Incentive value and spatial certainty combine additively to determine visual priorities

K.G. Garner, H. Bowman, & J.E. Raymond

School of Psychology, University of Birmingham

Correspondence to:

K.G. Garner

getkellygarner@gmail.com

Queensland Brain Institute (79)

University of Queensland

St Lucia, QLD, 4072

Key words: attention, prediction, expectation, reward, incentive

Word count: 8629

Acknowledgements

K.G. Garner would like to acknowledge Christopher Nolan, Andrew Clouter and Dragan Rangelov for the comments and insightful discussions about this work.

The authors declare no competing interests.

Abstract

How does the brain combine information predictive of the value of a visually guided task (incentive value) with information predictive of where task relevant stimuli may occur (spatial certainty)? Human behavioural evidence indicates that these two predictions are combined additively to bias visual selection (additive hypothesis), whereas neuroeconomic studies posit that they may be multiplicatively combined (expected value hypothesis). We sought to arbitrate between these two alternatives, and to test the possibility that both operations are available to visual prioritization mechanisms, but that their use is context dependent (mixed operations hypothesis). Participants viewed two coloured placeholders that specified the potential value of correctly identifying an imminent letter target if it appeared in that placeholder. Then, prior to the target's presentation, an endogenous spatial cue was presented indicating the target's more likely location. Spatial cues were parametrically manipulated with regard to the information gained (in bits). Across two experiments, response time and accuracy were greater for targets appearing in high versus low value placeholders and higher when targets appeared in validly cued locations, even under conditions designed to impinge the optimality of an additive operation. Interestingly, these factors did not interact; Bayesian model selection showed that the additive hypothesis clearly outperformed the expected value and mixed operations hypotheses in accounting for the observed data from both experiments. These findings refute theories that expected value computations are the singular mechanism driving the deployment of endogenous spatial attention. Instead, incentive value and spatial certainty seem to act independently to influence visual selection.

Incentive value and spatial certainty combine additively to determine visual priorities

Humans are good at learning that specific sensory information, or cues, can predict subsequent events. For example, we learn quickly that hearing a siren on the left predicts a fast moving emergency vehicle appearing from that direction, or that seeing a smile predicts a likely future opportunity to gain social approval. Knowledge about where and when new, important sensory information may appear or new reward opportunities may arise is only useful, however, if such knowledge can be deployed to modulate and control the neural processes that control behaviour. Our understanding of exactly how learning and experience act to modify such prioritisations of visual signals, i.e. *visual selection*, is still far from complete. Particularly, it remains unclear how multiple concurrent sensory cues, each associated with and therefore predictive of specific consequent outcomes are combined to influence goal-directed visual selection. In the current study, we seek to understand how two types of learned associations are combined by the brain to influence visual selection behaviours. These are (1) the association between sensory cues and available reward outcomes (incentive value) and (2) the association between sensory cues and response-defining information, in this case the probability of the upcoming target location (spatial certainty).

Many cognitive, computational, and neurobiological theories of visual selection (see Itti and Koch 2001; Buschman and Kastner 2015; Moore and Zirnsak 2017, for reviews), assume that competition for high-level neural representation of external objects is flexibly biased by goal-directed mechanisms. For example, if information is provided that indicates the most likely location of an imminent task-relevant visual target, visual processing of information appearing at the predicted location is faster than for that for information appearing at other locations. This clear demonstration of the influence of internal knowledge on visual selection has been widely studied using the spatial cueing paradigm (Posner, Snyder, and Davidson 1980). In such studies, including the current one, participants observe two locations, typically indicated by placeholders, and then are required to identify a briefly presented pre-specified target that suddenly appears in one of the placeholders (left or right with chance probability). At the same time, a distracting stimulus appears in the other placeholder, requiring active selection of the target. Preceding target onset, a symbolic cue such as

an arrow indicates one of the potential locations. Even though the probability that a target will appear on the left versus right is at chance, participants are faster to respond when the target appears in the cued location, relative to when it appears in the uncued location. The influence of symbolic spatial cues on performance is assumed to reflect the leveraging of information pertinent to current goals, such as using accrued knowledge that arrows are spatially informative, to bias visual selection towards the cued location.

Recent investigations have focused on the possibility of multiple parallel endogenous influences on visual selection. In particular, studies asked how previous experience regarding associations between sensory cues and meaningful outcomes can influence a cues' ability to gain visual priority (see Hutchinson and Turk-Browne 2012; Le Pelley et al. 2016, for reviews), even when processing the cue is at odds with the current goal as defined by the task-set (see Awh, Belopolsky, and Theeuwes 2012, for a review). One example of such an association is when cues predict information relevant to success in the current task, such as directional cues that provide probabilistic information regarding potential target locations (predictive spatial cues). Another, yet distinct, example is when a stimulus signals an opportunity to gain a reward for appropriately engaging in a visual task. Such incentive cues typically indicate the magnitude of rewards on offer but are uninformative about visual targets or the actions required to acquire the reward. Here, we asked how these two sources of influence, spatial cues and incentive cues, might be combined by the brain to prioritise visual selection.

The effects of predictive spatial cues on attention has been intensively studied and is very well established. Critically, spatial cueing effects, indexed by response time slowing when cues are invalid relative to valid, scale with the reliability of the cue; i.e., the higher the probability that the target will appear in a location, given the cue, the larger the difference in response times to valid versus invalidly cued targets (Lanthier et al. 2015; van der Heijden 1989; Jonides 1980; Kingstone 1992). This occurs, even though participants may be unaware of cue-target contingencies (Lanthier et al. 2015). Indeed, Prinzmetal et al. (2015) showed that spatial cueing effects follow the Hick-Hyman law of decision-time (Hick 1952; Hyman 1953); i.e., the size of response time (RT) benefits scale linearly with the spatial certainty gained (in bits) by a spatial cue. To compute spatial certainty they used

Shannon's (Shannon 1948) measure of entropy (H) as we did here (see Methods). It currently remains unknown how spatial certainty is combined with other predictive cues, such as incentive cues that provide information about upcoming potential reward magnitudes.

Multiple previous studies have shown that when stimuli probabilistically predict the reward magnitude available for correct performance on a task, they become value-associated and may influence subsequent visual selection even when rewards cease to be available (Raymond and O'Brien 2009; Chelazzi et al. 2014; although see Sha and Jiang 2016). Critically, it has also been shown that when such incentive cues are irrelevant to current task objectives they may nevertheless interfere with performance, indicating their power to influence visual prioritisation mechanisms (Anderson, Laurent, and Yantis 2011; Le Pelley et al. 2015). Of course, in many everyday situations (e.g., in computer games), incentive cues serve to energize or facilitate visual selection tasks, observations supported by controlled experiments (Kiss, Driver, and Eimer 2009; Small et al. 2005; Pessoa and Engelmann 2010). Indeed, the presentation of incentives cues prior to simple visual tasks has been shown to not only decrease response times but also to produce modulations of electrophysiological correlates of visual readiness and selection (Sawaki, Luck, and Raymond 2015). These and related data from selective attention tasks in monkeys (Stănişor et al. 2013) provide strong support for the notion that visual selection mechanisms are biased by brain mechanisms that code the incentive value of specific stimuli.

Considering the possibility that there may be multiple influences acting on goal-directed visual selection, a question that arises is how might incentive value and spatial certainty be combined to influence competition among sensory signals, and thus visual selection? It has recently been shown in a study by Stănişor et al (2013) that the mechanism that mediates the ultimate influence of incentive value and spatial certainty on visual selection may be unitary. This finding suggests that there may be a point where prioritising information from one source is either integrated with or overwritten by the other and that a singular mechanism ultimately determines visual selection. In the Stănişor et al (2013) study, monkeys gained different sized rewards by making saccades to one of two different coloured circles appearing at two different locations. Circle colour signalled reward magnitude, and each circle was equally likely to become a target prior to the presentation of a 100%

valid spatial cue (a line connecting fixation to one of the circles). Saccades to the cued target were selectively rewarded when made either 50ms or 400ms after spatial cue onset. When the response was made shortly after spatial onset (50 ms), V1 spiking activity measured by multiunit electrode recordings correlated with the relative reward values associated with the two circles. However, when the onset of the saccade was delayed (400 ms condition), V1 activity correlated with the spatial location of the cued target, regardless of its associated reward value. Across the short and long saccade onset conditions, responses to both value and spatial cues showed a similar latency of onset (~120 ms) and the strength of the response to incentive value cues predicted the strength of the response to spatial cues, suggesting overlap in the neurons sensitive to both signals. Finding that the incentive value response was abolished by the spatial certainty information motivated the conclusion that the symbolic, spatial cue served to reweight the relative incentive values of the two coloured circles (Stănişor et al. 2013). For example, if a spatial cue signalled that the low incentive value circle was the target on that trial, it nullified the incentive value offered by the now irrelevant high incentive value cue. However, as all cues were 100% predictive of target location in this study, it sheds no light on how variation in spatial certainty might be combined with incentive value to guide visual selection.

