

Towards an Understanding of How People Build Mental Representations

by

Derick Valadao

A thesis

presented to the University of Waterloo

in fulfillment of the

thesis requirement for the degree of

Doctor of Philosophy

in

Psychology

Waterloo, Ontario, Canada, 2016

© Derick Valadao 2016

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Navigating the environment and making everyday decisions is a process plagued by noise, uncertainty, and non-stationary contingencies. Efficient and effective action is predicated upon a stable internal representation of the environment that guides action without extensive or exhaustive observation, deliberation, and alteration at the slightest deviation from expected outcomes. The ability of individuals to build these mental models and update them as needed represents a critical component of everyday decision and action. The current thesis provides an in-depth exploration of this construct through a series of brain imaging and behavioural experiments examining the neural correlates of mental model building and updating focusing on how other cognitive abilities (i.e., working memory and attention) influence the speed and accuracy of these processes. Brain imaging results highlight a network of frontal, parietal, and subcortical areas that support mental model updating. Follow-up behavioural experiments reveal both working memory and attention to be important gating mechanisms to the processing of environmental stimuli that comprise a mental model. Taken together, the results point to a robust neural network coupled with working memory and attentional gating mechanisms that support this behaviour.

Acknowledgments

Firstly, I would like to thank my supervisor, Dr. James Danckert, who has been steadfast in his guidance, mentorship, and support of my academic and professional development. I owe so much to the trust he placed in me early on by giving me the opportunity to develop valuable research experience and collaborate with lab members. His patience and perspective also helped me cope with the highs and lows of my graduate training. I'd also like to thank my thesis committee, Drs. Tara McAuley, Britt Anderson, Michael Barnett-Cowan, and Daniel Meegan, for their time and thoughtful consideration of my work. I'm particularly grateful to the generosity of both Drs. Anderson and McAuley in providing me with counsel and encouragement throughout my research and clinical pursuits.

Thank you to all the members of the Danckert lab, both past and present, who at various times provided me with guidance, support, and camaraderie. I'd particularly like to thank Dr. Marc Hurwitz for his support and mentorship as I transitioned to graduate studies, Dr. Yael Goldberg for the advice and encouragement proffered throughout my clinical training, Alex Filipowicz for being such a talented collaborator and friend, Julia Isacescu for her friendship and support, Dr. Linda Carson for her sage advice and encouragement, and Nadine Quehl for supporting us and cultivating a community out of a disparate set of individuals who otherwise would have passed each other like ships in the night.

I'd also like to thank the two members of my clinical cohort, Ivana Lizdek and Anisha Varghese, who I am proud to call friends. I know that we have all leaned on each other at certain times as we traversed this long, difficult, and rewarding road together. Being able to work and develop alongside you, support each other through loss and hardship, and celebrate each other's victories has been an absolute privilege.

I'd like to thank my family. Thank you to my parents and brother for the years of unconditional love, support, and encouragement. Thank you to my grandparents, who gave me the courage to make my own leap of faith. Finally, a special thanks to my wife, Gloria. Thank you for believing in me when I couldn't, for the patience you've shown, and the love and support you've given me on this journey.

Table of Contents

List of Tables	ix
List of Figures.....	x
Chapter 1: Introduction	1
Chapter 2: Neural Correlates of Updating¹.....	4
2.1. Introduction	4
2.2. Method.....	10
Participants.....	10
Design	11
fMRI data collection.....	17
fMRI Data analysis.....	18
Behavioural Data Analysis.....	19
2.3. Results.....	21
2.4. Discussion	32
Evidence of modeling bias in environmental regularity.....	32
Contributions of right parietal cortex to evidence seeking and integration.....	33
Cingulate cortex may signal the need to update.....	34
Involvement of prefrontal cortex in model maintenance	36
Role of insular cortex in balancing risk and reward.....	37
General conclusion.....	38
Chapter 3: Working memory correlates of updating².....	41
3.1. Introduction	41
3.2. Experiment 1: Method	43
Participants.....	43
Apparatus	44
Procedure.....	44
Data Analysis	52

3.3. Experiment 1: Results	53
3.4. Experiment 2: Introduction	57
3.5. Experiment 2: Method	57
Participants	57
Procedure.....	58
3.6. Experiment 2: Results	58
3.7. Experiment 3: Introduction	60
3.8. Experiment 3: Method	61
Participants	61
3.9. Experiment 3: Results	61
3.10. Discussion	63
Chapter 4: Attentional Influences on Building and Updating Mental Representations	66
4.1. Introduction	66
4.2. Experiment 1: Introduction	69
4.3. Experiment 1: Method	69
Participants	69
Apparatus	70
Procedure.....	70
Data Analysis	74
4.4. Experiment 1: Results	76
4.5. Experiment 1: Discussion.....	79
4.6. Experiment 2: Introduction	79
4.7. Experiment 2: Method	80
Participants	80
Procedure.....	80
4.8. Experiment 2: Results	83

4.9. Experiment 2: Discussion.....	85
4.10. Experiment 3: Introduction.....	86
4.11. Experiment 3: Method	86
Participants.....	86
Procedure.....	87
4.12. Experiment 3: Results	89
4.13. Experiment 3: Discussion.....	94
4.14. Experiment 4: Introduction.....	94
4.15. Experiment 4: Method	95
Participants.....	95
Procedure.....	95
4.16. Experiment 4: Results	98
4.17. Experiment 4: Discussion.....	101
4.18. Discussion	101
Chapter 5: General Discussion	106
References	111

List of Tables

Table 2.1. Curve Estimation Statistics.....	21
Table 2.2. Whole-brain GLM Activations for computed contrasts	25
Table 3.1. Inverse Learning Curve Estimation Statistics.....	56

List of Figures

Figure 2.1. Rock Paper Scissors Stimuli.....	12
Figure 2.2. fMRI Run Conditions.	14
Figure 2.3. Rock Paper Scissors Behavioural Data.	22
Figure 2.4. Whole-brain GLM Contrast Volume Maps.....	28
Figure 3.1. Working Memory Prediction Stimuli.....	46
Figure 3.2. 2-Back Task.....	49
Figure 3.3. N-Back Task Accuracy.....	51
Figure 3.4. Experiment 1 Location Prediction Data.	54
Figure 3.5. Experiment 2 Shape Prediction Data.....	59
Figure 3.6. Experiment 3 Shape Prediction Data.....	62
Figure 4.1. Original Plinko Task.....	71
Figure 4.2. Experiment 1 Plinko Task.	73
Figure 4.3. Accuracy Estimation Example.	75
Figure 4.4. Experiment 1 Estimation Accuracy.....	77
Figure 4.5. Experiment 1 Estimation Accuracy by Condition.....	78
Figure 4.6. Simultaneous Plinko Task.	82
Figure 4.7. Experiment 2 Estimation Accuracy.....	83
Figure 4.8. Experiment 2 Accuracy by Condition.	84
Figure 4.9. Experiment 2 Overlapping Distribution Accuracy.....	85
Figure 4.10. Plinko SART Task.....	88
Figure 4.11. Experiment 3 SART Task Accuracy.	90
Figure 4.12. Experiment 3 Estimation Accuracy.....	91

Figure 4.13. Estimation Accuracy by Condition.	92
Figure 4.14. Estimation Accuracy for Wide vs. Narrow Distribution.	93
Figure 4.15. Experiment 4 Distribution Pairs.	97
Figure 4.16. Experiment 4 Estimation Accuracy by Distribution Width.	99
Figure 4.17. Experiment 4 Wide vs. Narrow Estimation Accuracy by Attention.	100
Figure 4.18. Experiments 1-3 Combined Attended vs. Unattended vs. Control.	104

Chapter 1: Introduction

Interacting with a complex and dynamic environment requires an ability to represent the environment and flexibility to adjust to changes amidst uncertainty and noise. Ideally, such a process would allow individuals to make optimal predictions about their environment and the consequences of their actions without excessive observation and deliberation. Throughout this thesis, these processes are collectively referred to as mental model building and updating.

Extensive research has proposed numerous underlying mechanisms that support mental model building and updating, from heuristics (Tversky & Kahneman, 1974; Gaissmaier & Schooler, 2008), statistical learning (Saffran, Johnson, Aslin, & Newport, 1999; Fiser & Aslin, 2001; 2002) to Bayesian cognition (Yu, 2007; Nassar, Wilson, Heasly, & Gold, 2010). Common among these is the axiom that individuals seek an expedient understanding of their environment through intuitive reasoning and observation for the purpose of decision-making and action (Griffiths & Tenenbaum, 2006; Tenenbaum, Kemp, Griffiths, & Goodman, 2011, Johnson-Laird, 2013).

Mental models, as defined here, are representations of the regularities in the environment that form a cognitive construct to inform decision making and action that facilitates predictions concerning the outcomes of those decisions/actions (i.e., “What will happen next?”; e.g., Danckert, Stöttinger, Quehl, & Anderson, 2012). For example, one could have a mental model of how long it takes to get to work based on repeated timings of their commute factoring in traffic patterns, distance, various routes, and ancillary tasks that need to be completed upon arrival (e.g., paying for parking, walking to the office from the parking lot). A considerable amount of variables and information gets condensed into a representation of time and variance that allow us to work backwards to ensure we rarely arrive late despite the confluence of factors that could otherwise conspire to delay us. There are numerous examples of such feedforward predictive

models in the literature. Anything from playing competitive games to finding the way to the bathroom in the middle of the night without stubbing your toe, require individuals to model their environment to make optimal predictions and motivate efficient action (Wolpert, 2007; Wolbers, Hegarty, Büchel, & Loomis, 2008; Zhu, Mathewson, & Hsu, 2012; Kwon & Knill, 2013; Stöttinger, Filipowicz, Danckert, & Anderson, 2014).

