certainty on performance. As the expected-value operation involves a multiplicative weighting of value by certainty, RTs to high spatial-values should decrease, relative to responses produced by an additive-bias operation, and RTs to low spatial-values should increase, relative to responses produced by an additive-bias operation. Therefore conditions in which faster responses can accrue higher rewards (that also outweigh the losses from slower responses to low spatial-values) would render an additive-bias operation less profitable than an expected-value one (see Figure 1d). We created a reward-condition where an exponential decay function was applied to spatial-values upon target onset, to provide the reward conditions required to perform this test. Again, and against the expected-value and mixed-operation hypotheses, an additive influence was found between spatial-certainty and spatial-value. Collectively the results favour the additive-bias hypothesis, suggesting that spatial-certainty and spatial-value act serially to influence visual-selection.

Results

Experiment 1

We adapted a spatial-orienting task[4] to test whether the influence of spatial-certainty and -value remains additive under conditions that approach maximal certainty (see Figure 1). Participants observed two small vertical lines presented in the centre of the screen. Two spatial-value cues of differing colour were placed on the horizontal meridian, to the left and right of fixation. After a short period, one of the lines would turn darker. This served as the spatial cue. Subsequently a letter would appear within each value cue. One would contain an ‘H’ or a ‘Z’, and the other a ‘K’ or a ‘Z’. Participants were asked to report whether an ‘H’ or a ‘Z’ had been presented, they were informed to ignore the ‘K’ or the ‘Z’.

Spatial-certainty of the cue was manipulated across blocks, and the location of the spatial-value cues was counterbalanced across trials. Participants were instructed that the value cues signified how many points were available should the target appear in that location (50 vs 1 point), and that the darkened line was a clue for where the target could appear. They were requested to keep their eyes at fixation, and to respond as accurately and as quickly as possible to the target. Participants were also informed that their points would be exchanged for cash at the end of the session (1000 = £1).

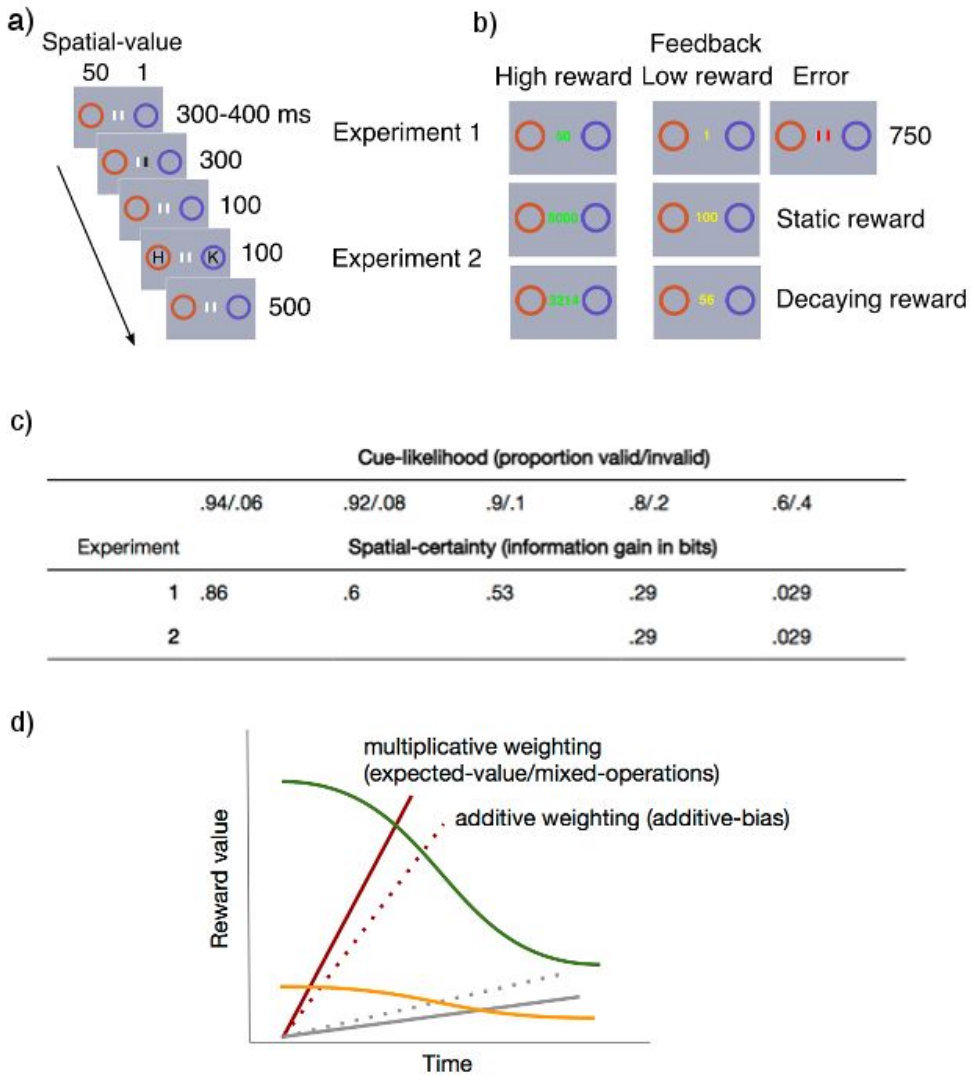

Figure 1: Study method. Task procedure and feedback conditions for Experiment 1 and Experiment 2. a) Trial structure: participants monitored the two value cues, and were informed that a central spatial cue would indicate in which of the two locations the target would probably appear. Participants were told to report whether the letter 'H' or

‘N’ appeared on each trial (a distractor, ‘K’ or ‘Z’ appeared in the other value cue). b) Reward feedback structure: At the end of each trial participants were informed as to whether they had gained a high reward, a low reward, or had made an error. In Experiment 2, reward feedback was either static (same reward value throughout the response period) or decaying (reward value began decreasing upon target onset). c) spatial-certainty was parametrically manipulated by increasing the information gained (in bits) from the cue. d) Logic of the decaying reward condition in Experiment 2. As both expected-value and mixed-operations computations (solid red and grey lines) involve a multiplicative weighting of spatial-certainty and spatial-value, responses towards high spatial-values (red lines) should be faster than an additive-bias operation (red dotted lines). Furthermore, responses to low spatial-values should be slower for expected-value/mixed-operations computations (grey solid line) than an additive-bias operation (grey dotted line). Applying an exponential decay function to the spatial-value (green line = high spatial-value, yellow line = low spatial-value) at target onset means that the extra rewards accrued by being faster towards high spatial-values would outweigh the losses accrued from being slower towards low spatial-values. Therefore any operation that favours this response pattern would accrue greater total rewards than an additive-bias operation.