In the current study, we examine three plausible alternatives for how incentive value and spatial certainty might be combined to control selective attention. The first of these is the expected value hypothesis, a concept that has its origins in economic theories of decision making (Morgenstern and Von Neumann 1953). According to this view, visual selection is biased by relative expected value, i.e., incentive value multiplied by spatial certainty for each outcome (which directly determines reward likelihood), and normalised across potential outcomes given the current trial context. (Traditionally, the expected value computation is assumed to draw from the probability of an incentive value offered by a single stimulus). Indeed, human saccadic response times have been shown to correlate with the relative expected value of potential targets (Milstein and Dorris 2007). In that study, participants were required to saccade to a single red dot that appeared within the left or right hemi-field of a blank display. Over blocks, the likelihood of a specific target location varied as well as its location-contingent reward value. The time taken to initiate a saccade to the target was more tightly correlated with the relative expected value of the potential target location than with either the

relative likelihood of the target location, or its relative reward value. However, these learned associations were tied to the appearance of the target rather than to *a priori* sensory cues, meaning they could not be used to bias selective attention before the target's appearance. Moreover, previous research has shown that similar interactive effects between spatial and non-spatial expectations regarding target identity occur when targets are presented in isolation, but become additive when presented with a concurrent distractor (Kingstone 1992). Therefore, it is unclear whether an expected value operation would hold when covert attention is required to arbitrate the competition between sensory signals.

The second hypothesis we tested is the additive hypothesis, an idea given tentative support by Stankevich and Geng (2014). They asked participants to detect as quickly and accurately as possible a simple target that could appear on the left or right within a coloured placeholder. In their experiment, placeholder colour indicated the magnitude of response-contingent rewards, as it does in the experiment we report here. However, they provided no explicit spatial cues as to target location. Instead, across blocks and without instruction, the probability of the target appearing on one side versus the other was varied. Greater performance benefits were observed when the target appeared in the more probable location and when that location corresponded to a high versus low value placeholder. Critically, these benefits were additive which suggests that incentive value and spatial certainty acted independently to influence visual selection, according to additive-factors logic (Sternberg 1998). However, in their experiment, spatial-certainty was built up over many trials, allowing predictions about target location to be generated well before each trial began and certainly in advance of location-specific incentive information. This may explain why incentive information provided an additive "top-up" effect. Such effects might not occur when location-specific incentive information is available first and spatial certainty cues providing task relevant information are presented subsequently, as in the Stanisor et al. (2013) study. Therefore, it remains unknown how two different endogenous cues (incentive and spatial) might be combined in humans when they are available only via new sensory information as each trial unfolds.

We also considered a third hypothesis, which we refer to here as the mixed operations hypothesis. Perhaps the selective attention system maintains independence of incentive value and

spatial certainty biases when these predictions are orthogonal, as is the case when experimental designs ensure that variables are unconfounded. Yet, there may be limits to the additive nature of these influences, after which a non-additive relationship emerges, given that the same mechanism is influenced by both information sources. If this were the case, the influence might be additive until the system is sufficiently stressed, after which an expected value operation may be revealed.

The central aim of the two experiments reported here was to directly compare predictions made by these three hypotheses, using data from behavioural experiments on humans that combined the methods of relevant previous studies. In Experiment 1, we specifically tested conditions where spatial-certainty approached maximum, densely sampling probabilities between .8 and 1. We reasoned that if incentive value and spatial certainty were combined serially and additively, then additivity should be maintained across all levels of spatial certainty (*additive hypothesis*). In contrast, if the two influences were integrated in some way, then 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 spatial cue becomes close to 100 % valid (*mixed operations hypothesis*). We reasoned this might occur because as the computation of spatial certainty becomes trivial, the mechanism using both information sources could have left-over resources to implement the influence of incentives. To anticipate, we find that the influence of incentive value and spatial certainty remains additive across all tested levels of spatial certainty, arguing against the expected value and mixed operations hypotheses.

In Experiment 2, we tested whether the reward structure of the task could influence the combination of spatial certainty and incentive biases. Previous studies show that learning can direct the sampling of sensory information to optimise reward accrual (Drugowitsch et al. 2015; Kiani and Shadlen 2009; Serences 2008). We reasoned that if an expected value computation is available to direct visual selection, then reward conditions that favour this operation should yield a non-additive influence of incentive and spatial certainty on performance. Given that expected value computations are multiplicative, they should produce super-additive effects when both incentive values and spatial certainty are high, and sub-additive effects when incentive value and spatial certainty are low. As RTs to targets appearing at specific spatial locations should scale inversely with the expected value for

that location, RTs driven by an expected value operation should be faster than those driven by an additive operation when spatial certainty and incentive value are high, and should be slower than an additive operation when both spatial certainty and incentive value are low. Therefore, reward conditions that preferentially reward fast RTs on high incentive/certainty trials, and that minimize costs incurred for slow RTs on low incentive/certainty trials should favour an expected value operation over an additive operation (see Figure 1d). To effect these reward conditions, in Experiment 2, participants completed the same task as in Experiment 1 (albeit sampling fewer levels of spatial certainty), with an added condition wherein reward value exponentially decayed after target onset. Again, and contrary to the expected value and mixed operation hypotheses, incentive value and spatial certainty additively combined to drive visual selection biases. Collectively the results favour the additive hypothesis, suggesting that incentive value and spatial certainty act independently to influence visual selection.

Experiment 1

Method

All the task and analysis code, and data from the current study are available online[1]. The trial sequence of the spatial-orienting task (Posner 1980) used to assess the combined influence of spatial certainty and location specific incentive cues, and the key manipulations for Experiments 1 and 2, are shown in Figure 1.

[1] https://github.com/kel-github/attention-value-certainty

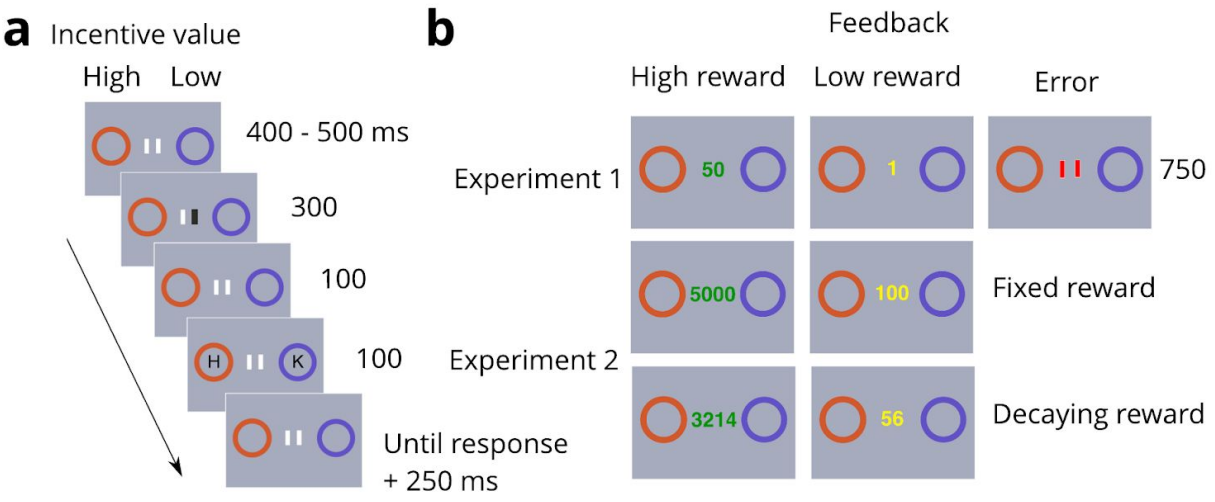

c

Cue likelihood (Proportion valid/invalid)					
	.96/.04	.92/.08	.9/.1	.8/.2	.6/.4
Spatial-certainty (information gain in bits)					
Exp 1	.86	.6	.53	.29	.029
Exp 2				.29	.029