While mental models are useful for making predictions and guiding efficient action, their utility is contingent on their ability to adapt when circumstances change or feedback demonstrates that the model needs to be updated (Danckert et al., 2012, Stöttinger et al., 2014, Valadao, Anderson, & Danckert, 2015). Using an earlier example, should the city decide to embark on an ambitious construction project in the middle of the winter, the internal model of how long it takes to get to work may change drastically even though the individual lives just as close to their job as they once did. Failure to update in this instance would lead to potentially costly delays. For successful updating to take place, one would first need to detect a mismatch between expected and observed outcomes in the environment, thereby triggering re-evaluation and accumulation of evidence in the pursuit of a newly formulated model (d'Acremont, Fornari, & Bossaerts, 2013a; 2013b). There exists ample evidence in the literature suggesting that healthy individuals are capable of making such shifts as environmental contingencies change (Brown & Steyvers, 2009; Berniker, Voss, & Kording, 2010; Nassar, Wilson, Heasly, & Gold, 2010; Green, Benson, Kersten, & Schrater, 2010) in addition to indications that diffuse brain damage impairs this ability (Danckert et al., 2012).

The current thesis examines how individuals build and update models of their environment by examining the neural and cognitive bases for how the process takes place. Chapter 2 examined the neural networks underlying model updating by having individuals play a

serial decision making game against a computer opponent using a variety of often-exploitable strategies, while being scanned in an fMRI scanner. Chapters 3 and 4 explored the influence of working memory (Chapter 3; WM; Baddeley, 1992, 1993) and attention (Chapter 4; Posner, Snyder, & Davidson, 1980) on how well individuals are able to build and update simple models of their environment. These chapters contribute to a broader literature implicating these abilities in the process of learning from regularity in the environment (Nissen & Bullemer, 1987; Corbetta & Shulman, 2002; Toro, Sinnet, and Soto-Faraco, 2005; Turk-Browne, Jungé, and Scholl, 2005; Unsworth & Engle, 2005; de Gardelle & Summerfield, 2011; Albrecht, Scholl, & Chun, 2012; Annac, Manginelli, Pollmann, Shi, Müller, & Geyer, 2013; Manginelli, Langer, Klose, & Pollmann, 2013; Zhao, Al-Aidroos, & Turk-Browne, 2013; Kabata, Yokoyama, Noguchi, and Kita, 2014).

Results in the following chapters suggest the presence of a neural network consisting of frontal, parietal, and subcortical areas that include the anterior insula and cingulate cortex that play a role in supporting the ability of individuals to mentally represent and exploit regularity in their environment. In addition, behavioural data highlight the importance of unconstrained WM and attentional resources in supporting this ability. The final chapter provides a summary of key findings across the current thesis and a discussion of limitations and suggestions for future research.

Chapter 2: Neural Correlates of Updating¹

2.1. Introduction

Interactions with our environment are characterized by uncertainty, noisy input, distractions, and non-stationary contingencies. As such, behaviourally and cognitively efficient interactions require stable and robust internal representations of those aspects of the environment most relevant to a desired goal. Such representations – or ‘mental models’ – allow us to make optimal decisions without having to waste precious time and resources considering all possible outcomes or oversampling the environment (Tenenbaum, Kemp, Griffiths, & Goodman, 2011). These mental models reflect the distillation of relevant sources of information and observed outcomes into an actionable hypothesis space that is capable of being learned, evaluated in real time, and aids in imposing structure over our chaotic environment (Johnson-Laird, 2013). Additionally, as the environment changes, a critical component of mental models is that they are capable of updating to ensure optimal behavior under mutable contexts.

Mental models, and their applications, have been demonstrated in numerous domains. As evidenced from the action and perception literature, models of space facilitate way-finding in the absence or degradation of visual input (Wolbers, Hegarty, Büchel, & Loomis, 2008). Baseball players learn the speed tendencies of pitchers to perfectly time and place swings within a time window smaller than the blink of an eye (Kwon & Knill, 2013). Indeed, the fluidity and economy of our movements have been suggested to rely partially on an existing model that facilitates feed-forward inferences about the desired goal state accounting for prior experience in performing requisite actions (Wolpert, 2007).

¹ A version of this chapter is under review as Valadao, D. F., Anderson, B., & Danckert, J. Parietal and Cingulate Cortex Involvement in Strategy Updating Using a Serial Competitive Zero-Sum Game. *Frontiers in Neuroscience*. It is reproduced here with permission.

Mental models have also been shown to be important across a range of cognitive domains. For example, visual search has been shown to be more efficient as the configuration of distractors becomes more consistent (i.e., contextual cueing; Jiang, Swallow, & Rosenbaum, 2013), suggesting the presence of a higher order representation that aids in efficient processing. Furthermore, a number of studies have begun to reconsider the influence of sequential effects (e.g., the “gambler’s fallacy” and “probability matching” behaviours), long believed to be evidence of suboptimal human decision making (Tversky & Kahneman, 1974). That is, these so-called ‘suboptimal’ strategies can be recast as a reflection of an individual’s attempt to model global changes in the environment by representing local changes in statistical probabilities (Yu & Cohen, 2009; Gaissmaier & Schooler, 2008; Green, Benson, Kersten, & Schrater, 2010; Griffiths & Tenenbaum, 2006).

Bayesian theorists have long suggested that individuals enter a given context with a set prior, namely an implicit but demonstrable belief about the elements of a task or the probability of relevant events (Tenenbaum et al., 2011). Such priors allow us to determine how discrepant a given outcome is relative to predictions based on the prior, which in turn can be used to evaluate the efficacy of an existing model (Yu, 2007; Nassar, Wilson, Heasly, & Gold, 2010). Thus, this mechanism represents one plausible way in which mental models can be updated to accommodate new information.

Research on statistical learning has also suggested that individuals are capable of representing regularities evident in a seemingly random and chaotic environment (Saffran, Johnson, Aslin, & Newport, 1999), and are capable of doing so practically at birth (Bulf, Johnson, & Valenza, 2011). By passively observing the environment, one is able to extract the statistical organization of objects or events through either its spatial configuration or temporal

order, by extracting the transitional probabilities that exist between a set of items or sounds (Fiser & Aslin, 2001; 2002). This process typically occurs without awareness and is demonstrated by higher than chance performance on forced-choice recognition tasks or by behavioural facilitation in speeded response tasks (Turk-Browne, Jungé, & Scholl, 2005). While models like those discussed above are undoubtedly useful and efficient means of interacting with our world, they would be of little utility if they were unable to change in step with changes evident in the environment. As new information contradicts the outcomes predicted by an existing model, the model should be updated to account for these events.

A number of studies have pointed to spontaneous updating of mental models based on participants' observation of new information that contradicts a prior model. Research from our lab indicates that healthy individuals are capable of detecting a change in a computer's play strategy in a zero-sum game and subsequently act to exploit it by choosing the option that increases their win rate (Danckert, Stöttinger, Quehl, & Anderson, 2012). Brown and Steyvers (2009) showed that participants optimally change judgments regarding the nature of an underlying distribution of events in step with true changes to the distribution. Indeed, recent research has shown that individuals are capable of dynamically shifting their priors over time in order to compensate for a shifting environment (Berniker, Voss, & Kording, 2010). Bayesian computational models corroborate these results with models that utilize non-stationary priors, which allow the model to modify and rebuild an existing prior as a function of the extent to which a prediction error exceeds acceptable limits of uncertainty (Nassar et al., 2010; Wilder, Jones, & Mozer, 2009).

Aside from updating in response to discrepant observations, models can also change when more information is gleaned about the underlying generative process of seemingly random

data. Green and colleagues (2010) explicitly cued participants to the underlying probability of outcomes being generated by manipulating a physical aspect of the task (i.e., changing the shape of a pinwheel such that larger sections reflected greater probability of a target landing in that area). When the physical change indicated higher target probability for a given region, participants abandoned the typical ‘probability matching’ behavior (i.e., choosing a given target region as often as they observed targets appearing there; Geng, Soosman, Sun, DiQuattro, Stankevitch, & Minzenberg, 2013) in favour of the more optimal ‘maximization’ strategy (i.e., choosing the larger region 100% of the time; see Koehler and James, 2010 for review). Taken together, these studies suggest that individuals are constantly seeking to accurately represent their environment; and when given discrepant information or cued to additional relevant information, they update their model and shift behaviour accordingly. The research discussed above highlights the well-established prevalence and importance of mental models. Furthermore, numerous learning mechanisms have been proposed for how mental models of the environment are generated. However, the neural correlates of this behavior have yet to be clearly laid out.