The key aim of each analysis was to determine whether a model that included an interaction term between spatial-certainty and spatial-value was more probable, given the data, than one that only included main effects. We therefore sought to quantify the evidence in favour of one model relative to another. To achieve this, we used a Bayesian approach. First, we fit all possible linear mixed models, and computed log inverse Bayes Factors (BFs), which quantifies evidence for each possible linear mixed effects model against the null model (intercept plus random effects of participant). We then identified the 6 best performing models. We report the BF of the winning model relative to the null model, and the BF ratios between the best model and the next 5 best models, to quantify the evidence in favour for the winning model.

RT

In support of the additive-bias hypothesis, the RT data show evidence in favour of main effects for spatial-value and spatial-certainty across the increasing levels of certainty. Against the expected-value and mixed-operations hypotheses, we did not observe evidence in favour of an interaction between spatial-certainty and spatial-value.

The RT data (dots) and winning model fit (lines) are presented in Figure 2a. Against the expected-value and mixed-operations hypotheses, the preferred model included only main effects of each factor, and a spatial-certainty x cue-validity interaction term (BF = 1.74E+58, ± .87 %, see Figure 2B). The main effect of value demonstrates that RTs were slowed by approximately 50 ms ± .3 (SE) for low spatial-values, relative to high spatial-values. The influence of cue-validity increased with

cue-certainty; the difference between valid and invalid trials increased by approximately 90 ms ± 29 (SE) across the levels of certainty. Importantly, there was positive evidence that this model was preferred over the next best model (BF = 3.8, ± 1.45%), which was identical to the winning model except that it also included the spatial-certainty x spatial-value interaction term. Therefore the evidence favours a model that does not include an interaction between spatial-value and spatial-certainty.

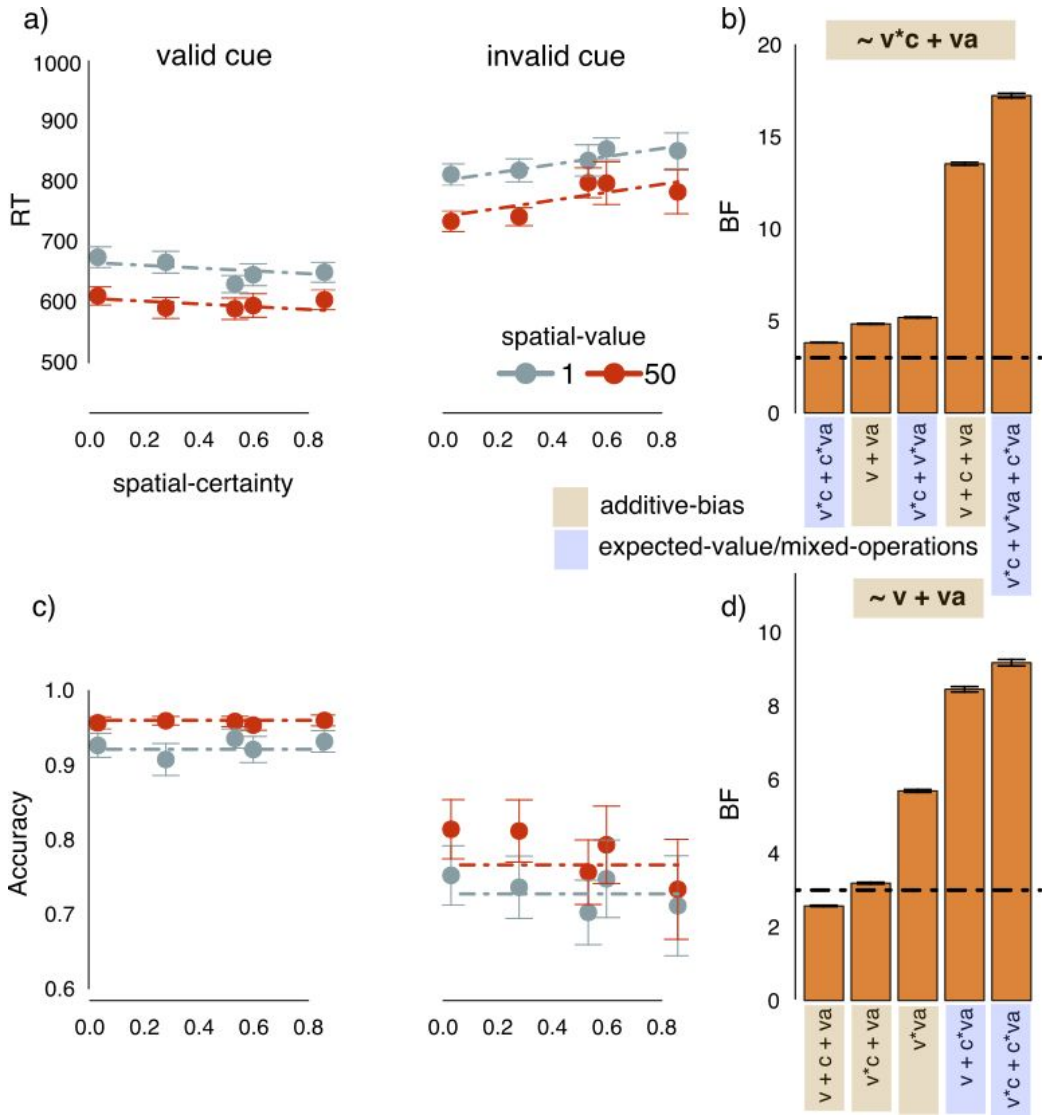

Figure 2: Overall results from Experiment 1. a) Observed RTs in ms (points) and 95 % within-subject confidence intervals calculated using the Cousineau-Morey method[43] and fit of the winning model (dashed lines). The winning model included a validity x spatial-certainty interaction and a main effect of spatial-value. b) BFs for the winning RT model (title) relative to the 5 next best models (see x-axis). Crucially there was positive evidence (boundary indicated by black dotted line) in favour of the winning model relative to models that allowed spatial-value to interact with the spatial-certainty factor. Models that allow an interaction between spatial-value and spatial-certainty (or cue-validity), and thus would provide support for the expected-value or mixed operations hypotheses, are shaded in blue. Models that would provide support for the additive-bias model are shaded in cream. c) Accuracy data showed a main effect of spatial-value and a main effect of cue-validity. Crucially there was positive evidence in favour of this

model relative to models that allowed spatial-value to interact with either cue-validity or spatial-certainty (d). RT = response-time, BF = Bayes Factor, v = cue-validity, c = spatial-certainty, va = spatial-value.