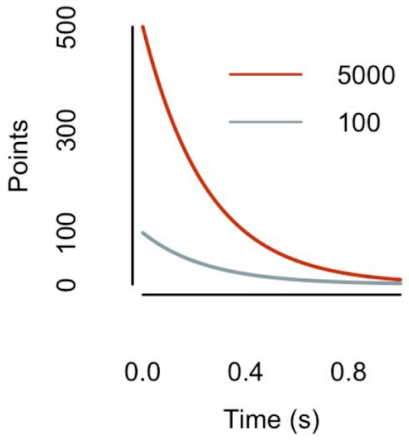

Figure 1: Study method. Task procedure and feedback conditions for Experiment 1 and Experiment 2. **a.** Trial structure: Participants monitored two different coloured circular placeholders (incentive cues). Colour indicated the magnitude of a performance-contingent reward for correct target ("H" or "N") identification, should the target subsequently appear within that placeholder. Then, one of two central bars darkened, indicating the more likely target location (left, right). **b.** Reward feedback structure: After response + 250 ms, performance feedback and response-contingent rewards were presented as shown. In Experiment 2, reward feedback was either independent of response time (fixed) or decremented exponentially after target onset until response (decaying). **c.** Spatial certainty was parametrically manipulated across blocks by increasing the information gained (in bits) from the spatial cue. **d.** Logic of the decaying reward condition in Experiment 2. Figure shows reward value available as a function of time from target onset. As both expected value and mixed operations computations involve a multiplicative weighting of spatial certainty and incentive value, responses should be super-additive or sub-additive depending on the spatial certainty/incentive value combination. As response times should reflect the inverse of this weighting, responses should be faster in a high certainty/high incentive scenario than responses based on an additive operation, and slower in a low certainty/low incentive scenario than an additive operation. Applying an exponential decay function to the incentive value at target onset means that the extra rewards accrued by being faster towards high incentive value cues (change in the high (5000) value on the y-axis, while moving leftward on the x-axis) would outweigh the losses accrued from being slower towards low incentive value cues (change in the low (100) value on the y-axis, while moving rightward on the x-axis). Therefore, any operation that favours this response pattern would accrue greater total rewards than an additive operation, and therefore may emerge under such reward conditions.

Participants

As larger samples protect against spurious findings (Button et al. 2013; Lorca-Puls et al. 2018), we opted to double the sample size of previous work correlating human performance with expected-value (N=10) (Milstein and Dorris 2007), 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 minimum sample size (6 weeks with 4 people per week, allowing recruitment of >20 participants in order to protect against drop outs). Participants were recruited if they were aged 18 years or over and reported normal or corrected-to-normal vision, with 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. 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.

Apparatus

Experiments 1 and 2 All experimental procedures took place in a room with a single testing station, under conditions of low ambient light. All tasks were programmed in Matlab (Mathworks, Natick, MA, 2013a), using the Psychophysics Toolbox extension (Brainard 1997; Pelli 1997). 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 white [RGB: 200, 200, 200] vertical lines (0.5° w x 1° h) 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 (2.2° deg in diameter), one in purple [87, 75, 80], the other in orange [120, 86, 1] (matched for luminance) served as value cues. They were aligned along the horizontal meridian and

positioned 4.5° from the centre. Targets ('H' or 'N') and distractors ['Z' or 'K'] were presented in light grey [90, 90, 90] Helvetica font, encompassed 1°, and were centred on the disc's centre. 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. All stimuli were presented on a grey [RGB: 118, 118, 118] background.

Procedure

As shown in in Figure 1, each trial began with the simultaneous presentation of both incentive value cues and two centrally presented vertical lines. After a pseudo-randomly chosen duration of 400-500 ms, the left or right fixation line darkened for 300 ms. After a further 100 ms, the target and distractor were presented for 100 ms (target identity was equiprobable for each incentive value x spatial certainty x cue validity condition). Participants pressed with the 'v' or 'g' key on a standard keyboard to indicate the target identity. After 500 ms, 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 correct trials to prevent feedback signals becoming redundant. The high and low incentive value cue locations and the target location were equally likely to be on the left or right; all conditions (cue value location, target location, and target identity) were fully crossed within each session. Colour/value pairings (e.g. purple = 50 points/orange = 1 point) as well as target-response mappings were counterbalanced across participants.

Across blocks, the likelihood of cue validity was varied to be either .6 valid/.4 invalid, .8/.2, .9/.1, .92/.08, .94/.06, resulting in information gains (spatial-certainty) of .029, .29, .53, .6 and .86 bits. Each block contained 100 trials. At each of 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 explicitly instructed how many points were available should the target appear in the location of the high and low value incentive value cues (50 vs 1 point), and were instructed that the cues signified that points were available most of the time, but not all of the time.

They were also instructed 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) (R Core Team 2013), and R Studio (v1.0.44) (RStudio Team 2016). RT data were rejected if they were greater than +/- 2.5 s.d. from the mean for that participant in that condition. As 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 specification and selection

The aim of the study was to compare whether an additive 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 incentive value and spatial certainty was more probable, given the data, than one that only included main effects (i.e., an additive model). To quantify evidence, we used a Bayesian approach that provides the advantage of offering the ability to quantify evidence against a specific model, which is not possible using null hypothesis significance testing approaches (Wagenmakers 2007; Nickerson 2000). Additionally, Bayesian approaches protect against the problem of model complexity: although more complex models may predict with high likelihood a greater range of values,

if these predictions are uninformative, this will result in a more diffuse marginal likelihood distribution when integrating across prior distributions for the parameters, thereby penalising the resulting model evidence. First, we fit all possible linear mixed models, with the regressors being (a) incentive value of the target location, (b) cue validity, and (c) spatial certainty offered by the cue. Spatial certainty was computed in line with Prinzmetal et al (Prinzmetal et al. 2015). Specifically, Shannon's (Shannon 1948) measure of entropy (H) measures the amount of uncertainty in a probability distribution and is at maximum when the cue is completely unpredictable with regard to the target location. Therefore, spatial certainty gained by an informative cue can be calculated as:

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

when *H* is defined in the standard manner:

$$H = -\sum_i p_i \log_2 p_i \quad (2)$$

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 (corresponding to complete uncertainty, i.e. .5/.5), then the information gained by the cue is 1 - .72 $\approx$ .28 bits.

After fitting all possible models to the RT and accuracy data obtained in each experiment, we computed Bayes Factors (BFs) to quantify evidence for each linear mixed effects model against the null model (intercept plus random effects of participant) using the Bayes Factor package (Morey, Rouder, and Jamil 2014), and implementing the default Jeffreys-Zeller-Siow (JZS) prior on slope estimates (Liang, Paulo, and Molina 2008). We then identified the six 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 five best models, to demonstrate the strength of evidence in favour of the winning model. We follow the guidelines of Kass and Rafferty (Kass and Raftery 1995) when interpreting the strength of evidence. This was typically sufficient to determine whether the evidence favoured a model that included only main effects, or an incentive value x spatial certainty interaction. However, in

the one case where this was not sufficient (accuracy measures of Experiment 2), further targeted comparisons were also made. All BFs are reported along with the proportional error of the estimate. For readers interested in confirming that an null hypothesis significance testing (NHST) approach yields the same conclusions, please refer to the online repository for this study.

Results

We adapted a spatial-orienting task (Posner 1980) to test whether the influence of spatial certainty and incentive value remains additive under conditions that approach maximal certainty (see Figure 1). In Experiment 1, for both the RT and accuracy data, we find strong evidence is in favour of an additive influence of incentive value and spatial certainty, thereby providing support for the additive hypothesis, against the expected value and mixed operations hypotheses.

RT

Group mean 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 (incentive value, spatial certainty, and cue validity), and a spatial certainty x cue validity interaction term (BF = 1.74E+58, ± .87 %, see Figure 2B). The main effect of incentive value was to speed RT by approximately 50 ms ± .3 (SE) for high versus low incentives. Spatial certainty served to increase the effect of cue validity; the difference between valid and invalid trials increased by approximately 90 ms ± 29 (SE) across 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 an incentive value x spatial certainty interaction term. Therefore, the evidence favours a model that does not include an interaction between incentive value and spatial certainty.