We previously examined the ability of brain-damaged individuals to develop and update an internal model of a non-stationary biased strategy of a computer player in a zero-sum decision-making game (i.e., “rock, paper, scissors”; Danckert et al., 2012). Participants played against a computer that, while initially playing randomly, began to adopt a biased strategy, playing one option more often than others (Danckert et al., 2012). By the end of the experiment the computer opponent had adopted a strong bias (i.e., one option played on 80% of trials). Unsurprisingly, healthy controls (age-matched to our groups of neurological patients) effectively detected and exploited the bias by playing the optimal choice to beat the opponent more frequently. In fact, healthy controls demonstrated the commonly observed, although suboptimal,

probability matching behavior. That is, they chose the play option that would beat their opponent's biased choice approximately 80% of the time (which would lead on average to only a 64% win rate). In contrast, left brain damage (LBD) patients actually outperformed controls by adopting a probability maximizing strategy – i.e., by choosing the option that would beat the most likely computer choice (played on 80% of trials) 100% of the time thereby leading to an 80% win rate. Right-brain damaged (RBD) patients took significantly longer to begin exploiting the bias and failed to fully capitalize on the bias at the levels obtained by the either controls or LBD patients, with the bulk of the RBD group continuing to play randomly even when faced with an opponent's strong bias. Lesion overlay analyses in the RBD group indicated that the superior temporal gyrus (STG), insula and putamen were commonly involved and are thus likely to be important structures for updating a mental model. Those same structures were again implicated in a second group of RBD patients we studied who failed to update play strategy in a different version of rock, paper, scissors and took longer to update perceptual representations of ambiguous figures (Stöttinger et al., 2014).

Functional neuroimaging in healthy individuals highlights similar brain regions to those lesioned in our two RBD groups. Paulus and colleagues (2005) used an event-related fMRI paradigm to explore the brain regions activated when participants played 'rock, paper, scissors', against a computer opponent. While prior studies have examined 'rock, paper, scissors' in the context of fMRI research, they have mainly focused on action and perception (Dinstein, Hasson, Rubin, & Heeger, 2007) or reinforcement signals (Kadota, Nakajima, Miyazaki, Sekiguchi, Kohno, & Kansaku, 2009; Vickery, Chun, & Lee, 2011) on a trial-by-trial basis. In contrast, Paulus and colleagues looked at both outcome and subsequent action selection processes following wins, ties, and losses. Results showed that the insula was activated primarily when

evaluating outcomes—particularly when the outcome was negative. Further, the STG became active primarily during play choice. The authors suggest that these areas are important for evaluating outcomes in the service of formulating a strategy.

While the work of Paulus and colleagues highlights the role these areas play in a zero-sum game, what remains to be seen is what role these areas play in adapting to a *change* in context requiring the updating of one's model. Paulus and colleagues merely changed which strongly preferred play choice the computer tended towards over the course of their experiment (participants exploiting the preferred choice would win 90% of trials if played on every trial) rather than altering the weight of the bias over time. In our view, this limits the parameters of potential biases the participant must model and exploit, making the task akin to a set-shifting exercise, akin to the Wisconsin Card Sorting task (Buchsbaum, Greer, Chang, & Berman, 2005) which we showed in prior work was unrelated to deficits in exploiting an opponent's bias in 'rock, paper, scissors' (Danckert et al., 2012; Stöttinger et al., 2014). That is, once the fact that one choice yields success is discovered, all that remains is to determine whether the preferred choice (set) has shifted. Furthermore, while Paulus and colleagues focused on establishing the regions involved in each component of the decision making process on a trial-by-trial basis (e.g., feedback processing, action selection), the current study focuses on an overall network involved in modeling regularity over time relative to randomness. The current study also utilizes a more complex computer algorithm with multiple bias probabilities, such that participants playing the preferred choice would not always result in a high rate of success. Furthermore, the current study extends these findings to determine whether the relevant regions are sensitive to different types of bias (i.e., no-bias, heavily biased to one option, moderately biased towards two options).

In addition, the influence of reinforcement and rewards has not yet been explicitly examined in this context. Insofar as previous studies have examined the ability to update one's model, participants also experienced greater rewards as their behavior changed. Thus, it may be possible that the experience of having received explicit rewards and the inference that a mental model has been updated may have been conflated (e.g., Paulus et al., 2005). It is conceivable that participants, whose behavior is taken as representing an updated mental model, are simply driven by transient changes in reward. For example, they may be changing strategy in response to seeking or securing increased wins as opposed to actually figuring out something new about their environment, with all the risk that following through on one's beliefs entail. In the current study we attempted to decouple the reinforcement process from the responses that generate them.

The current study employed a protocol adapted from Danckert and colleagues (2012) in which participants played rock, paper, scissors against a computer opponent that utilized a variety of strategies. The protocol was modified for fMRI in order to elucidate a network of areas that support the updating of an individual's mental model as evidenced by behavioural exploitation of the shifting biases presented in the computer's play. It was expected that individuals would demonstrate the ability to model and exploit the relatively simplistic computer stratagem and that exploration of neural activations under this context would reflect a network of right-hemisphere areas supporting the building and updating of mental models.

2.2. Method

Participants

Twelve (5 male; mean age = 29.9, SD = 5.2) neurologically healthy participants with no history of psychiatric or neurological illness participated in the experiment. Each participant

completed 5 experimental runs. Three of the 60 runs were excluded due to technical reasons. All participants had normal or corrected-to-normal visual acuity. All participants provided informed consent prior to participation and the research protocol was approved both by the Office for Research Ethics at the University of Waterloo and the Tri-Hospital Research Ethics Board of the Region of Waterloo in Ontario, Canada. Participants were given \$20 remuneration for their involvement in the current study.

Design

Participants played a zero-sum game of rock, paper, scissors. In this task, participants were to choose one of three options (i.e., rock, paper, or scissors) such that any chosen option could either beat the computer opponent, be beaten by the opponent, or tie the opponent. Participants were informed that the goal for each trial was to defeat the computer opponent's choice.

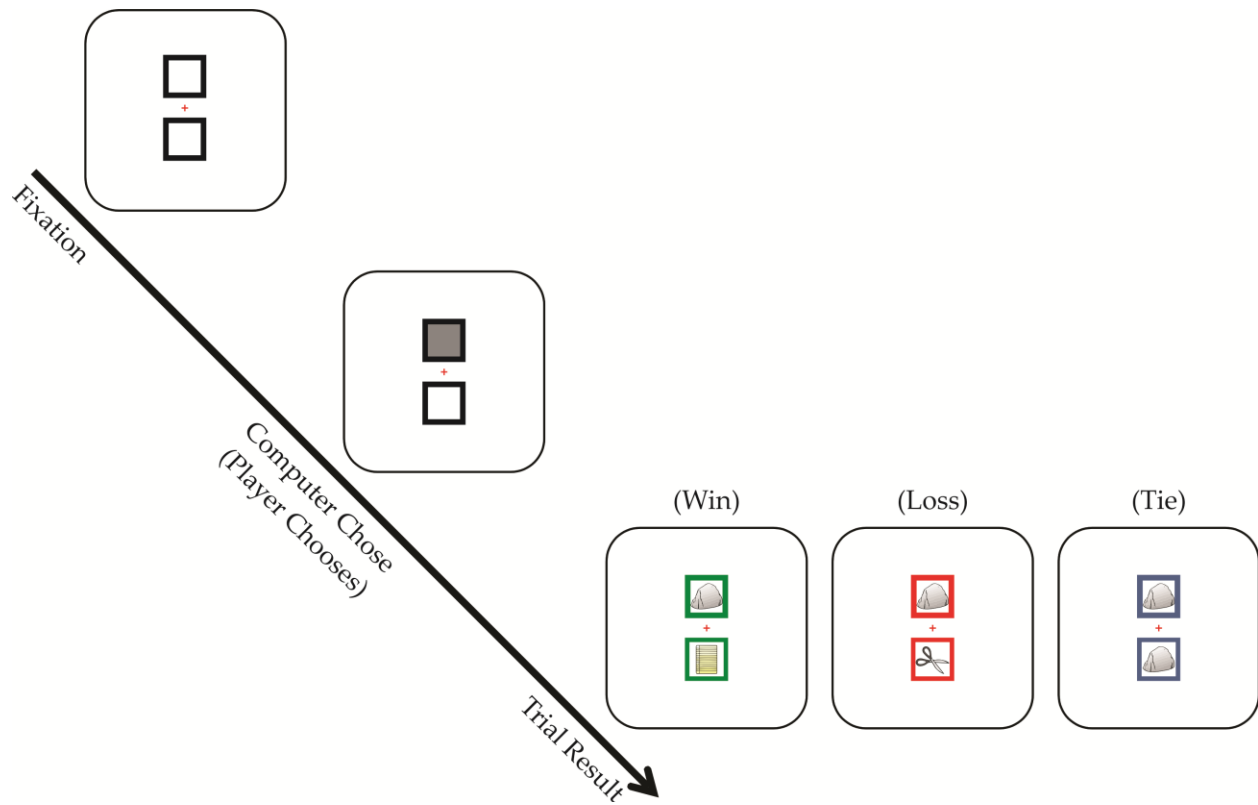


Figure 2.1. Rock Paper Scissors Stimuli. Schematic representation of the experimental stimuli, display, and procedure that took place on each trial regardless of condition. Computer (top box) would choose first prompting user input. Trial results show computer and player's choice in the top and bottom boxes, respectively. Box outline would change to green, red, or blue to depict trial win, loss, or tie, respectively.