Accuracy

Accuracy data did not suggest that the results were due to a speed accuracy trade-off (see Figure 2c). The preferred model included only main effects of cue-validity and spatial-value (BF = 8.83E+47, ± .56 %, Figure 2d). For the main effect of cue-validity, accuracy on invalid trials decreased by approximately .16 ± .02 (SE), relative to valid trials. For the main effect of spatial-value, accuracy was decreased by around .05 ± .02 (SE) for low spatial-values, relative to high spatial-values. The evidence for this model against the next best model, which also included a main effect of spatial-certainty, was weak (BF = 2.57 ± .68 %). The evidence in favour of the preferred model was more positive relative to the third best model, which also included a cue-validity x spatial-certainty interaction (BF = 3.19 ± .78 %). This suggests that spatial-certainty may have influenced the accuracy measure, but below the extent that we were able to reliably detect. Interestingly, there is strong evidence that the model that includes main effects of cue-validity and spatial-value is preferred to any model that allows these factors to interact (see Figure 2d). Therefore, although we did not reliably detect an influence of spatial-certainty on the accuracy data, these results corroborate the notion that spatial-value is a distinct source of influence on the visual-selection system.

Experiment 1 shows that models posing an additive influence of spatial-certainty and spatial-value outperform those allowing an interaction between the two, across all levels of spatial-certainty. This goes against the expected-value hypothesis and suggests that visual-selection mechanisms maintain a distinction between spatial-certainty and spatial-value, even when approaching the limits of certainty. However, the second best model for RTs did include an interaction between spatial-certainty and spatial-value, suggesting that this interaction is not entirely implausible. We reasoned that it is still possible that an additive influence may not hold if another form of appropriate pressure is applied to visual-selection. A reward structure that favours an expected-value operation may be sufficient to modulate the additive influence of spatial-certainty and spatial-value. The aim of

Experiment 2 was to provide this test. In the decay reward-condition, the spatial-values began to decrease at target onset, thereby faster responses would accrue greater reward gains. In this condition, performance based on a multiplicative influence of spatial-value and spatial-certainty would fare better than performance motivated by an additive operation, particularly if the gains outweigh the losses that an expected-value operation would yield for low-value locations (relative to an additive operation, see Figure 1d). This is in contrast to the static reward condition where the total amount of reward received is independent of the operation used to produce visual-selection. If the expected-value or mixed-operations hypotheses are correct, then we would expect to see a non-additive influence in the decay reward-condition, in contrast to an additive influence in the static reward-condition.

Experiment 2

Results

RT

The RT data show that the influence of spatial-value is additive to cue-validity, even under conditions where it is suboptimal for reward accrual. Interestingly evidence no longer favoured an influence of spatial-certainty, suggesting that these effects may be more subtle at the range of certainties tested in the current experiment. However, we did find an influence of spatial-certainty on the accuracy data that is reported below.

The RT data are presented in Figure 3a. If the additive relationship between spatial-certainty and spatial-value is modulated by reward potential, then we would expect to find an interaction between these two factors in the decay reward-condition, that is larger than anything observed for the static reward-condition. First we identified the most likely model given the data. The winning model included main effects of cue-validity, spatial-value, and reward-condition (BF = 2.18E+67 ± .69 %, relative to null model). There was good evidence that this was the best model for the data, as it was positively preferred to the next best model, which included an additional spatial-value x cue-validity interaction term (BF = 4.65 ± 2.43 %). As spatial-certainty was found to interact with cue-validity in Experiment 1, we tested the

evidence for the winning model against one that also included a cue-validity x spatial-certainty interaction term. Again, there was positive evidence that the winning model provided a better fit to the data (BF = 8.97 ± 1.79 %). Therefore, the RT data reflect additive influences of cue-validity (RTs were on average 33 ms ± 12 (SE) slower on invalid trials relative to valid trials), spatial-value (RTs were approximately 21 ms ± 17 (SE) slower to low cue-values relative to high cue-values), and reward-condition (RTs were approximately 54 ms ± 17 (SE) faster in the decay reward-condition than the static reward-condition), without a detectable influence of spatial-certainty. Collectively, the results show that even when an additive operation is disadvantageous, the additive-bias model is still a better account of the data. However, given we did not observe an influence of spatial-certainty on the RT data, we can only say that spatial-value remains additive to the binary influence of spatial cues (i.e. valid/invalid) under conditions of decaying reward-value.