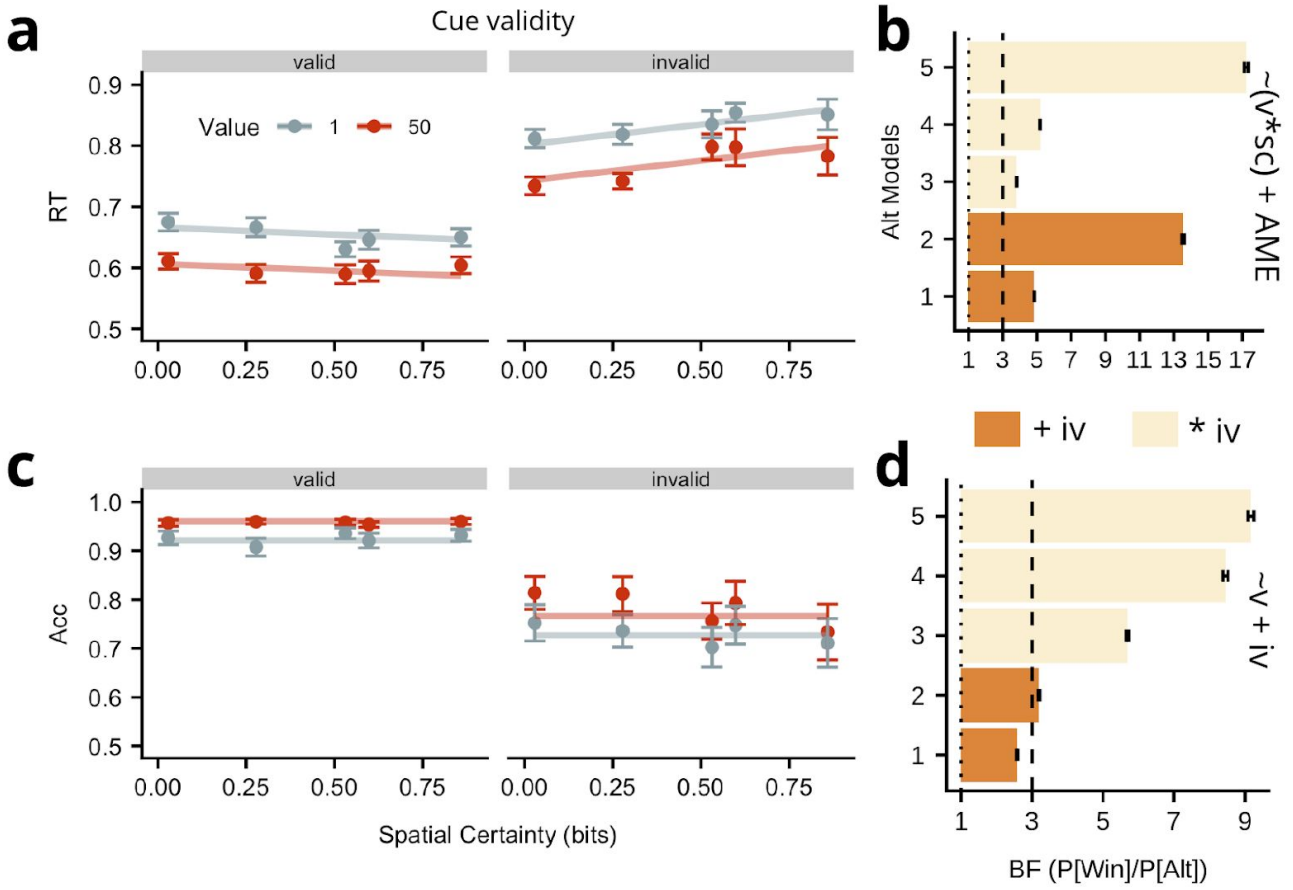

Figure 2: Results from Experiment 1. **a.** Observed group mean RTs in ms for targets appearing in high (dark circles) or low (light circles) value placeholders plotted as a function of spatial certainty (x-axis) for valid and invalid spatial cues (panels). Vertical lines indicate ± 1 within-subject s.e. Lines represent fit of the winning model. The winning RT model (shown on the right of panel b) did not involve any interaction of incentive value (iv) and spatial certainty (sc) supporting an additive hypothesis, although it did indicate an interaction of sc and validity (v). **b.** BFs for the probability of the winning RT model (P[Win]: (v * sc) + AME) relative to the 5 next best models (Alternative, P[Alt], models, y-axis). The larger the BF value, the stronger evidence for the winning model. Any values lower than 1 (dotted line) support P[Alt]. BF values over 3 (dashed line) constitute strong evidence for the winning model. Dark bars indicate P[Alt] contains only an additive influence of incentive value; light bars indicate P[Alt] involves a multiplicative influence of incentive value and either spatial certainty or validity. The Alt RT models were as follows: 1) ~v + iv, 2) ~AME, 3) ~(v * sc) + (sc * iv) + AME, 4) ~(v * sc) + (v * iv) + AME, 5) ~(v * sc) + (v * iv) + (sc * iv) + AME. **c.** Observed group mean accuracy plotted as in panel a. The winning model (shown on the right of panel d) involved only a main effect of incentive value and a main effect of cue validity, again supporting the additive hypothesis. **d.** BFs for the probability of the winning accuracy model (P[Win]: iv + v). Alt models: 1) ~AME, 2) ~(v * sc) + AME, 3) ~(v * iv) + v + iv, 4) ~(sc * iv) + AME, 5) ~(v * sc) + (iv * sc) + AME. RT = response-time, BF = Bayes Factor, v = cue-validity, sc = spatial certainty, iv = incentive value, AME = all main effects.

Accuracy

Accuracy data do not support the possibility that effects were due to a speed accuracy trade-off (see Figure 2c). The preferred model for these data, relative to the null model included only main effects of incentive value and cue validity (BF = 8.83E+47, ± .56 %, relative to the null model, Figure 2d). The probability of an accurate response on invalid trials was .16 ± .02 (SE) lower than on

valid trials, and was .05 ± .02 (SE) less for targets appearing in low versus high incentive value cues. Compared to the preferred model, evidence for the next best model, which additionally included a main effect of spatial certainty, was weak (BF = 2.57 ± .68 %), implying that we cannot rule out an influence of spatial certainty on the accuracy data. The evidence in favour of the preferred model was more positive relative to the third best model, which, akin to the RT data, included a cue validity x spatial certainty interaction (BF = 3.19 ± .78 %). Crucially for the additive hypothesis, there is strong evidence that the model that includes main effects of incentive value and cue validity is preferred to any model that allows these factors to interact (see Figure 2d), therefore any models with an incentive value x spatial certainty interaction were a poorer account for the data than the winning model. Therefore, although we did not reliably detect an influence of spatial certainty on the accuracy data, these results corroborate the notion that incentive value confers an additive influence on visual selection.

Discussion

Experiment 1 shows that models posing an additive influence of incentive value and spatial certainty outperform those allowing an interaction between the two. This goes against both the expected value and the mixed operations hypotheses and suggests that visual selection mechanisms do not integrate incentive value and spatial certainty, even when approaching the limits of certainty. Nevertheless, the second-best model to account for the RT data included an interaction between spatial certainty and incentive value, suggesting that this interaction is not entirely implausible. Perhaps additive effects would fail to provide the best description of the data if another form of appropriate pressure is applied to visual selection. Previous studies show that learning can direct the sampling of sensory information to optimise reward accrual (Drugowitsch et al. 2015; Kiani and Shadlen 2009; Serences 2008). Therefore, a reward structure that favours an expected value operation may be sufficient to modulate the additive influence of incentive value and spatial certainty. The aim of Experiment 2 was to provide this test.

Experiment 2

Experiment 2 sought to test whether a reward contingency optimised for an expected value operation could induce a non-additive influence of incentive value and spatial certainty. This was achieved by comparing two reward conditions: a decay condition and a fixed condition (the latter as used in Experiment 1). In the decay condition, the reward magnitude decremented exponentially after target onset so that faster responses could accrue greater rewards. In this condition, performance based on a multiplicative influence of incentive value and spatial certainty would gain higher reward values than performance motivated by an additive operation. Furthermore, the gains would outweigh the losses that an expected value operation would yield for low value locations (relative to an additive operation, see Figure 1d). 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 fixed reward condition.