Each trial began with a vertical array of two boxes; one above and below a fixation cross. Fixation was not strictly enforced. Participants were informed that the upper box represented the computer's choice while the lower box represented their own choice. After 200 ms, the computer opponent's selection box turned gray, indicating that the computer had made its choice and was awaiting the participant's decision. Participants then indicated their choice by pressing one of three buttons on a response box. Immediately after participants made their choice the trial results would appear. Pictorial representations of the choices made were revealed in the upper and lower boxes for the computer and participant respectively (Figure 2.1). The outlines of the boxes, previously black, would change colour to green if the participant won, red if they lost, and blue for a tie. Each trial, in total, lasted 2500 ms. In the rare event that participants took longer than 2 seconds to respond from the beginning of the trial, they would receive feedback that they were too slow by being shown both boxes with a gray interior and red outline instead of the typical trial result described above. This occurred infrequently (0.7% of all trials across all participants), and when slow responses did occur they were predominantly evident at the very beginning of the experiment (i.e., 64% of all slow responses occurred during the first epoch while participants became familiar with the time window with which they were required to respond). Only one participant had more than 1% of trials in which they responded too slowly (2%).

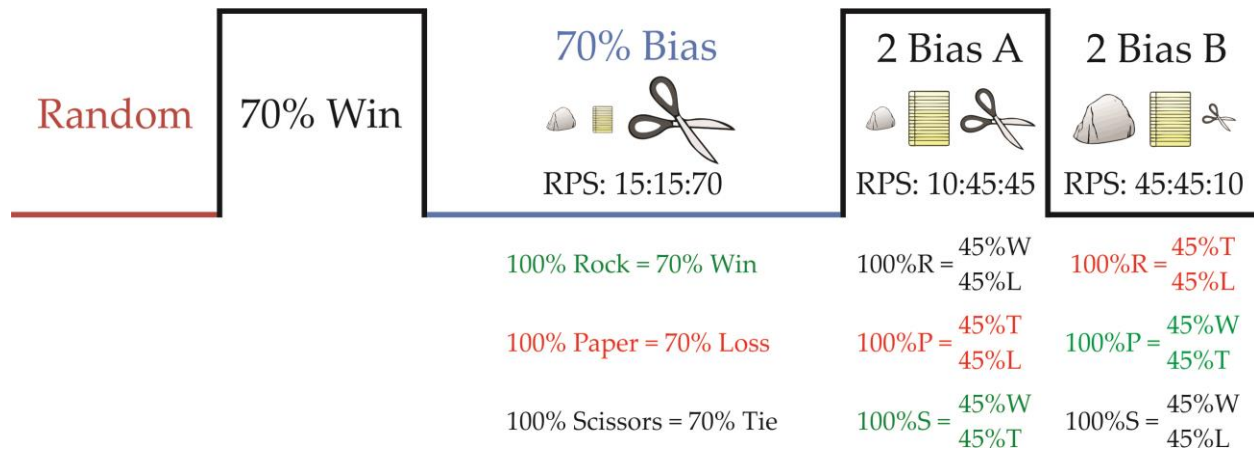


Figure 2.2. fMRI Run Conditions. Schematic representation of a single experimental run within the experimental protocol. Different epochs are denoted by the square wave and are labelled by the bias adopted by the computer opponent (with example ratios for each choice indicated where appropriate). Below this, possible player strategies for each condition and the anticipated outcomes corresponding to adherence with each respective strategy are shown (e.g., for the 70% bias if a participant chose rock 100% of the time this would lead to a 70% win rate). Optimal strategies are indicated in green while suboptimal strategies are depicted in red.

While the goal of each *individual* trial was to beat the computer opponent, the strategy of the computer was not stationary. Participants were not informed of this aspect of the task design. Throughout each experimental run the computer utilized five different strategies for a minimum of 30 trials (Figure 2.2). Paulus and colleagues (2005) demonstrated that 20 trials was sufficient to demonstrate learning of relatively simple biases, and reported much of the learning observed was found in the first 10 or so trials. Therefore, it was determined that 30 trials would allow ample opportunity to model the more complex computer behaviours while still getting sufficient data collected within a confined window of time. In order to set a behavioural baseline, as well as a functional control for subsequent contrasts, the computer started by playing randomly for 30 trials (75s per run); whereby each of the three possible choices were played on one third of the trials (R, P, S: 33%, 33%, 33%). This was followed by a condition in which the computer ‘let’ the participant win for 30 trials by simply choosing the option that would be beaten by the participant’s choice on 70% of trials. This arbitrary win rate was intended to provide participants with positive reinforcement and increased wins relative to the random condition in a manner that did not require the participant to develop a model of the computer opponent’s tendencies by detecting and exploiting an observed bias in its play (i.e., strictly speaking, there was no observable bias). This allowed us to explore the question of whether merely winning at a certain rate would activate regions thought to be necessary for updating. This condition involved some level of deception in that participants were told prior to the experiment that when the computer’s choice square changed colour this meant that the choice had been “locked in” and would not change for the duration of that trial. While true for all other conditions, this arbitrary win rate condition involved the computer’s choice box changing colour prior to the participant choosing,

with the actual choice of the computer being made *after* the participant's choice was entered to ensure the participant won on 70% of trials.

From there, the computer began playing according to a specific bias, whereby it chose one option on 70% of trials, choosing each of the other two options 15% of the time (e.g., R, P, S: 15%, 70%, 15%) for a total of 60 trials (150 s per run).

In order to evaluate participants' ability to navigate more complex scenarios requiring mitigation of losses (as opposed to maximizing reward as in the 70% bias), the computer then switched to what we call a "two-bias" strategy for 30 trials; whereby the computer chose two options on a combined 90% of the trials (i.e., each respective option was chosen on 45% of trials), with a third option being chosen on only 10% of trials. This 'two-bias' structure continued the prior option favoured in the 70% bias albeit to a lesser degree (45% of trials vs. 70%). The second bias of 45% was then chosen to defeat plays that would have been considered optimal under the 70% bias. For example, if the computer had previously been biased towards playing rock, thereby reinforcing increased play of paper, the computer would then play rock and scissors most often. If the participant continued to prefer paper choices – that is, the optimal play under the previous 70% bias structure – they would experience an increase in losses. However, if they were able to formulate a new model of the computer's strategy, and act to minimize losses, they would switch to preferring rock; a strategy that would provide a similar amount of wins but fewer losses (i.e., more ties). To the extent that participants engaged in this change of strategy, their play could be considered optimal in that it minimized losses and maximized wins (Figure 2.2). It is also important to note that participants, technically, could exhibit a suboptimal strategy whereby they experience primarily ties and losses (e.g., choosing scissors). As such, for each

permutation of the “two-bias” strategy that the computer could adopt, there was both an optimal and suboptimal strategy that participants could adopt.

The computer then switched to 30 trials of a final two-bias strategy that, again, counteracted the previously optimal play strategy. As with the previous two-bias strategy, there was an optimal participant strategy (wins and ties) and a suboptimal strategy (ties and losses; Figure 2.2). This final condition enabled us to directly examine brain regions responsible for updating representations in order to mitigate or avoid losses by directly comparing each two-bias condition; whereas typical comparisons in behavioural decision making tasks focus on requiring participants to simply maximize rewards values. To prevent participants from merely memorizing the correct choice over the course of the five experimental runs, each run had the computer exhibit a bias towards a different set of choices than that of the previous run.

fMRI data collection

Functional data were collected using T2*-weighted images collected on a 1.5 Tesla Phillips machine located at the Grand River Hospital medical imaging department in Waterloo, Ontario (TR = 2500 ms; TE = 40ms; slice thickness = 5mm with no gap; 26 slices/volume; FOV = 220 x 220 mm²; voxel size = 2.75 x 2.75 x 5mm³; flip angle = 90°). Each experimental run began and ended with 5 volumes of fixation. At the beginning of each session, a whole-brain T1-weighted anatomical image was collected for each participant (TR = 7.5ms; TE = 3.4ms; voxel size = 1 x 1 x 1mm³; FOV = 240 x 240mm²; 150 slices with no gap; flip angle = 8°). Experimental stimuli were presented to participants in the magnet using an Avotec Silent Vision™ fibre-optic presentation system using binocular projection glasses (Model SV-7021). The experimental protocol was programmed using E-Prime experimental presentation software

(v1.1 SP3; Psychology Software Tools, Pittsburgh, PA). The onsets of trials were synchronized to each functional volume collection using trigger-pulses from the magnet.

fMRI Data analysis

Functional imaging data were analyzed using Brain Voyager (v2.6, Brain Innovation B.V., Maastricht, Netherlands) with each participant's low-resolution functional data aligned to their respective high-resolution anatomical image and subsequently transformed into Talairach space (Talairach & Tournoux, 1988). All functional runs were subject to trilinear interpolation to correct for motion artefacts. To determine whether significant motion occurred during any one run, pre- and post-correction images were compiled into a virtual movie that was played to look for instances of movement which would render the data inappropriate for the sample (Culham et al., 2003). No serious motion artefacts were present in the sample. All functional data were then pre-processed, including the use of linear trend removal and a spatial smoothing filter of 8mm FWHM. Predicted activations were convolved to the Boynton hemodynamic response function. For each run, linear predictors were applied corresponding to the timeline of the computer opponent's strategy. The predictors were then entered into a General Linear Model where the significance for each voxel was evaluated as being significant below threshold of .001, which was below a False Discovery Rate threshold of .05. Additionally, a cluster size threshold of 20 voxels was applied to each contrast.