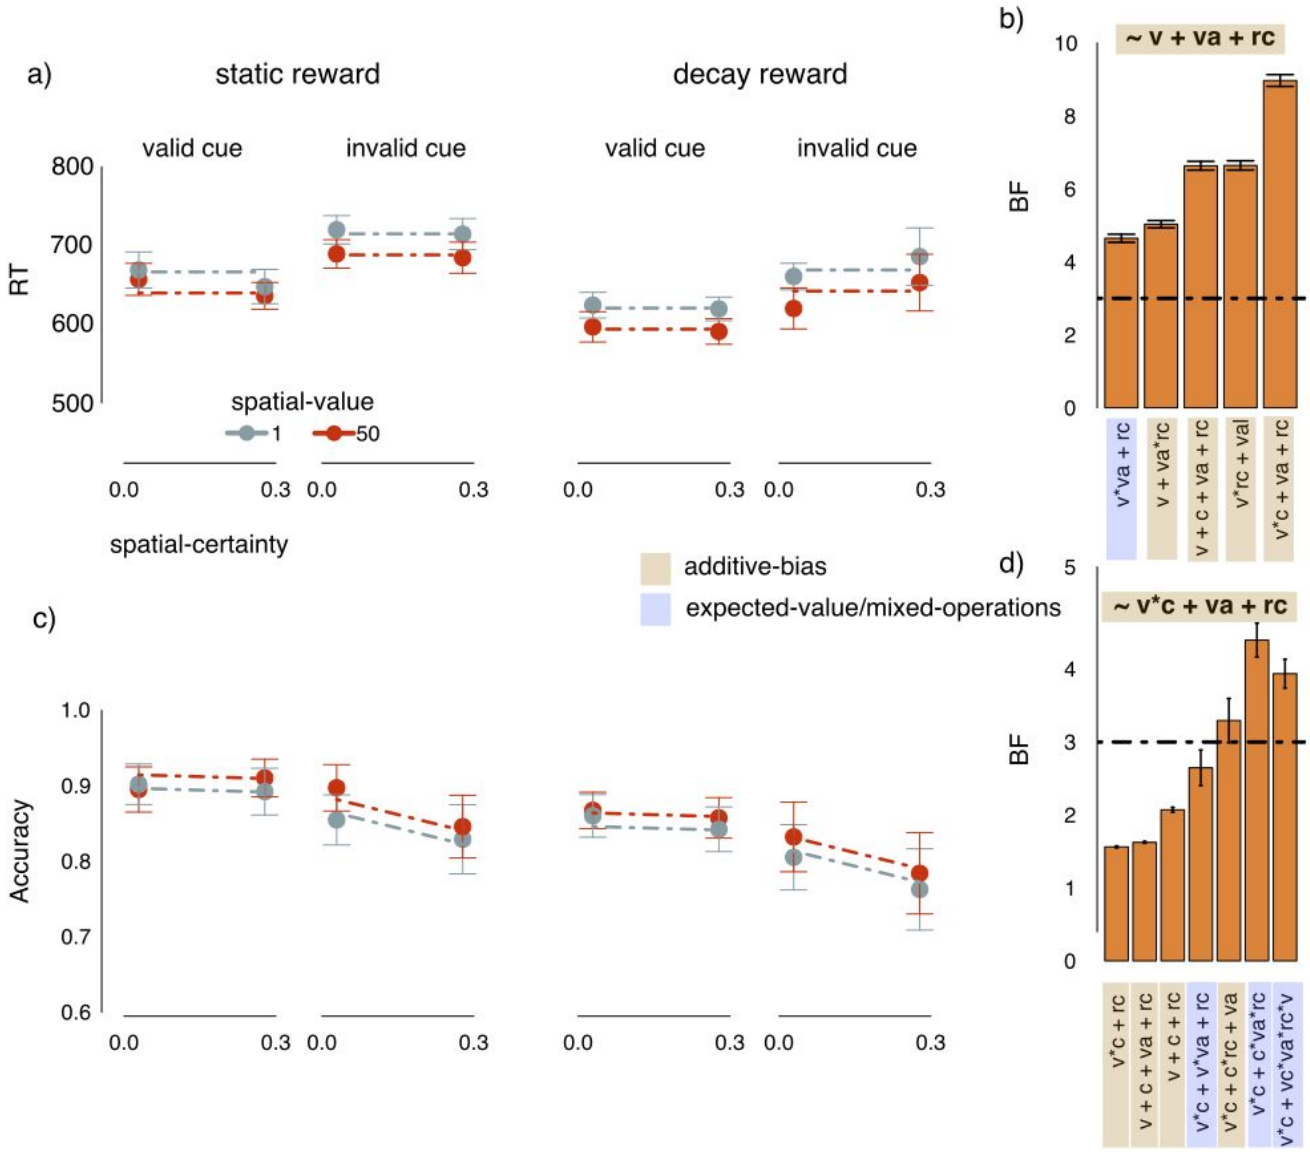

Figure 3: Overall results for Experiment 2. a) Observed RTs in ms (points) and 95 % within-subject confidence intervals calculated using the Cousineau-Morey method[43] and fit of the winning model (dashed lines) for the static and decaying reward-conditions. The winning model included main effects of cue-validity, spatial-value and reward-condition. Crucially, there was positive evidence in favour of this model relative to models allowing an interaction between spatial-value and the other experimental factors (b). c) Accuracy data showed an influence of spatial-certainty, cue-validity, spatial-value and reward-condition. Although the winning model was less clear cut, there was positive evidence in favour of the top performing model relative to a model that allowed an interaction between spatial-value and spatial-certainty (d). RT = response-time, BF = Bayes Factor, v = cue-validity, c = spatial-certainty, va = spatial-value.

Accuracy

Interestingly, and in contrast to the RT data, the accuracy data did show an influence of spatial-certainty. Importantly for the central question, the models that best accounted for the accuracy data did not show a spatial-certainty x spatial-value interaction (see Figure 3c/d).

Selecting which model *per se* best accounts for the accuracy data was less simple, as evidence for a single winning model was not conclusive. Comparable to the RT data in Experiment 1, the best model contained a cue-validity*spatial-certainty interaction, and main effects of reward-condition, cue-validity, spatial-value, and spatial-certainty. Differences in accuracy performance between valid and invalid trials grew larger as spatial-certainty increased (approximately .15 ± .07 (SE)). Furthermore, accuracy performance was slightly higher for high relative to low spatial-values (approximately .0003, ± .0001 (SE)). Accuracy was also higher for the static relative to the decay reward condition (.05, ± .008 (SE)). However, evidence in favour of this model was weak, relative to the next best model, which dropped the main effect of spatial-value (BF = 1.56, ± 0.75 %). Given this, and to address the key theoretical question, we directly tested whether spatial-value interacted with spatial-certainty by comparing the winning model to those that additionally included either a spatial-certainty x spatial-value x reward-condition interaction or a cue-validity x spatial-certainty x spatial-value x reward-condition interaction. Against the expected-value and mixed-operations hypotheses, evidence for the winning model relative to these two was positive (BFs = 4.62 ± 0.93 %, BFs = 3.9 ± 0.93 %). Collectively, the data support the additive-bias model and imply that spatial-value and spatial-certainty act distinctly to bias visual-selection.