Overall and once again, evidence supports the additive hypothesis and is against the expected value or mixed operations hypothesis.

Method

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

In addition to that used for 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. This was used to ensure that eye movement were not contributing to results,

even though participants were clearly instructed not to move their eyes (replicating instructions used in Experiment 1 that were not, however, verified for compliance).

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 (Stankevich and Geng 2014).

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 incentive value. In the decay reward-condition, an exponential decay function (reward value = points*(e^{-4*RT}), RT = Response Time) was applied to each value at target onset. 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 four spatial certainty/reward contingency 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 fixed 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 single line. However,

we found no evidence that cue-likelihood influenced placement of the placeholders (p = .96), and this separate aspect of the study is discussed no further. Participants also completed a BIS/BAS questionnaire (Carver and White 1994) 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.

Results

RT

RT data (Figure 3a) show that the influence of incentive value is additive with cue validity, even under conditions where it is suboptimal for reward accrual, i.e. in the decay reward condition. If the additive relationship between incentive value and spatial certainty is modulated by reward potential, then we would expect to find an interaction between these two factors in the decay reward condition. First, we identified the most likely model given the data. The winning model included main effects of cue validity, incentive value, and reward condition (BF = 2.18E+67 ± .69 %, relative to null model), but did not include an influence of spatial certainty (although see accuracy data). 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 incentive value x cue validity interaction term (BF = 4.65 ± 2.43 %, see Figure 4a). 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 spatial certainty x cue validity 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 incentive

value (RTs were approximately 21 ms ± 17 (SE) slower to low cue-values relative to high cue-values), cue validity (RTs were on average 33 ms ± 12 (SE) slower on invalid trials relative to valid trials), and reward condition (RTs were approximately 54 ms ± 17 (SE) faster in the decay reward-condition than the fixed reward-condition), without a detectable influence of spatial certainty. Collectively, the results show that even when an additive operation is disadvantageous, an additive model is still a better account of the data.

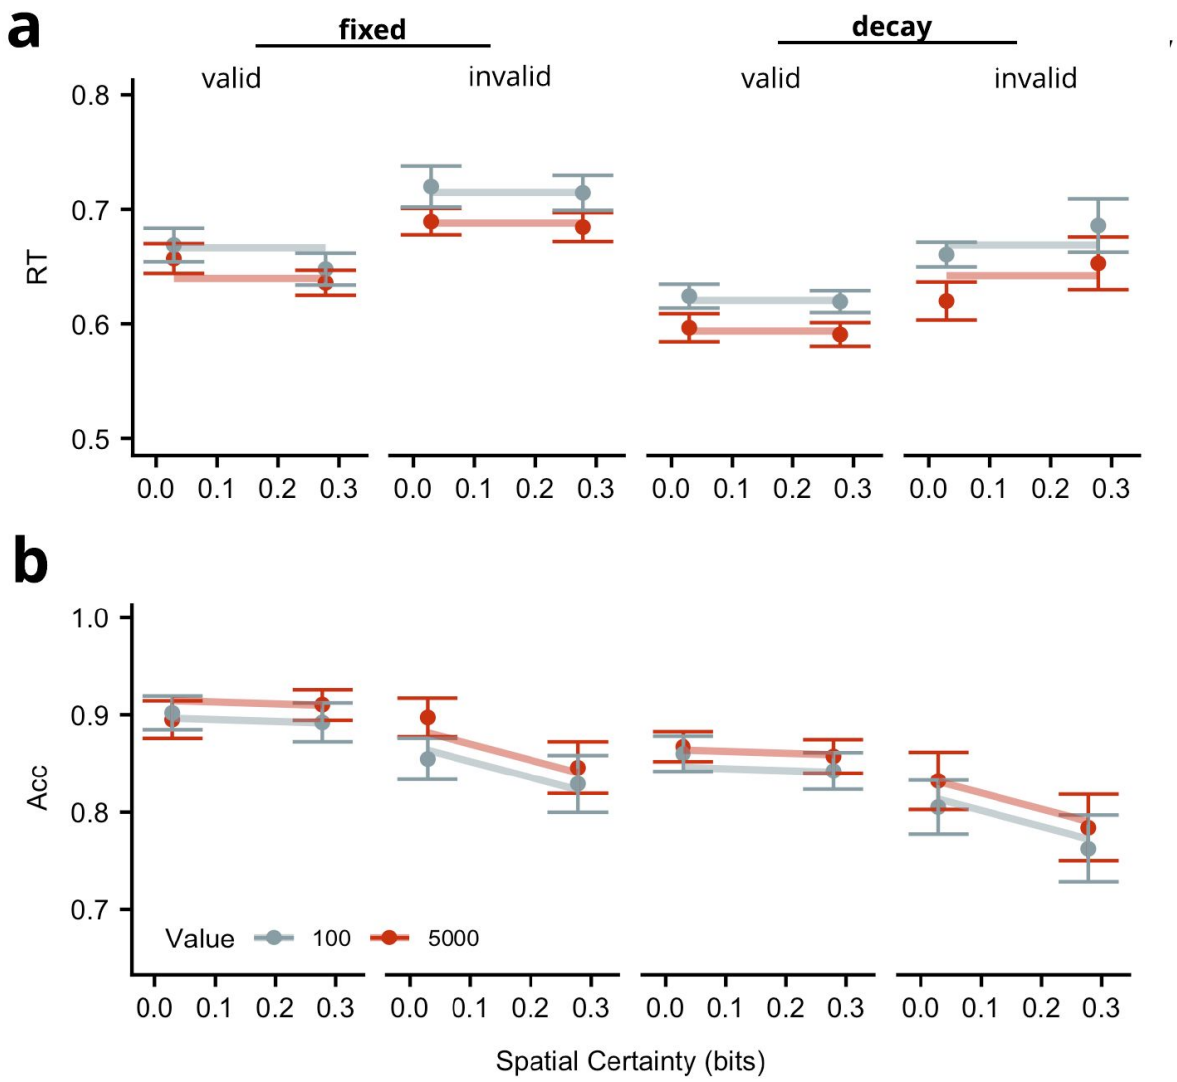

Figure 3: RT and Accuracy for Experiment 2. **a.** Observed group mean RTs in ms for targets appearing in high (dark circles) or low (light circles) value placeholders plotted as a function of spatial certainty (x-axis) for valid and invalid spatial cues (panels) for each reward condition. Vertical lines indicate ± 1 within-subject SE. Lines represent fit of the winning model. The winning model involved main effects of incentive value, cue validity, and reward condition. **b.** Group mean accuracy plotted as in panel a. The winning model (solid lines) was ~(cue validity x spatial certainty) + AME. RT = response-time. Acc = accuracy

Accuracy

Interestingly, and in contrast to the RT data, accuracy scores showed an influence of spatial certainty. However, and importantly for the central question, the models that best accounted for accuracy data did not show an incentive value x spatial certainty interaction (see Figure 3b/4b).

Selecting which model *per se* best accounts for the accuracy data was less simple, as evidence for a single winning model was not conclusive. Consistent with the RT data in Experiment 1, the best model contained a cue-validity x spatial certainty interaction, and main effects of incentive value, spatial certainty and cue validity. Differences in accuracy between valid and invalid trials grew larger as spatial certainty increased (approximately by .15 ± .07 (SE)). Furthermore, accuracy performance was slightly higher when targets appeared in high relative to low incentive value placeholders (by approximately .0003, ± .0001 (SE)). Accuracy was also higher for the fixed 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 incentive value (BF = 1.56, ± 0.75 %). Given this, and to address the key theoretical question, we directly tested whether incentive value interacted with spatial certainty by comparing the winning model to those that additionally included either an incentive value x spatial certainty x reward condition interaction or an incentive value x spatial certainty x cue-validity 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 model.