To extract relevant areas of activation in the current study, seven contrasts were computed. First, to identify brain regions involved in exploiting a biased computer strategy, we contrasted the BOLD signal of the 70% frequency biased condition to the random condition. To investigate the role of non-contingent rewards on how individuals model their environment, we

contrasted the 70% arbitrary win condition to the random condition. To examine the brain regions involved in model building above and beyond non-contingent reinforcement, we contrasted the 70% bias condition to the 70% arbitrary win condition. To examine the regions sensitive to changes in potential loss rates we contrasted each two-bias condition to the random condition, independently. We also contrasted the second two-bias condition against the first to determine whether there were any additional areas preferentially involved in switching from one model to another. Finally, to investigate the possibility of difficulty-sensitive areas, we contrasted the two-biased and one-bias conditions. In all contrasts, significant regions of interest are reported with centroid co-ordinates, cluster size and peak t-scores.

Behavioural Data Analysis

In order to determine the neural regions supporting the building and updating of a mental model, it was first necessary to demonstrate that participants were indeed capable of building an accurate model in the first instance. To demonstrate this, participants' play choices on each trial of the experiment were coded as being optimal, suboptimal, or neither. For example, if the computer was biased towards playing scissors, the optimal choice would be rock, suboptimal would be paper, and scissors would be neither. It was then possible to examine whether participants chose optimally above chance, and minimized suboptimal play, for each condition of the experiment. For conditions where no clear optimal choice was possible (i.e., random and arbitrary win conditions), participant's modal choice during that epoch was tracked to determine how participants developed a preference for that particular play choice. Behavioural data were subjected to both linear and non-linear curve estimations regressed over each trial for each condition. The Durbin-Watson statistic was used to determine the extent to which these data

were auto-correlated. Examination of fit statistics suggested that a linear trend provided the best fit to these data under all conditions except for the two-bias conditions where a quadratic estimation provided significantly better fit.

2.3. Results

Examination of participant performance when playing against the 70% bias revealed proportions of optimal plays that were significantly above chance, $t(11) = 7.89$, $SE = .03$, $p < .001$. Further, participants simultaneously minimized suboptimal play choices such that the proportion of suboptimal play was significantly below chance, $t(11) = 15.23$, $SE = .01$, $p < .001$. Examination of the time course of participant's ability to adapt to the computer's strong bias revealed a significant trend whereby participants were capable of adapting to the biased computer strategy (Table 2.1), reaching probability matching thresholds of optimal play within the 30 trial window (Figure 2.3). Consequently, it can be inferred that participants successfully represented and exploited the bias in computer play choices (Figure 2.3).

Table 2.1. Curve Estimation Statistics.

Condition	Linear				Quadratic			
	b_{lin}	Fit	F_{sig}	D-W	b_{lin}	b_{quad}	Fit	F_{sig}
Arbitrary Win Modal Choice	.08	12.02	<.01	1.34	-.25	.35	12.14	<.001
1-Bias Early	.22	83.29	<.001	1.46	.87	-.66	42.69	<.001
1-Bias Late	.09	14.29	<.001	1.54	-.04	.14	8.18	<.001
1-Bias Combined	.22	168.97	<.001	1.49	.48	-.28	93.24	<.001
2-Bias A	.10	4.38	<.001	1.58	1.80	-1.70	27.46	<.001
2-Bias B	.03	1.32	.251	1.62	-2.96	2.99	21.19	<.001

Linear and quadratic estimation statistics where participant play choices were regressed over the course of each experimental epoch. Arbitrary win condition play choice indicates the rate of frequency with which participants played their most common choice during that epoch over the course of the trials. Remaining conditions were analyzed using the frequency with which participants selected the most optimal choice for each trial of each condition.




Figure 2.3. Rock Paper Scissors Behavioural Data. Depicts the proportion of optimal choices across each experimental condition. For all but the superstitious condition, the optimal choice was the move that carried the best chance of an optimal result (Figure 2.2). The superstitious condition, where participants were guaranteed to win on 70% of the trials, did not have a particular bias that participants could exploit. Therefore, the modal choice for each participant throughout the condition was plotted over time. Performance within the shaded area indicates failure to reach probability-matching thresholds. Dashed red line depicts random chance threshold. Dashed blue lines about the curves indicate 95% confidence intervals.

Contrasting the 70% bias condition against the random control condition revealed a network of activations including bilateral STG, right anterior insula, right inferior parietal lobule (IPL), and anterior cingulate cortex (ACC; Table 2.2). With respect to the 70% arbitrary win condition, the intent was to provide participants with reinforcement signals that were disconnected from the requirement to model and exploit the computer's play. Ideally, participants would receive greater amounts of wins relative to the random condition, but would do so regardless of their own play strategy. A priori we assumed that participants would play randomly both in the random and arbitrary win-rate conditions. They did not. When playing in the 70% arbitrary win condition, participants played *as if* they had discovered a bias in the computer's play and acted to exploit it. That is, rather than playing randomly, participants began to favour one option over another much as they did in the 70% bias condition (Figure 2.3). The play choice made most often by participants in this condition (i.e., the choice participants putatively believed would result in the most wins) was played significantly greater than chance, $t(11) = 5.65$, $SE = .05$, $p < .001$. Examining the time course of play (Figure 2.3) suggests that participants believed they had "discovered" a bias in play and began to prefer the option that would win at an increasing rate over time, such that they adopted a play choice that mimicked probability matching (Table 2.1; Figure 2.3). It is important to note that this is not merely a carry-over effect from the previous trials (i.e., the random condition). The play option that participants heavily favoured in the 70% win condition was played significantly more in that condition than it was in the previous random condition, $t(11) = 4.82$, $SE = .05$, $p < .001$. Further, participants' preferred play option in the 70% win condition was not chosen at a greater-than-chance rate in the previous random condition, $t < 1$.

Given the similarities between participants' performance in the non-contingent 70% win condition and how they performed when facing a computer that exhibited a true bias, it would not be altogether surprising to find similar areas of activation when computing a similar contrast. That is, in the 70% win condition participants played as if their opponent had adopted a frequency bias. Thus, when we contrasted the 70% win condition to the random control, we found a similar network of cortical areas including the right STG, right ACC, and right MFG (Table 2.2). Comparing the 70% win condition with the 70% bias condition yielded no significant activations. This finding is consistent with the behavioural data indicating that the participants treated the strategy played by the computer similarly in both conditions.

Table 2.2. Whole-brain GLM Activations for computed contrasts (p<.001).

Contrasts	Co-ordinates			Activated Region(s)				
	X	Y	Z	R/L	Region	BA	Cluster Size	Peak t-score ^a
1. 70% Bias > Random	60	-11	8	R	STG	22	1483	5.77
	54	-10	-11	R	MTG	21	1029	4.55
	45	-6	-27	R	ITG	20	756	5.58
	30	-42	53	R	IPL	40	2548	6.62
	34	-20	53	R	PCG	4	3410	7.46
	10	-46	56	R	Precuneus	7	2115	5.35
	9	-46	31	R	PCC	31	2683	6.14
	12	-1	-12	R	Ventral striatum	--	687	5.45
	42	-4	2	R	Post. insula	13	3257	5.14
	38	-9	14	R	Ant. insula	13	1095	5.38
	1	-13	41	R	ACC	24	2384	5.66
	25	21	-15	R	OFC	47	2506	5.89
	-15	0	-13	L	Ventral striatum	--	1998	6.00
	-8	-49	27	L	Precuneus	31	4686	6.41
	-3	-32	57	L	SMA	5	1785	6.52
	-57	-6	-9	L	MTG	21	4931	8.22
	-52	-55	24	L	STG	22	3468	5.83
2. 70% Win > Random	58	-12	1	R	STG	22	4825	6.60
	56	-34	4	R	MTG	22	1875	5.6
	27	-39	57	R	IPL	40	1181	5.11
	25	-30	61	R	PCG	3	3858	5.92
	1	-3	36	R	ACC	24	3634	5.90
	5	-50	28	R	PCC	31	3574	5.86
	7	14	6	R	Caudate	--	1445	5.25
	32	16	-13	R	OFC	47	9514	7.03
	0	61	7	R	vmPFC	10	2872	6.38
	6	45	39	R	dmPFC	8	2094	4.77
	10	7	-8	R	Ventral striatum	--	4199	7.28
	-52	-59	24	L	TPJ	39	5490	7.83
	-57	-12	-9	L	MTG	21	7128	8.18
	-17	0	-13	L	Ventral striatum	--	2767	6.33