General Discussion

Over two experiments we tested whether the additive-bias hypothesis would outperform the expected-value and mixed-operations accounts, even under conditions expected to stress the optimality of additivity. In Experiment 1, we assumed that if spatial-certainty and spatial-value share an underlying resource, conditions where spatial-certainty is trivial to compute (i.e. very high certainty) may free up resources available for the influence of spatial-value, thereby motivating a non-additive influence on performance. We pitted spatial-certainty and spatial-value against each other in a spatial-orienting task,

where endogenous cues signalled the likely location of upcoming letter targets. An additive influence of spatial-certainty and spatial-value was observed, even under conditions of very high certainty. Spatial-certainty increased the size of the cueing-effect (i.e. the difference between invalid and valid-cue trials), whereas spatial-value had a comparable influence on both valid and invalid trial types.

In Experiment 2, we reasoned that if an expected-value operation is available to the mechanisms underlying visual-selection, then a reward structure that favours a multiplicative weighting of spatial-certainty and spatial-value may be sufficient to reveal it. We applied a decay function to spatial-values at target onset, so that the gains made by faster RTs on high-value trials would outweigh the losses made by slower RTs on low-value trials. Although the influence of spatial-certainty was manifest differently to Experiment 1; i.e. by influencing accuracy, rather than by increasing the size of RT cueing-effects, we observed that the influence of spatial-value remained additive to spatial-certainty and to the other experimental factors. Again, the findings support the additive-bias hypothesis.

Overall the findings imply that the mechanism mediating the influence of spatial-value and spatial-certainty on visual-selection maintains a separation between the two information-sources. The finding of additive factors implies that at some point, the visual system engages in a serial readout of the two information sources[37]. This is potentially surprising, given that there is much overlap in the neural activity that correlates with value and certainty information during visual decision-making[26,32,33,35], and that relative-expected value computations assume a multiplicative relationship between the likelihood of an outcome and the value of that outcome. However, the current findings are in line with previous observations that monkeys are more likely to saccade to a distractor location (i.e. make an error), in the presence of a 100 % valid central cue, when V1 activity that correlates with spatial-value remains detectable in the signal[35]. This suggests that V1 neurons may need to reverse a spatial-value signal in order to represent the location indicated by the spatial-cue. The current results extend these findings by demonstrating that the implementation of visual-selection biases remains additive, or serial, even when the spatial-value cue is no longer made redundant by the central cue (i.e. when the central cue does not provide complete certainty regarding upcoming target locations). Thus, the current results show that

relative expectancies must be computed separately for spatial-certainty and spatial-values, and that visual-selection reflects the serial access to each information source.

It therefore must be of greater benefit to maintain this separation than it is to combine spatial-certainty and spatial-value information to guide behaviour towards a target. If we believe that the goal of the system is to read out the target as quickly and accurately as possible, to ensure the accrual of higher reward gains, this appears somewhat non-intuitive. If there is high spatial-certainty that a target will appear at a given location, and that location is associated with high spatial-value, it makes sense to combine these probabilities multiplicatively to maximally increase sensitivity to the upcoming target location. However, we have shown here that the system appears to apply a summation of these information sources, even when given cause not to do so. If however, we assume that the overarching goal is instead to accurately model an ever changing environment, then it seems reasonable that the computations pertaining to spatial-certainty and spatial-value are kept separate. Serial computation of various outcome-associations buys the advantage that the output of any given computation can be fed as an input into any other subsequent computation. Therefore, it could be that serial processing of learned-associations buys flexibility in terms of which information sources can be learned about and combined in order to guide behaviour.

The current finding of additivity across learned-associations also accord with computational[3] and neurophysiological[44] models of visual-selection. Although these models are aimed at understanding distinct properties of selective-attention, one convergent principle is that additivity across feature dimensions can predict a range of visual-selection phenomena. The current work suggests that additivity across feature dimensions can also apply to the associations made between physical stimuli and their consequent outcomes. When the current findings are interpreted in this framework, it suggests that at a computational level, the visual system treats learned associations comparably to physical features. However, within a physically constrained system, learned-associations and physical features cannot sum ad infinitum. Future work must identify the limits of additivity across feature dimensions in order to understand where lies the information processing constraints of the system.

Our results are in apparent contradiction to previous work showing that saccadic onset latencies correlated with expected-value[39]. There are at least two possible reasons for this difference. First, the predictions made by the additive-bias and expected-value hypotheses are similar, so it may be that the additive-bias hypothesis can still provide a better model for the saccadic onset times observed by Milstein and Dorris[39], and that either the additive-bias or expected-value model would perform better than one based on spatial-certainty or spatial-value alone. A good first test would be to apply the current analysis to the saccadic response time data, to directly pit the additive-bias and expected-value models in this context.

A second alternative that can unite these findings relates to the timing with which spatial-certainty and spatial-value expectancies were resolved in this and previous work. In the current study, distinguishing the target from the distractor was required to inform whether or not the likely location and value outcomes had occurred, therefore uncertainty pertaining to both associations were resolved close together in time (or concurrently), once the target was detected. In contrast, Milstein and Dorris[39] presented single targets in the visual field. Thus uncertainty regarding spatial location was immediately resolved at target onset. This may have conferred a small temporal advantage for the resolution of spatial uncertainty relative to value uncertainty. Previous findings show that this temporal advantage can influence whether expectancies regarding location and orientation interact or are additive[25], suggesting that confirmation of one expectation (location) acts to magnify the influence of subsequent expectations (orientation). This suggests that if spatial-certainty and spatial-value influence one another in the same manner, then an additive influence would be predicted in the current study, and a multiplicative influence would be predicted for the saccade task used by Milstein and Dorris[39]. According to this account, it should be possible to change the influence of spatial-certainty and spatial-value on performance by varying the time within which uncertainty pertaining to both is resolved. However, for this to be true, each learned-association must act serially on the mechanisms underlying visual-selection, which itself constrains the possible use of an expected-value computation to guide performance.

It should also be noted that the approach we used here assumes that the influence of spatial-certainty and spatial-value on performance is one that can be sufficiently modelled using a linear combination of predictors. This is reasonable, given a lack of theoretical or empirical evidence to assume a more complex relationship. However, examination of figures 2 and 3 show that model fits can potentially be improved. What we know from the current analysis is that it cannot be improved by adding any linear combination of the existing terms. However, we also know that even though we had positive evidence for the additive-bias model, models allowing interactions between location-certainty and location-value appeared within the top 6 models across both experiments. We therefore cannot rule out whether modelling more complex relationships would reveal evidence for an expected-value computation in visual-selection. However, given the current work, we can demonstrate that if the assumption of linearity holds, then expected-value computations are unlikely to drive visual-selection.