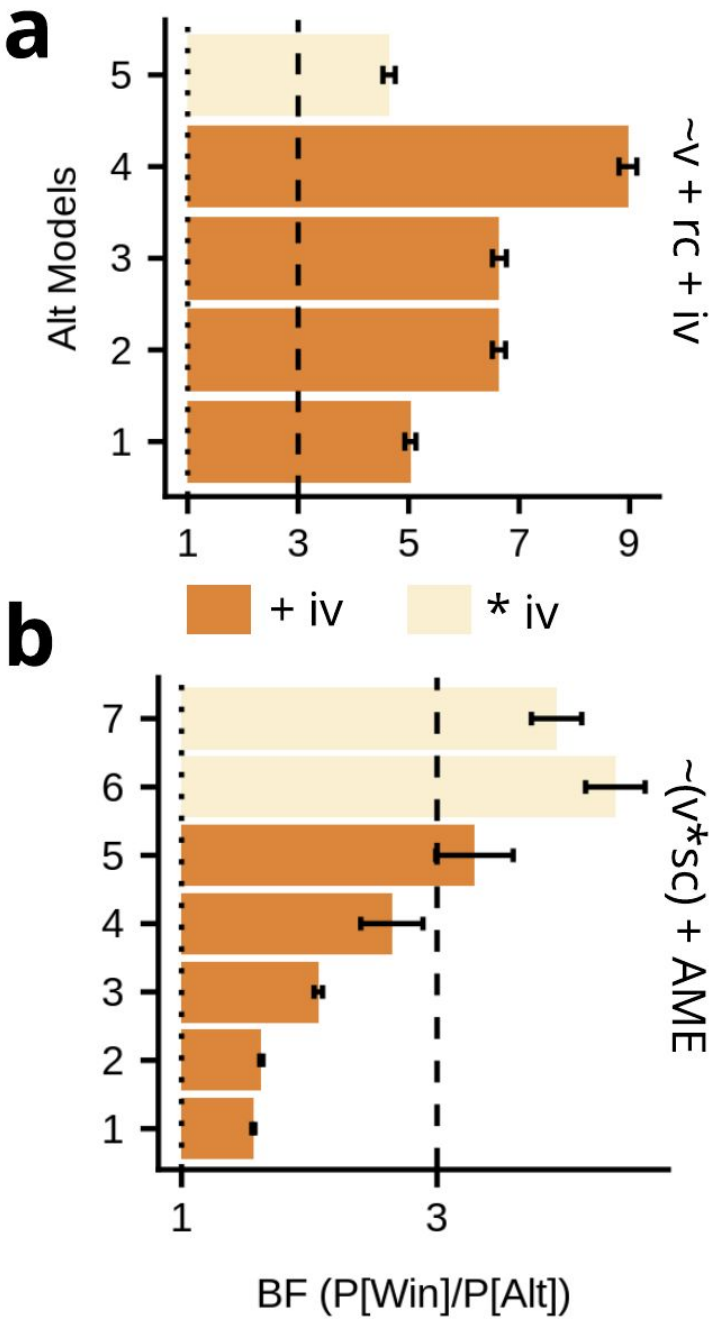

Figure 4: *Results of the Bayes Factor Analysis for Experiment 2*. **a.** BFs for the probability of the winning RT model (P[Win]: v + rc + iv), relative to the 5 next best models (Alternative, P[Alt], models, y-axis). The larger the BF value, the stronger evidence for the winning model. Any values lower than 1 (dotted line) support P[Alt]. BF values over 3 (dashed line) constitute strong evidence for the winning model. Dark bars indicate P[Alt] contains only an additive influence of incentive value; light bars indicate P[Alt] involves a multiplicative influence of incentive value and either spatial certainty or validity. The Alt RT models were as follows:
1) ~(rc x iv) + v + iv + rc, 2) ~AME, 3) ~(rc x v) + iv + v + rc, 4) ~(v x sc) + AME, 5) (v x iv) + v + iv + rc. **b.** BFs for the probability of the winning accuracy model (P[Win]: (v*sc) + AME. Alt models: 1) ~(rc * v * sc) + rc + v + sc, 2) ~AME, 3) ~rc + v + sc, 4) ~(v * sc) + (v * iv) + AME, 5) ~(rc * sc) + (v * sc) + AME, 6) ~(v * sc) + (rc * sc * iv) + AME, 7) ~(v * sc) + (rc * v * sc * iv) + AME. BF = Bayes Factor. valid = cue-validity, sc = spatial certainty, iv = incentive value, rc = reward condition, AME = all main effects.

General Discussion

Over two experiments we tested whether the additive hypothesis would outperform the expected value and mixed operations accounts, even under conditions expected to challenge the optimality of additivity. In Experiment 1, we hypothesised that if incentive value and spatial certainty influence a common underlying mechanism, then conditions wherein spatial certainty is trivial to compute (i.e., very high certainty) might best reveal non-additive effects because these conditions should be minimally taxing to central resources and thus be more likely to enable an influence of incentive value on visual selection. We created this condition by using very high spatial certainty cues and then pitted incentive value and spatial certainty against each other in a spatial-orienting task, where endogenous cues signalled the likely location of upcoming letter targets. Interestingly, an additive influence of incentive value and spatial certainty was observed, even under conditions of very high certainty. Spatial certainty increased the size of the cueing-effect (i.e. the difference in performance between invalidly and validly cued trials), whereas incentive value had a comparable influence on both valid and invalid trial types.

In Experiment 2, we reasoned that if an expected value operation can bias visual selection, then a reward structure that favours a multiplicative weighting of incentive value and spatial certainty may reveal it. We applied an exponential decay function to incentive values at target onset; this ensured that if RTs were driven by a multiplicative weighting rather than an additive weighting, then reward gains accrued by faster RTs under high incentive value/certainty conditions would outweigh the losses incurred by the slower RTs under low incentive value/certainty conditions, relative to RTs predicted by an additive model. Although the influence of spatial certainty was manifest differently than in Experiment 1, i.e., by modulating accuracy, rather than RT, we observed that the effect of incentive value remained additive to spatial certainty and to other experimental factors. Again, the findings support the additive hypothesis.

What kind of mechanism or computation could result in a robust additive influence between incentive value and spatial certainty? In concert with recent theoretical and empirical developments suggesting that cognitive control processes are offset by subjective and computational costs of effortful control (Braver 2012; Yee and Braver 2018), we believe the current data can be interpreted

as reflecting trial by trial adaptations aimed at the conservation of effort. If we assume that the maintenance of the task set, i.e., *a priori* preparedness to identify a data-limited target at two locations, requires energetic resources from the underlying selection mechanism, it is of benefit to the brain to predict conditions where effort can be relaxed, in order to conserve energy expenditure. For example, by learning the energetic range over which target identification mechanisms can be adjusted, to ensure good enough target detection, given the task parameters. Acoording to this view, a cost-benefit analysis could inform how much energy could be saved, given an acceptable decrement to accuracy and response time.

Applied to current context, after the onset of the incentive value cues, selection mechanisms should maintain a steady level of task preparation favouring the high value location for example, by increasing the excitability of neuronal assemblies whose collective receptive fields correspond to detecting lines or edges at that location (Desimone and Duncan 1995; Roelfsema, Lamme, and Spekreijse 2000; Carrasco 2011; Schmitz and Duncan 2018), thereby biasing the system towards a stronger response to the upcoming stimulus (Buschman and Kastner 2015). This presumably allows for more rapid detection of the stimulus at that location, and a consequent reduction in the period during which sensory evidence needs to be evaluated to identify the letter at that location. Concurrently, the excitability of neuronal assemblies directed towards encoding information from the low value location should be relaxed, as the cost of sometimes missing the target at that location, given the energy needed to detect it, should become negligible. Similarly, upon spatial cue onset, preparation of such target detection mechanisms could be further relaxed for the unlikely location, proportionally to how unlikely that location is to possess a target. Importantly, this could be performed incrementally to the previous adjustment. This would ensure that the system is most ready to encode stronger representations of items with higher value and higher certainty, while sacrificing the representation strength of items imbued with lesser value and lesser certainty to minimize costs. This interpretation predicts that the degree to which incentive value or spatial certainty can influence performance is dependent upon the range over which preparatory processes can be titrated and still yield acceptable performances. For example, reducing the duration of data-limited target presentation should likewise reduce the influence of spatial certainty and value, as there will be a lower range

within which the energetic effort expended on the cognitive operation can be titrated before performance starts dropping unacceptably. This would account for why, in Experiment 2, under a context with greater pressure on RT performance, we observed an influence of cue validity but not spatial certainty. Presumably it had become too costly to modulate RT performance by spatial certainty and meet the perceived demands of the task.