	-40	23	-15	L	OFC	47	953	5.06
	-15	44	44	L	dmPFC	8	1655	5.21
	-5	-52	26	L	PCC	31	7215	7.22
	-19	-47	55	L	Precuneus	7	651	5.24
3. 70% Bias > 70% Win	--	--	--	--	No activations	--	--	--
4. 2-Bias (A) > Random	--	--	--	--	No activations	--	--	--
5. 2-Bias (B) > Random	--	--	--	--	No activations	--	--	--
6. 2-Bias (B) > 2-Bias (A)	43	-57	37	R	TPJ	39	4417	5.41
	45	-45	38	R	IPL	40	3162	5.42
	9	31	45	R	dmPFC	8	4332	5.86
	37	15	46	R	dIPFC	8	2482	4.47
	29	57	14	R	vmPFC	10	4533	6.53
	48	21	-9	R	OFC	47	1950	5.50
7. 2-Bias (A+B) > 70% Bias	15	-18	7	R	Thalamus	--	4588	6.80
	50	-35	-11	R	MTG	20	1929	6.22
	43	-73	-13	R	Fusiform Gyrus	19	1584	6.43
	-28	41	14	L	vmPFC	10	1553	5.21
	-38	-44	38	L	IPL	40	1622	5.14

Rand, control condition where participant played against computer playing without bias towards any of the three choices. *70% Bias*, condition where computer played one option on 70% of the trials, with each other choice being played 15% of the time. *70% Win*, condition where computer rigged game to allow participant to win on 70% of trials regardless of their choice. *Two-Bias A/B*, conditions where computer was biased towards two of three options (45% of trials for each) such that sticking to previous optimal strategy would yield reduces success (i.e., wins, but increased losses). *STG*, superior temporal gyrus; *STS*, superior temporal sulcus; *MTG*, medial temporal gyrus; *IPL*, inferior parietal lobule; *SMA*, supplementary motor area; *SMG*, supramarginal area; *ACC*, anterior cingulate cortex; *PCC*, posterior cingulate cortex; *PCG*, precentral gyrus; *MFG*, medial frontal gyrus.

^a The peak t-statistic within a specific area of activation. All voxels within specified area surpass significance threshold of $p < .001$.

When participants played against the two-biased computer opponent, a strategy requiring a more complex representation factoring in an increased risk of losses, there was not a clear-cut choice that participants were able to arrive at. None of the three options would guarantee success on over half the trials, forcing participants to model the likelihood not only of winning but also of losing (or tying) with their opponent. Arguably then, the most optimal choice in these epochs would be the one that offers a greater chance at wins, while minimizing exposure to losses. Participants were successfully able to model the complex play of the opponent in these circumstances and chose the response that maximized wins while minimizing losses at a greater than chance level in both the first, $t(11) = 3.75$, $SE = .02$, $p < .01$, and second epochs, $t(11) = 7.07$, $SE = .02$, $p < .001$. As with the other conditions, participants also avoided the worst possible choice and played it at a level significantly lower than chance in both the first, $t(11) = 14.03$, $SE = .02$, $p < .001$, and second epochs, $t(11) = 5.67$, $SE = .02$, $p < .001$. Examining the time-course of participant responses also reveals a significant quadratic trend in their response style that reached probability matching thresholds by the end of the epoch on this more difficult task, consistent with what was observed in the other conditions (Table 2.1; Figure 2.3).

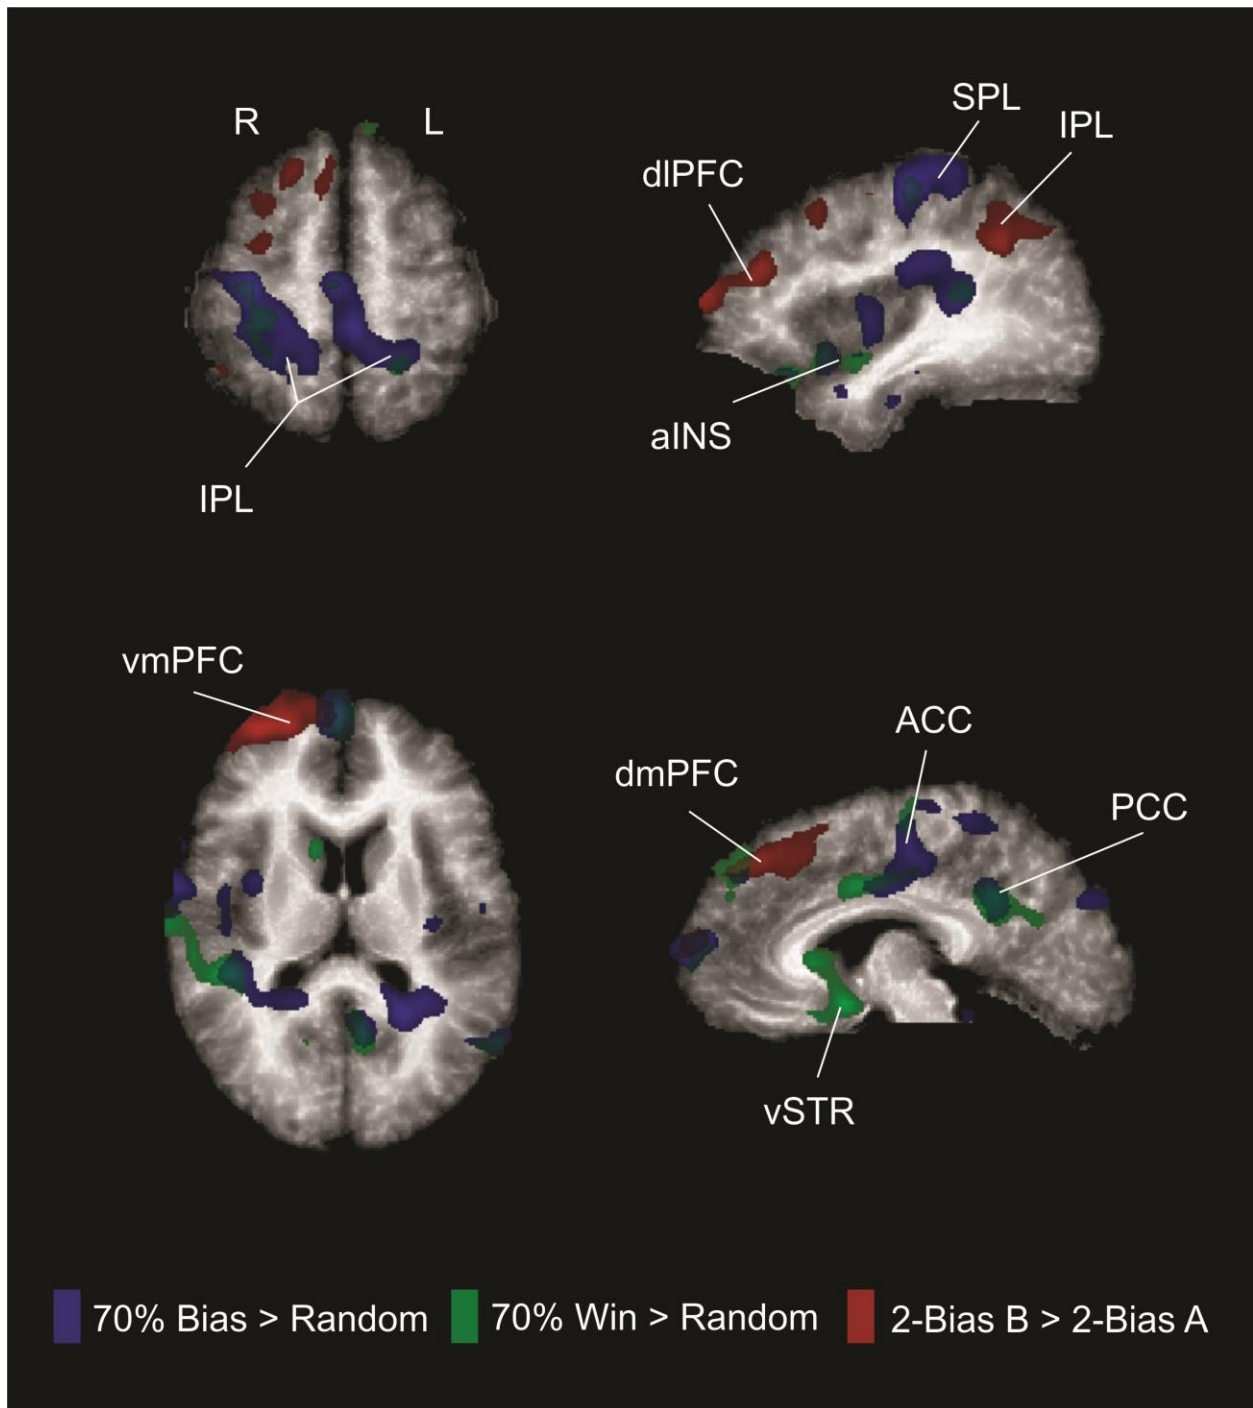


Figure 2.4. Whole-brain GLM Contrast Volume Maps. Significant regions activated when comparing the 70% bias vs random condition are depicted in blue, 70% win vs random condition in green, and 2-bias B vs 2-Bias A in red.

However, compared to the previous conditions, participants were not able to demonstrate consistent optimal play that persisted throughout the epoch as with the easier conditions. Examination of each respective 2-bias epoch shows that participants' optimal play choice dipped significantly below random chance thresholds following the beginning of each epoch. Critically, participants' optimal play fell below random chance thresholds at the end of the first 2-bias epoch, suggesting that they weren't able to effectively exploit the computer's strategy by the end of this difficult condition despite playing at above probability matching thresholds earlier. Conversely, participants optimal play choices were significantly above probability matching thresholds by the end of the later 2-bias condition, suggesting that participants needed additional experience with the more difficult computer strategy to learn its tendencies, but that with sufficient experience they were able to effectively exploit this bias. Given that participants did avoid the most obvious suboptimal choice (i.e., the choice that resulted in predominantly losses and ties), the variability in optimal play under the two-bias conditions most likely reflects pattern seeking, whereby participants knew what two options were best and sought out the right balance in which to play them for maximal success.