Conclusions

Over two experiments, we sought to arbitrate between competing theories for how learned associations pertaining to spatial-value and spatial-certainty combine to influence visual-selection. Specifically, we asked whether this influence was additive (*additive-bias hypothesis*), multiplicative (*expected-value hypothesis*) or both (*mixed-operations hypothesis*). We tested these hypotheses by pitting spatial-certainties and spatial-values against one another in a spatial-cueing task under conditions expected to challenge the optimality of an additive operation. In Experiment 1 we tested whether approaching maximal spatial-certainty modulated the extent to which spatial-value acted on visual-selection performance. In Experiment 2 we tested whether changing the task reward structure to favour a multiplicative operation could result in a non-additive influence. Over both experiments we observed unambiguous support for the additive-bias hypothesis. The data support the notion that visual-selection mechanisms show serial sensitivity to spatial-certainty and spatial-value information. We interpret our results in accordance with computational models of visual-selection and suggest that the visual system treats learned associations comparably to physical features when prioritising information processing.

Method

Experiment 1

Participants

As larger samples protect against spurious findings[45], we opted to double the sample size of previous work correlating human performance with expected-value (N=10)[39], and recruit a minimum of 20 participants. We calculated the stopping rule for data collection as the number of weeks where testing at maximum capacity would bring us to at least the minimal sample size (6 weeks with 4 people per week).

A total of 23 participants were recruited. Of these, 1 was excluded due to technical failure and a second due to experimenter error. The remaining 21 participants (19 female, 18 right-handed, mean age = 20.3, sd 4.5) completed all the procedures. Participants were recruited if they were aged 18 years or over and reported normal or corrected-to-normal vision; no history of psychiatric or neurological illness, injury, or disorder. Participants earned either course credit or payment (£7 per session), and any additional rewards accrued during the session (~£10). All procedures were approved by the University of Birmingham Human Research Ethics Committee and were within the guidelines of the the National Statement on Ethical Conduct in Human Research.

Apparatus

All tasks were programmed in Matlab (Mathworks, Natick, MA, 2013a), using the Psychophysics Toolbox extension[46,471]. The tasks were run on a Stone SOFREP-144 computer with a 23-inch Asus VG278HE monitor (1920 x 1080 pixels, 60-Hz refresh) viewed from 57 cm.

Stimuli

Two vertical lines [RGB: 200, 200, 200] were presented in the centre of the screen. A darkening of one of the lines [50, 50, 50] served as the endogenous spatial cue. Two coloured discs matched for luminance (purple [87, 75, 80] and orange [120, 86, 1]) served as value cues. They were presented on

[1] All code (experimental and analysis) and data are available online at github.com/kelly-garner-git-stuff/ADDBIAS

the horizontal meridian, 4.5° from the centre. All stimuli were presented on a grey [RGB: 118, 118, 118] background. Each value cue was 2.2° in diameter. Targets ['H' or 'N'] and distractors ['Z' or 'K'] were presented in Helvetica font [90, 90, 90] and encompassed 1°. Feedback was presented in green [0, 255, 0] for high reward values, amber [255, 191, 0] for low reward values, and red [255, 0, 0] for errors.

Procedure

The procedures are presented in Figure 1. Each trial began with two value cues and two centrally presented vertical lines for a pseudo-randomised duration of 300-500 ms. Value-cue location was counterbalanced across trials, and colour/value pairings (e.g. purple = 50 points/orange = 1 point) were counterbalanced across participants. One of the fixation lines then darkened for 300 ms. After a further 100 ms, the target and distractor were presented for 100 ms (target identity was equiprobable for each cue-certainty x cue-value x cue-validity condition). Participants responded with the 'v' and 'g' keys to indicate the target identity. Target and response-mappings were counterbalanced across participants. 500 ms after the response, feedback was presented for 750 ms; either the central fixation was replaced with the high reward value, the low reward value, an error signal (fixation lines turned red), or the fixation remained the same (no reward). Rewards were awarded on 80 % of trials, given the participant made the correct response. This reinforcement schedule was selected to ensure that the value feedback signals did not become redundant - i.e. they could not be fully predicted.

Across blocks, the likelihood of cue-validity was varied to be either .6 valid/.4 invalid, .8/.2, .9/.1, .92/.08, .96/.04, resulting in information gains (spatial-certainty) of .029, .29, .53, .6 and .86 bits. Each block contained 100 trials. Over 4 sessions, participants completed 4 blocks for each level of spatial-certainty. Participants took between 4 days and 1.5 weeks to complete the experiment (block order was pseudo-randomised for each session). Target-value contingencies were split equally within each set of valid and invalid trials for each cue-likelihood condition

Participants were instructed that the value cues signified how many points were available should the target appear in that location (50 vs 1 point), and that the darkened line was a clue for where the

target could appear. They were requested to keep their eyes at fixation, and to respond as accurately and as quickly as possible to the target. Participants were also informed that their points would be exchanged for cash at the end of the session (1000 = £1). At the start of the first session, participants practiced until they achieved at least 16/20 correct responses.

Statistical Approach

Data pre-processing

All data was analysed using the R programming language (v3.3.2)[48], and R Studio (v1.0.44)[49]. RT data were rejected if they were greater than +/- 2.5 standard deviations from the mean for that participant in that condition. Participants were not explicitly informed when there was a change in spatial-certainty. We assumed that trials immediately subsequent to changes in spatial-certainty would be contaminated by learning effects. To remove the contaminated trials for each participant, we collapsed the data across spatial-certainty blocks, and ordered the data according to trial number. We then fit piecewise linear regressions to find the break point that minimized the mean square error (MSE). Trials occurring prior to the breakpoint were removed (mean = 12.3, sd 8.0). However, when we performed the analyses without removing these trials, the pattern of results was the same.