The current findings shed further insights into the Stănişor et al.'s (2013) proposed unitary selection mechanism that biases competition between visual cortical representations of stimuli, in the presence of both incentive and explicit spatial cues. To recap they showed that overlapping clusters of V1 neurons were sensitive to both incentive value cues and 100 % informative explicit spatial cues. The authors proposed that the explicit spatial cue served to reweight the relative value between the target and the distractor, and that this reweighting was instantiated by a unitary selection mechanism. It remained unknown whether the information offered by the spatial cue overwrote, or was integrated with the relative incentive value. The current study indicates that the spatial certainty offered by the explicit cue does not overwrite incentive value, as we revealed an additive relationship between the two. Furthermore, this additive influence shows that the spatial certainty derived from explicit cues is not entered into an expected value operation to reweight the relative value between the two items. Rather, an additive influence points to the repeated invocation of the selection mechanism, for example, maintaining or decreasing excitability for neuronal assemblies with receptor fields over the target area (Desimone and Duncan 1995; Roelfsema, Lamme, and Spekreijse 2000; Schmitz and Duncan 2018), on the basis of updates from separable information sources. However, as incentive value and spatial certainty have been added, rather than multiplied, this suggests that the two have been transformed into a common representational space, or unit, prior to influence on the visual selection mechanism.

An additive influence of incentive value on visual selection was also observed by Stankevich and Geng (2014), when value was pitted against the varying probability that a target would appear on one side versus the other, in the absence of explicit spatial cues. Interestingly, what our results show is that even when explicit information is available on each trial regarding the likely location of an upcoming target, the influence on visual selection is independent to that provided by incentive value

information. Following the line of reasoning above, this suggests that the information offered by explicit cues may be encoded into the same representational space, or units, as incentive value, just as occurs for non-cued spatial information that has been acquired over multiple exposures. This suggestion is supported by the observation that explicit and non-cued spatial knowledge also share an additive influence on visual selection (Geng and Behrmann 2005). Thus these studies suggest that regardless of how associations are signalled (i.e. explicitly or non-explicitly), the information offered by them are converted into a common representational space, or unit, to be summed over to exert control over visual selection.

A visual comparison of the current data and the data from Stankevich and Geng (2014) also yields some interesting points of difference concerning the influence of spatial certainty in the presence or absence of explicit spatial cues. With the current explicit cues, we observe RTs that are comparable across spatial certainty conditions for valid cues, whereas RTs on invalid trials increase as spatial certainty decreases. In contrast, Stankevich and Geng (2014) observed decreasing RTs towards more likely locations, and stable RTs towards less likely locations. Moreover, we observe a large main effect of the spatial cue, even in low certainty conditions, whereas Stankevich and Geng (2014) found small RT differences between valid and invalid locations when uncertainty was high (i.e. in a p = .6/.4 condition). Therefore, our results suggest that the explicit spatial cue we used resulted in preparation towards the cued location that did not vary with the certainty offered by the cue, coupled with a relaxation of preparation towards the invalid location that scaled with certainty. In contrast, non-explicit spatial knowledge appears to cause a strengthening of preparation towards the more likely location, with no concomitant relaxation towards the unlikely location. This suggests that spatial certainty influences visual selection differently dependent on how it is learned. This may reflect a contextual selection of the most salient behaviour as a baseline response. For example, a long history of arrows serving as useful directional cues could motivate a strong response to the directional stimulus, against which other useful behaviours can be adapted. In contrast, the absence of any explicit cues could result in a conservative preparation towards possible locations that is adapted or strengthened with exposure to spatial contingencies, relative to improbable locations. In either case, it

appears that visual selection behaviours adapt to environmental information in relation to the most contextually relevant baseline behaviour.

The current finding of additivity across learned-associations also accords with computational (Itti and Koch 2001) and neurophysiological (Deco and Rolls 2005) models of visual selection. Although these previous 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 also applies to the various associations made between physical stimuli and their consequent outcomes. When the current findings are interpreted within a feature additivity framework, they suggest that at a computational level, the visual system makes use of learned associations in the same way as it uses diverse stimulus feature information. Perhaps what unites feature dimensions and learned associations is that they jointly account for unique variance in the visual scene that is relevant for the current task-set. Thus, visual selection mechanisms could make use of an additive weighting of components derived from a dimension reduction of incoming sensory data combined with associated internal data. This suggests that activity from neurons responding to features is combined with the activity from neurons responding to associations for normalisation into a common space.

Our results are in apparent contradiction to previous work showing that saccadic onset latencies correlated with expected-value (Milstein and Dorris 2007). There are at least two possible reasons for this difference. First, the predictions made by the additive and expected value hypotheses are similar, so it may be that the additive hypothesis can still provide a better model for the saccadic onset times observed by Milstein and Dorris (2007), and that either the additive or expected value model would perform better than one based on incentive value or spatial certainty alone (as was the comparison made by (Milstein and Dorris 2007). A good first test would be to apply the current analysis to their saccadic response time data to directly pit the additive and expected value models in this context. Second, it may be that expected value computations are leveraged to influence visual selection, but that the computations are specific to single explicit cues in the environment. If a single cue, such as an arrow, were systematically varied to signal different potential reward outcomes with

varying degrees of certainty, then perhaps a multiplicative relationship between incentive value and certainty could be observed. Future experiments should disambiguate between these possibilities.

Conclusions

Over two experiments, we sought to arbitrate between competing theories for how learned associations pertaining to incentive value and spatial certainty combine to influence visual selection. Specifically, we asked whether this influence was additive (*additive hypothesis*), multiplicative (*expected value hypothesis*) or both (*mixed operations hypothesis*). We tested these hypotheses by pitting incentive values and spatial certainties against one another in a spatial cueing task under conditions expected to challenge the optimality of an additive operation. The data from two experiments support the notion that visual selection mechanisms show independent sensitivity to incentive value and spatial certainty information, and that both information sources are converted to a common representational space, or unit, in order to influence visual selection. We interpret our results as suggesting that the mechanism leveraging visual selection dynamically leverages distinct information sources to reflexively conserve effort within a range that allows acceptable performance given the task parameters. We also 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.

References

Anderson, Brian A., Patryk A. Laurent, and Steven Yantis. 2011. “Value-Driven Attentional Capture.” *Proceedings of the National Academy of Sciences of the United States of America* 108 (25): 10367–71.

Awh, Edward, Artem V. Belopolsky, and Jan Theeuwes. 2012. “Top-down versus Bottom-up Attentional Control: A Failed Theoretical Dichotomy.” *Trends in Cognitive Sciences* 16 (8): 437–43.

Brainard, D. H. 1997. “The Psychophysics Toolbox.” *Spatial Vision* 10 (4): 433–36.

Braver, Todd S. 2012. “The Variable Nature of Cognitive Control: A Dual Mechanisms Framework.” *Trends in Cognitive Sciences* 16 (2): 106–13.

Buschman, Timothy J., and Sabine Kastner. 2015. “From Behavior to Neural Dynamics: An Integrated Theory of Attention.” *Neuron* 88 (1): 127–44.

Button, Katherine S., John P. A. Ioannidis, Claire Mokrysz, Brian A. Nosek, Jonathan Flint, Emma S. J. Robinson, and Marcus R. Munafò. 2013. “Power Failure: Why Small Sample Size Undermines the Reliability of Neuroscience.” *Nature Reviews. Neuroscience* 14 (5): 365–76.

Carrasco, Marisa. 2011. “Visual Attention: The Past 25 Years.” *Vision Research* 51 (13): 1484–1525.

Carver, Charles S., and Teri L. White. 1994. “Behavioral Inhibition, Behavioral Activation, and Affective Responses to Impending Reward and Punishment: The BIS/BAS Scales.” *Journal of*

Personality and Social Psychology 67 (2): 319.

Chelazzi, Leonardo, Jana Eštočinová, Riccardo Calletti, Emanuele Lo Gerfo, Ilaria Sani, Chiara Della Libera, and Elisa Santandrea. 2014. "Altering Spatial Priority Maps via Reward-Based Learning." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 34 (25): 8594–8604.

Deco, Gustavo, and Edmund T. Rolls. 2005. "Neurodynamics of Biased Competition and Cooperation for Attention: A Model with Spiking Neurons." *Journal of Neurophysiology* 94 (1): 295–313.