As with the conditions involving higher win rates or stronger biases (i.e., 70% arbitrary win and 70% bias respectively), contrasting the second two-Bias condition with the first two-biased condition revealed a similar network of right-hemisphere cortical areas including the TPJ, IPL, vmPFC, and dmPFC (Table 2.2; Figure 2.4). However, comparing each respective two-biased condition against the random control yielded no significant activations.

Given the relative heterogeneity of our conditions and the sparse pattern of activations throughout the explored contrasts, we submitted the dataset to a probability mapping procedure where we compared all contrasts that reflected behavioural evidence of modeling and exploiting

the computer's bias. This procedure overlays all volume maps that show significant activations and displays activated regions common to all contrasts that reflect behavioural exploitation by participants. This allowed us to determine the most commonly activated areas during relevant behavioural change across all contrasts (Figure 2.5).

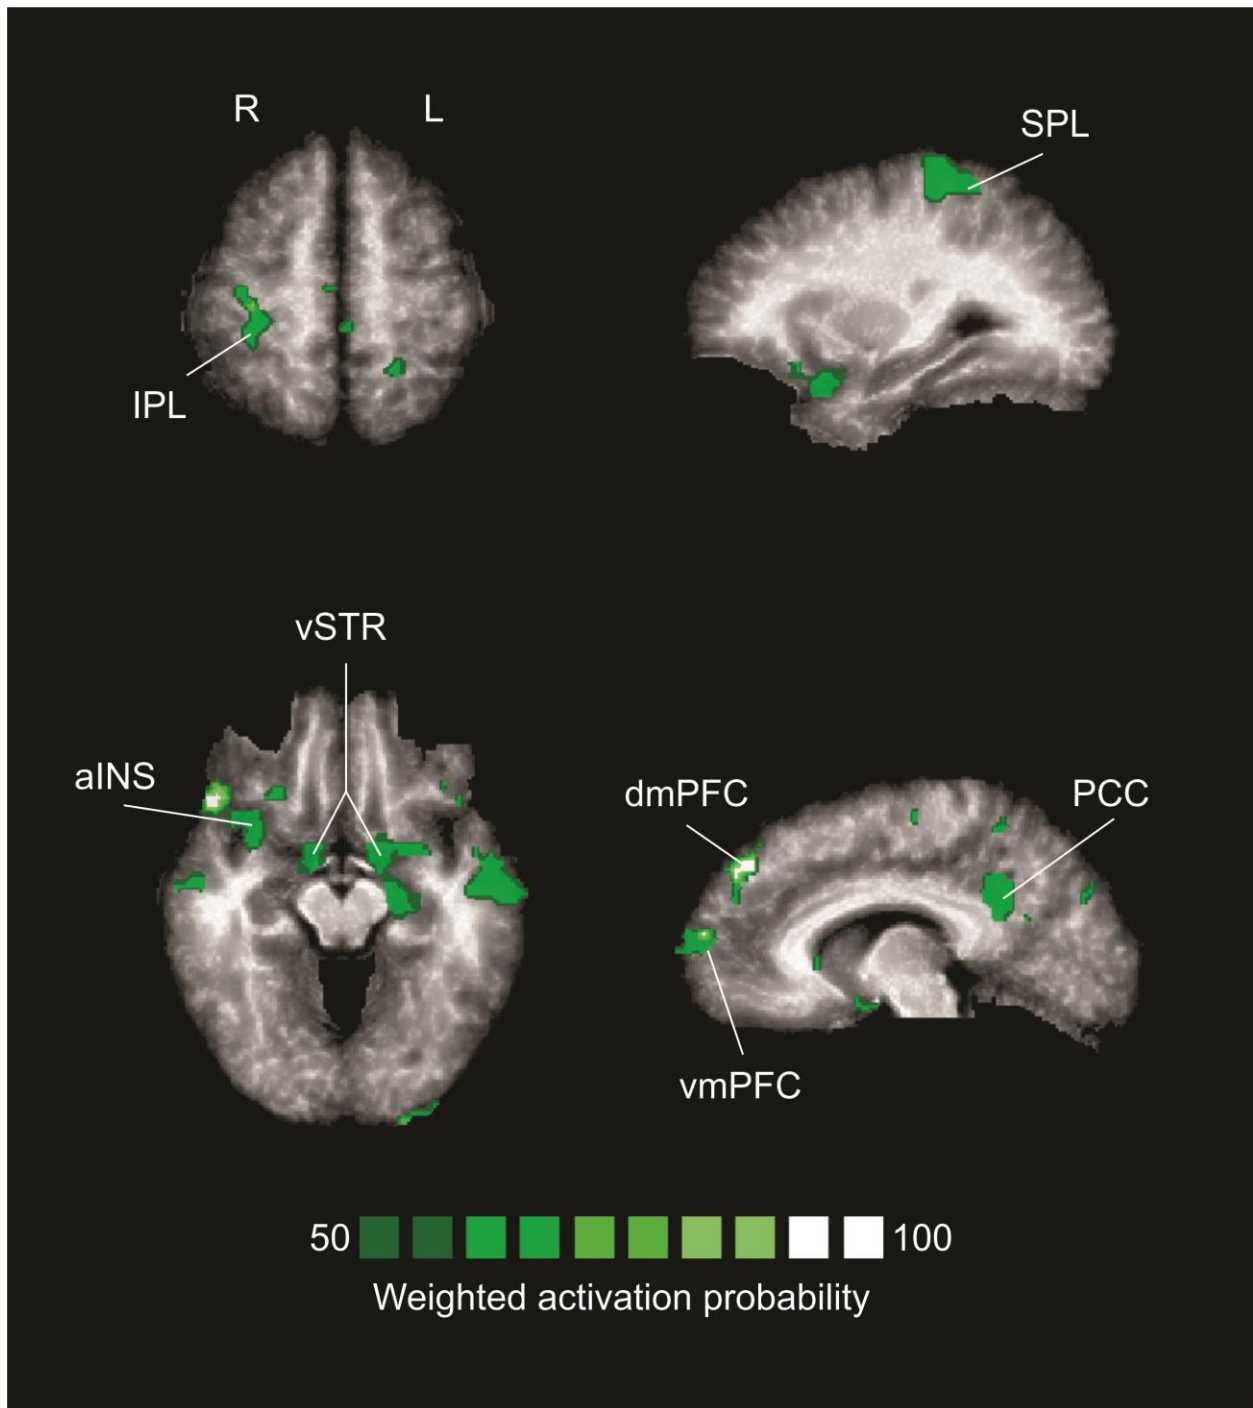


Figure 2.5. Probability Threshold Map. Depicts areas of significant overlap across conditions. Greater weighted activation probability signifies regions with shared voxels across a greater number of contrasts.

For the purpose of this analysis, we focused only on the arbitrary win condition, the 70%-bias condition, and the second 2-bias condition as these were the only instances in which optimal play choice was significantly above probability matching thresholds by the end of the epoch (Figure 2.3). Probability mapping analysis revealed a predominantly right-hemisphere network of areas. Most notably among these are the parietal cortex, cingulate cortex, anterior insula, prefrontal cortex including vmPFC and dmPFC, and reward processing areas including the ventral striatum (see Balleine, Delgado, & Hikosaka, 2007 for review).

2.4. Discussion

Evidence of modeling bias in environmental regularity

The current study examined the brain regions supporting the ability to build and update a mental model of an opponent's bias in a zero-sum game. As predicted, and consistent with previous research (Danckert et al., 2012), participants were consistently able to model the shifts in their computer opponent's play regardless of the complexity of the strategy. Participants consistently showed significantly greater proportions of optimal plays while simultaneously minimizing suboptimal plays, consistently reaching probability-matching thresholds by the conclusion of each respective condition. These behavioural results are consistent with previous research showing that individuals are capable of identifying the underlying distributions, and changes therein, of events over time despite the apparent noise and complexity of trial-to-trial events (Brown & Steyvers, 2009). Whereas some research has highlighted the capacity for shifts in models through explicit environmental cues (Green et al., 2010), the current results emphasize that individuals are capable of this despite the fact that no explicit cue or warning is given to foreshadow a fundamental change in their environment.

Examination of the neural correlates of this behavior focused on contrasting BOLD signal between conditions where participants played against a computer utilizing a biased strategy vs. a control where the computer played randomly. While individual areas may have differed somewhat based on particular contrasts, activations were generally observed in a predominantly right hemisphere network of areas including the rIPL, cingulate cortex, medial prefrontal cortex, and the insula.