Model selection

The aim of the study was to compare whether an additive-bias model remained the best model, given the data, even under conditions where an additive relationship could be expected to break down. The key aim of each analysis was to determine whether a model that included an interaction term between spatial-certainty and spatial-value was more probable, given the data, than one that only included main effects. We therefore sought to quantify the evidence in favour of one model relative to another. To achieve this, we used a Bayesian approach. First, we fit all possible linear mixed models, and computed log inverse Bayes Factors (BFs), which quantifies evidence for each possible linear mixed effects model against the null model (intercept plus random effects of participant) using the Bayes Factor package[50], and implementing the default Jeffreys-Zeller-Siow (JZS) prior on slope estimates[51]. We then

identified the 6 best performing models. We report the BF of the winning model relative to the null model, and the BF ratios between the best model and the next 5 best models, to quantify the evidence in favour for the winning model. We follow the guidelines of Kass and Rafferty[52] when interpreting the strength of evidence. This was typically sufficient to determine whether the evidence favoured a model that included only main effects, or a certainty x value interaction. However, in cases where this was not sufficient, further targeted comparisons were also made. All BFs are reported along with the proportional error of the estimate.

Experiment 2

Participants

We calculated the stopping rule for data collection as the number of weeks where testing at maximum capacity would bring us over the minimal sample size (3 weeks with 10 people per week). Of the 28 participants recruited, 1 was excluded due to technical difficulties with the eyetracker. A second participant was excluded as they did not meet the criterion required to terminate the practice. The remaining 26 (mean age = 19.5 years, sd = 1.03, 24 F, 26 right-handed) completed all the study procedures. Two of these participants had also completed Experiment 1.

Apparatus

Although participants were clearly instructed not to move their eyes in Experiment 1, we wanted to ensure that we could replicate our results (for the static condition at least) when we have an objective check that eye-movements did not occur, and thus did not contribute to the results. Therefore, in addition to Experiment 1, an Eyelink® 1000 desktop-mounted eye-tracker (SR Research Ltd., Ottawa, Ontario, Canada) recorded movements of the left eye with a sampling frequency of 500 Hz.

Stimuli

The stimuli were the same as in Experiment 1, except that the value cues were presented at 5.7°. This change was made to match the exact layout used in previous work[36].

Procedure

The procedure was the same as Experiment 1 with the following exceptions. Participants' eyes were monitored on every trial. If the participant's eyes moved more than 50 pixels from the fixation at cue-offset, text appeared to notify participants they had been "too-fast". The trial was then terminated. Terminated trials accounted for 3 % of all trials.

Cue-values were increased from Experiment 1 to 5000 vs 100 points, so that participants could gain at least 1 point when a decay was applied to the low spatial-value. In the decay reward-condition, an exponential decay function was applied to each value at target onset. The relative value between the two was maintained throughout the decay period. The monetary value of points was adjusted so that participants received the same rate of cash payments as Experiment 1 (100,000 = £1). Participants were informed at the start of the decay blocks that the value available to them would begin to run out upon appearance of the target.

Participants completed 200 trials for each of 4 spatial-certainty/reward-conditions (.29/ fixed, .29/decay, .029/fixed, .029/decay; block order was counterbalanced across participants). We included only these two levels of spatial-certainty as we wanted to avoid any possible floor or ceiling effects when testing the influence of reward-condition.

We also tested the separate hypothesis that individuals may mentally represent the high and low-incentive placeholders differently in terms of their relative value, when their value can be obtained more reliably (i.e. in the static reward-condition, relative to the decay reward-condition), and that this may be expressed via physical placement on a linear space. Every 50 trials, participants were instructed to use a mouse to drag the two placeholders wherever they liked on a visual analogue scale. However, we found no evidence that cue-likelihood influenced placement of the placeholders ($p = .96$), and this aspect

of the study is discussed no further. Participants also completed a BIS/BAS questionnaire [53] that was used to test a hypothesis for a separate study not reported here.

Statistical Approach

We followed the same data cleaning procedures as Experiment 1. Again, piecewise linear functions were fit to the data to isolate the trials contaminated by spatial-certainty learning effects. The number of trials removed from the start of each block were similar to Experiment 1 (mean = 14.7, sd 8.5).

We also used the same model comparison approach, with the exception that we added the reward-condition term to the linear mixed effects models that were fit to the data.

References

1. Malebranche, N. & Olscamp, P. J. *Malebranche: The Search After Truth: With Elucidations of The Search After Truth*. (Cambridge University Press, 1997).
2. Buschman, T. J. & Kastner, S. From Behavior to Neural Dynamics: An Integrated Theory of Attention. *Neuron* **88,** 127–144 (2015).
3. Itti, L. & Koch, C. Computational modelling of visual attention. *Nat. Rev. Neurosci.* **2,** 194–203 (2001).
4. Posner, M. I. Orienting of attention. *Q. J. Exp. Psychol.* **32,** 3–25 (1980).
5. Le Pelley, M. E., Mitchell, C. J., Beesley, T., George, D. N. & Wills, A. J. Attention and Associative Learning in Humans: An Integrative Review. *Psychol. Bull.* (2016). doi:10.1037/bul0000064
6. Raymond, J. E. & O'Brien, J. L. Selective Visual Attention and Motivation. *Psychol. Sci.* (2009). doi:10.1111/j.1467-9280.2009.02391.x
7. Chelazzi, L. *et al.* Altering spatial priority maps via reward-based learning. *J. Neurosci.* **34,** 8594–8604 (2014).
8. Anderson, B. A., Laurent, P. A. & Yantis, S. Value-driven attentional capture. *Proc. Natl. Acad. Sci.*