Desimone, R., and J. Duncan. 1995. "Neural Mechanisms of Selective Visual Attention." *Annual Review of Neuroscience* 18: 193–222.

Drugowitsch, Jan, Gregory C. DeAngelis, Dora E. Angelaki, and Alexandre Pouget. 2015. "Tuning the Speed-Accuracy Trade-off to Maximize Reward Rate in Multisensory Decision-Making." *eLife* 4 (June): e06678.

Geng, Joy J., and Marlene Behrmann. 2005. "Spatial Probability as an Attentional Cue in Visual Search." *Perception & Psychophysics* 67 (7): 1252–68.

Heijden, A. H. van der. 1989. "Probability Matching in Visual Selective Attention." *Canadian Journal of Psychology* 43 (1): 45–52.

Hick, W. E. 1952. "On the Rate of Gain of Information." *The Quarterly Journal of Experimental Psychology* 4 (1): 11–26.

Hutchinson, J. Benjamin, and Nicholas B. Turk-Browne. 2012. "Memory-Guided Attention: Control from Multiple Memory Systems." *Trends in Cognitive Sciences* 16 (12): 576–79.

Hyman, R. 1953. "Stimulus Information as a Determinant of Reaction Time." *Journal of Experimental Psychology* 45 (3): 188–96.

Itti, L., and C. Koch. 2001. "Computational Modelling of Visual Attention." *Nature Reviews. Neuroscience* 2 (3): 194–203.

Jonides, J. 1980. "Towards a Model of the Mind's Eye's Movement." *Canadian Journal of Psychology* 34 (2): 103–12.

Kass, Robert E., and Adrian E. Raftery. 1995. "Bayes Factors." *Journal of the American Statistical Association* 90 (430): 773–95.

Kiani, Roozbeh, and Michael N. Shadlen. 2009. "Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex." *Science* 324 (5928): 759–64.

Kingstone, Alan. 1992. "Combining Expectancies." *The Quarterly Journal of Experimental Psychology Section A* 44 (1): 69–104.

Kiss, Monika, Jon Driver, and Martin Eimer. 2009. "Reward Priority of Visual Target Singletons Modulates Event-Related Potential Signatures of Attentional Selection." *Psychological Science* 20 (2): 245–51.

Lanthier, Sophie N., David W-L Wu, Craig S. Chapman, and Alan Kingstone. 2015. "Resolving the Controversy of the Proportion Validity Effect: Volitional Attention Is Not Required, but May Have an Effect." *Attention, Perception & Psychophysics* 77 (8): 2611–21.

Le Pelley, Mike E., Chris J. Mitchell, Tom Beesley, David N. George, and Andy J. Wills. 2016. "Attention and Associative Learning in Humans: An Integrative Review." *Psychological Bulletin*, August. https://doi.org/10.1037/bul0000064.

Le Pelley, Mike E., Daniel Pearson, Oren Griffiths, and Tom Beesley. 2015. "When Goals Conflict with Values: Counterproductive Attentional and Oculomotor Capture by Reward-Related Stimuli." *Journal of Experimental Psychology. General* 144 (1): 158–71.

Liang, F., R. Paulo, and G. Molina. 2008. "Mixtures of G Priors for Bayesian Variable Selection." *Journal of the American Statistical Association* 103 (481): 410–23.

Lorca-Puls, Diego L., Andrea Gajardo-Vidal, Jitrachote White, Mohamed L. Seghier, Alexander P. Leff, David W. Green, Jenny T. Crinion, et al. 2018. "The Impact of Sample Size on the Reproducibility of Voxel-Based Lesion-Deficit Mappings." *Neuropsychologia* 115 (July): 101–11.

Milstein, David M., and Michael C. Dorris. 2007. "The Influence of Expected Value on Saccadic Preparation." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 27 (18): 4810–18.

Moore, Tirin, and Marc Zirnsak. 2017. "Neural Mechanisms of Selective Visual Attention." *Annual Review of Psychology* 68 (January): 47–72.

Morey, R. D., J. N. Rouder, and T. Jamil. 2014. “BayesFactor: Computation of Bayes Factors for Common Designs.” *R Package Version 0. 9* 8.

Morgenstern, Oskar, and John Von Neumann. 1953. *Theory of Games and Economic Behavior*. Princeton university press.

Nickerson, R. S. 2000. “Null Hypothesis Significance Testing: A Review of an Old and Continuing Controversy.” *Psychological Methods* 5 (2): 241–301.

Pelli, D. G. 1997. “The VideoToolbox Software for Visual Psychophysics: Transforming Numbers into Movies.” *Spatial Vision* 10 (4): 437–42.

Pessoa, Luiz, and Jan B. Engelmann. 2010. “Embedding Reward Signals into Perception and Cognition.” *Frontiers in Neuroscience* 4 (September). https://doi.org/10.3389/fnins.2010.00017.

Posner, M. I. 1980. “Orienting of Attention.” *The Quarterly Journal of Experimental Psychology* 32 (1): 3–25.

Posner, M. I., C. R. Snyder, and B. J. Davidson. 1980. “Attention and the Detection of Signals.” *Journal of Experimental Psychology* 109 (2): 160–74.

Prinzmetal, William, Kelly L. Whiteford, Joseph L. Austerweil, and Ayelet N. Landau. 2015. “Spatial Attention and Environmental Information.” *Journal of Experimental Psychology. Human Perception and Performance* 41 (5): 1396–1408.

Raymond, Jane E., and Jennifer L. O’Brien. 2009. “Selective Visual Attention and Motivation.” *Psychological Science*, August. https://doi.org/10.1111/j.1467-9280.2009.02391.x.

R Core Team. 2013. “The R Project for Statistical Computing.” *Available at Www. R-Project. Org/. Accessed October*.

Roelfsema, Pieter R., Victor A. F. Lamme, and Henk Spekreijse. 2000. “The Implementation of Visual Routines.” *Vision Research* 40 (10): 1385–1411.

RStudio Team. 2016. “RStudio: Integrated Development Environment for R.” Boston, MA: RStudio, Inc. http://www.rstudio.com/.

Sawaki, Risa, Steven J. Luck, and Jane E. Raymond. 2015. “How Attention Changes in Response to Incentives.” *Journal of Cognitive Neuroscience* 27 (11): 2229–39.

Schmitz, Taylor W., and John Duncan. 2018. “Normalization and the Cholinergic Microcircuit: A Unified Basis for Attention.” *Trends in Cognitive Sciences* 22 (5): 422–37.

Serences, John T. 2008. “Value-Based Modulations in Human Visual Cortex.” *Neuron* 60 (6): 1169–81.

Sha, Li Z., and Yuhong V. Jiang. 2016. “Components of Reward-Driven Attentional Capture.” *Attention, Perception & Psychophysics* 78 (2): 403–14.

Shannon, C. E. 1948. “A Mathematical Theory of Communication.” *Bell System Technical Journal* 27 (3): 379–423.

Small, Dana M., Darren Gitelman, Katharine Simmons, Suzanne M. Bloise, Todd Parrish, and M-Marsel Mesulam. 2005. “Monetary Incentives Enhance Processing in Brain Regions Mediating Top-down Control of Attention.” *Cerebral Cortex* 15 (12): 1855–65.

Stănişor, Liviu, Chris van der Togt, Cyriel M. A. Pennartz, and Pieter R. Roelfsema. 2013. “A Unified Selection Signal for Attention and Reward in Primary Visual Cortex.” *Proceedings of the National Academy of Sciences of the United States of America* 110 (22): 9136–41.

Stankevich, Beth A., and Joy J. Geng. 2014. “Reward Associations and Spatial Probabilities Produce Additive Effects on Attentional Selection.” *Attention, Perception & Psychophysics* 76 (8): 2315–25.

Sternberg, Saul. 1998. *Discovering Mental Processing Stages: The Method of Additive Factors*. The MIT Press.

Wagenmakers, Eric-Jan. 2007. “A Practical Solution to the Pervasive Problems Ofp Values.” *Psychonomic Bulletin & Review* 14 (5): 779–804.

Yee, Debbie M., and Todd S. Braver. 2018. “Interactions of Motivation and Cognitive Control.” *Current Opinion in Behavioral Sciences* 19 (February): 83–90.