U. S. A. **108,** 10367–10371 (2011).

9. Le Pelley, M. E., Pearson, D., Griffiths, O. & Beesley, T. When goals conflict with values: counterproductive attentional and oculomotor capture by reward-related stimuli. *J. Exp. Psychol. Gen.* **144,** 158–171 (2015).
10. Awh, E., Belopolsky, A. V. & Theeuwes, J. Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends Cogn. Sci.* **16,** 437–443 (2012).
11. Posner, M. I., Snyder, C. R. & Davidson, B. J. Attention and the detection of signals. *J. Exp. Psychol.* **109,** 160–174 (1980).
12. Desimone, R. & Duncan, J. Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* **18,** 193–222 (1995).
13. Carrasco, M. Visual attention: the past 25 years. *Vision Res.* **51,** 1484–1525 (2011).
14. Thompson, K. G. & Bichot, N. P. A visual salience map in the primate frontal eye field. *Prog. Brain Res.* **147,** 251–262 (2005).
15. Treue, S. Visual attention: the where, what, how and why of saliency. *Curr. Opin. Neurobiol.* **13,** 428–432 (2003).
16. Bisley, J. W. & Goldberg, M. E. Attention, intention, and priority in the parietal lobe. *Annu. Rev. Neurosci.* **33,** 1–21 (2010).
17. Ipata, A. E., Gee, A. L., Bisley, J. W. & Goldberg, M. E. Neurons in the lateral intraparietal area create a priority map by the combination of disparate signals. *Exp. Brain Res.* **192,** 479–488 (2009).
18. Gottlieb, J. P., Kusunoki, M. & Goldberg, M. E. The representation of visual salience in monkey parietal cortex. *Nature* **391,** 481–484 (1998).
19. Deco, G., Pollatos, O. & Zihl, J. The time course of selective visual attention: theory and experiments. *Vision Res.* **42,** 2925–2945 (2002).
20. Deco, G. & Zihl, J. Top-down selective visual attention: A neurodynamical approach. *Vis. cogn.* **8,** 118–139 (2001).
21. Chica, A. B., Martín-Arévalo, E., Botta, F. & Lupiáñez, J. The Spatial Orienting paradigm: how to

design and interpret spatial attention experiments. *Neurosci. Biobehav. Rev.* **40,** 35–51 (2014).

22. Lanthier, S. N., Wu, D. W.-L., Chapman, C. S. & Kingstone, A. Resolving the controversy of the proportion validity effect: Volitional attention is not required, but may have an effect. *Atten. Percept. Psychophys.* **77,** 2611–2621 (2015).
23. van der Heijden, A. H. Probability matching in visual selective attention. *Can. J. Psychol.* **43,** 45–52 (1989).
24. Jonides, J. Towards a model of the mind's eye's movement. *Can. J. Psychol.* **34,** 103–112 (1980).
25. Kingstone, A. Combining Expectancies. *The Quarterly Journal of Experimental Psychology Section A* **44,** 69–104 (1992).
26. Platt, M. L. & Glimcher, P. W. Neural correlates of decision variables in parietal cortex. *Nature* **400,** 233–238 (1999).
27. d'Acremont, M., Fornari, E. & Bossaerts, P. Activity in inferior parietal and medial prefrontal cortex signals the accumulation of evidence in a probability learning task. *PLoS Comput. Biol.* **9,** e1002895 (2013).
28. Prinzmetal, W., Whiteford, K. L., Austerweil, J. L. & Landau, A. N. Spatial attention and environmental information. *J. Exp. Psychol. Hum. Percept. Perform.* **41,** 1396–1408 (2015).
29. Hick, W. E. On the rate of gain of information. *Q. J. Exp. Psychol.* **4,** 11–26 (1952).
30. Hyman, R. Stimulus information as a determinant of reaction time. *J. Exp. Psychol.* **45,** 188–196 (1953).
31. Shannon, C. E. A Mathematical Theory of Communication. *Bell System Technical Journal* **27,** 379–423 (1948).
32. Sugrue, L. P., Corrado, G. S. & Newsome, W. T. Matching behavior and the representation of value in the parietal cortex. *Science* **304,** 1782–1787 (2004).
33. Dorris, M. C. & Glimcher, P. W. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* **44,** 365–378 (2004).
34. Louie, K. & Glimcher, P. W. Separating value from choice: delay discounting activity in the lateral

intraparietal area. *J. Neurosci.* **30,** 5498–5507 (2010).

35. Stănişor, L., van der Togt, C., Pennartz, C. M. A. & Roelfsema, P. R. A unified selection signal for attention and reward in primary visual cortex. *Proc. Natl. Acad. Sci. U. S. A.* **110,** 9136–9141 (2013).
36. Stankevich, B. A. & Geng, J. J. Reward associations and spatial probabilities produce additive effects on attentional selection. *Atten. Percept. Psychophys.* **76,** 2315–2325 (2014).
37. Sternberg, S. *Discovering mental processing stages: The method of additive factors*. (The MIT Press, 1998).
38. Geng, J. J. & Behrmann, M. Spatial probability as an attentional cue in visual search. *Percept. Psychophys.* **67,** 1252–1268 (2005).
39. Milstein, D. M. & Dorris, M. C. The influence of expected value on saccadic preparation. *J. Neurosci.* **27,** 4810–4818 (2007).
40. Drugowitsch, J., DeAngelis, G. C., Angelaki, D. E. & Pouget, A. Tuning the speed-accuracy trade-off to maximize reward rate in multisensory decision-making. *Elife* **4,** e06678 (2015).
41. Kiani, R. & Shadlen, M. N. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* **324,** 759–764 (2009).
42. Serences, J. T. Value-based modulations in human visual cortex. *Neuron* **60,** 1169–1181 (2008).
43. Morey, R. D. Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutor. Quant. Methods Psychol.* **4,** 61–64 (2008).
44. Deco, G. & Rolls, E. T. Neurodynamics of biased competition and cooperation for attention: a model with spiking neurons. *J. Neurophysiol.* **94,** 295–313 (2005).
45. Button, K. S. *et al.* Power failure: why small sample size undermines the reliability of neuroscience. *Nat. Rev. Neurosci.* **14,** 365–376 (2013).
46. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10,** 433–436 (1997).
47. Pelli, D. G. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* **10,** 437–442 (1997).
48. R Core Team. The R project for statistical computing. *Available at www. R-project. org/. Accessed*

October (2013).

49. RStudio Team. RStudio: Integrated Development Environment for R. (2016).
50. Morey, R. D., Rouder, J. N. & Jamil, T. BayesFactor: Computation of Bayes factors for common designs. *R package version 0. 9* **8,** (2014).
51. Liang, F., Paulo, R. & Molina, G. Mixtures of g Priors for Bayesian Variable Selection. *J. Am. Stat. Assoc.* **103,** 410–423 (2008).
52. Kass, R. E. & Raftery, A. E. Bayes Factors. *J. Am. Stat. Assoc.* **90,** 773–795 (1995).
53. Carver, C. S. & White, T. L. Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS Scales. *J. Pers. Soc. Psychol.* **67,** 319 (1994).