Contributions of right parietal cortex to evidence seeking and integration

Previous research on the role of right parietal cortex have commonly suggested the area to be partially responsible for the allocation of attentional resources (Hopfinger, Buonocore, & Mangun, 2000), in line with work in neurological patients showing that damage to these areas results in failures of attentional orienting towards contralesional space (Ferber & Danckert, 2006; Heilman & Edward Valenstein, 2011; Mesulam, 1981). However, more recent evidence has begun to implicate right parietal areas such as the IPL, in higher order cognitive processes, such as the perception of causality (Straube, Wolk, & Chatterjee, 2011), decision making under uncertainty (Paulus, Hozack, Zauscher, McDowell, Frank, Brown, & Braff, 2001; Paulus, Feinstein, Tapert, & Liu, 2004), theory of mind and signaling empathic responses (Decety & Lamm, 2007; Geng & Vossel, 2013). The common thread through these discrete and disparate abilities appears to be less related to orienting attention than to making sense of what is being experienced to support appropriate goal-directed action.

Furthermore, recent meta-analytic research suggests that areas surrounding the right temporoparietal junction (TPJ; namely IPL and STG) are more likely to be responsible for internal model updating rather than stimulus driven attention orientation as it is typically

regarded (Geng & Vossel, 2013; McGuire, Nassar, Gold, & Kable, 2014). This interpretation is consistent with work in our lab implicating the right hemisphere in impaired updating performance. Specifically, when given a similar rock, paper, scissors task, patients with right brain damage typically had greater difficulty modelling and exploiting the simple biases of the opponent relative to age-matched healthy controls (Danckert et al., 2012; Stöttinger et al., 2014).

This conceptualization of TPJ functioning is also consistent with neuroimaging research on healthy participants that has also shown the right IPL to be preferentially activated during evidence seeking and exploration while completing probabilistic decision making tasks (Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; d'Acremont, Schultz, & Bossaerts, 2013; d'Acremont, Fornari, & Bossaerts, 2013; Furl & Averbeck, 2011). In addition, event-related fMRI evidence from Paulus and colleagues (2005) suggest that the right STG may be responsible for hypothesis space generation and strategy integration based on evaluating outcomes (Paulus et al., 2005). The activations of these two areas – namely the IPL and STG – in the current study lend further credence to the notion that they share a critical role in model building by seeking out and integrating evidence into a model over time.

Cingulate cortex may signal the need to update

Functional imaging research examining the role of anterior cingulate cortex has long suggested it to be partially responsible for error detection (Carter, Braver, Barch, Botvinick, & Noll, 1998). As errors occur or tasks become more difficult, suggesting error rates are likely to increase, it is typical to observe a rise in ACC activation (e.g., MacDonald, Cohen, Stenger, & Carter, 2000). However, recent research has begun to examine the role of the ACC from a different perspective. Increasingly, ACC activations have begun to be associated with errors in

“belief-based” decision making. In a recent study, Zhu and colleagues (2012) had participants play a competitive investment game where they were asked to bid from an endowment in order to win a larger prize, but would lose both their initial investment and the prize in the event their opponent outbid them (Zhu, Mathewson, & Hsu, 2012). It was possible to succeed in this task by adopting a purely reinforcement learning scheme; slowly determining the strategic payoff of different levels of investment based on prior outcomes. Conversely, players could adopt a belief-based approach and actively attempt to model the aggressiveness or passivity of their opponent. Doing so would allow participants to rapidly predict the minimum investment necessary to win, resulting in much larger payoffs (i.e., the prize plus the extra endowment they did not need to invest in order to win). Critically, both of these approaches carried the possibility of experiencing errors in decision making. Results showed that ACC activation was higher when the errors were made whilst employing a belief-based decision making approach to the task, relative to a reinforcement learning approach. Further, ACC activations closely tracked individual differences in the level of engagement in the belief-based strategy. The authors concluded that the ACC could be better characterized as responsible for belief-based error signals, rather than merely error detection in the absence of a pre-existing model of the environment.

Aside from simply error processing, research has also implicated the ACC in the process of updating, where volatility in the environment exists relative to prior stability such that individual trial outcomes, positive or negative, must be weighed against expected levels of uncertainty (Walton, Crosson, Behrens, Kennerley, & Rushworth, 2007; Woolrich et al., 2007). Furthermore, Behrens and colleagues (2007) suggest that the ACC is involved in the integration of trial-by-trial outcomes, rather than simply evaluating a single trial in isolation. Indeed, numerous studies demonstrate a compelling relationship between ACC activation and increased

learning rates (Woolrich et al., 2007; McGuire et al., 2014), suggesting that the ACC is partly responsible for signaling when additional hypotheses need to be explored or tested when the environment becomes more volatile compared to epochs of relative stability. Furthermore, posterior cingulate cortex also appears to play a significant role in change detection by keeping track of a cumulative learning rate in a complex environment, which may help to trigger a later shift in strategy as new evidence reaches a critical threshold for change (Pearson, Heilbronner, Barack, Hayden, & Platt, 2011).

As it pertains to the current study, the presence of cingulate cortex activity in our sample may suggest that participants approached the task with the express purpose of modeling the behavior of the opponent; as the computer shifted strategies participants experienced belief-based errors that triggered a re-evaluation of the computer's strategy. Furthermore, under the context of a changing strategy, the cingulate cortex activation may reflect an integration of trial outcomes signaling the need to explore additional hypotheses once the play of the computer diverged from expected outcomes.

Involvement of prefrontal cortex in model maintenance

Activity within the medial prefrontal cortex and related PFC regions (specifically vmPFC, dmPFC, and dlPFC) has been consistently implicated in executive control functions (Geng, 2013; Hare, Camerer, & Rangel, 2009). More specifically, the medial prefrontal cortex (notably the vmPFC) are classically considered to be important for signaling the subjective reward value of stimuli and potential decisions in the context of ambiguity (Kable & Glimcher, 2007; Hare et al., 2009; Rangel & Hare, 2010). This is consistent with lesion symptom mapping research implicating medial prefrontal cortex in reward and outcome valuation in the context of

reinforcement learning (Gläscher et al., 2012). Recent research has also implicated vmPFC functioning in monitoring the reliability of currently utilized strategies (Domenech & Koechlin, 2015). Furthermore, prior research has shown that the ACC and insular cortex appear to modulate vmPFC activity during more complex decision making tasks where action costs and prospective rewards are not entirely clear (see Rangel & Hare, 2010 for review). Activity in the mPFC and ACC seen in the current study may reflect joint monitoring processes to maintain the application of an effective strategy for as long as it remains viable (Donoso, Collins, & Koechlin, 2014).

Role of insular cortex in balancing risk and reward

Research on the role of the insula has consistently implicated this area in tasks that involve risky decision making that typically involves the potential for gains and losses (d'Acremont, Lu, Li, van der Linden, & Bechara, 2009; see Levin, Xue, Weller, Reimann, Lauriola, & Bechara, 2012 for a recent review). Research on the insula has also implicated it as a key component in probabilistic learning in conjunction with other areas that were activated in the current study (e.g., parietal cortex, dmPFC; McGuire et al., 2014). This seems reasonable given the broad spectrum of co-activation observed between insular cortex and cognitive processing areas and its involvement in numerous cognitive disorders (Uddin, 2014). Indeed, previous studies have implicated insular cortex at different stages of the decision making process through the use of event-related designs that titrate the components of risk analyses such as the anticipation of a risky decision (i.e., prospects for success; Furl & Averbeck, 2011), and in the evaluation of the outcome of a risky decision (Paulus et al., 2005). The results of the current study lend further evidence suggesting that the insula is involved in learning and adapting to

change through its conjunction with other regions to code subjective risk and reward values for maintaining a current set of behaviours.

General conclusion

While the areas mentioned above have been separately implicated in different cognitive processes, their combined presence in the current study implicates a broad network of areas subsuming important component processes critical to the development and updating of a mental model. Areas in the frontal and parietal cortices appear to co-operate in the exploration and accumulation of additional evidence in the environment, as well as integrating these experiences into a causal model.

Once built, such a model is only effective as long as the predictions they inform hold true. As the model is tested against the environment, metrics that inform the system of the viability of the current model are persistently gathered. This may include iterative calculations of volatility and subjective ratings of risk and reward given the expected outcome of each decision involving prefrontal cortex and the ventral striatum (Table 2.2). However, as the environment changes, decisions made upon existing models begin to falter. As this continues, areas of cingulate cortex would signal that the prospects of success whilst maintaining the existing belief-set become increasingly worse. Past some threshold that allows for some expected stochasticity in observations (Yu & Dayan, 2005), the existing model becomes invalid, which would signal the resumption of evidence seeking and information integration into a new and updated model.

Although our data showed predominantly more right-hemisphere activation in terms of activated voxels, bilateral activations were also observed (e.g., STG, MTG, IPL, Precuneus). This is consistent with past research using split brain patients suggesting that the left hemisphere

does play a role in integrating the outcomes of previous events (Wolford, Miller, & Gazzaniga, 2000). This suggestion is further bolstered by fMRI examinations of healthy individuals completing prediction tasks that also reveal similar left brain area involvement including the STG alongside right temporal and parietal areas (Miller, Valsangkar-Smyth, Newman, Dumont, & Wolford, 2005). Taken together with the findings of the current study, it is reasonable to suggest that areas in the left hemisphere also contribute to an interpretive process critical to building models of one's environment. However, the specific nature of their contribution, aside from inferences that can be made by interpreting comparable activity in the right brain, has yet to be fully explored.

It is important to consider the results of the current study in the context of its relatively limited sample size. While literature exploring sample size recommendations (Friston, 2012) suggest a sample size of 16 as appropriate, Ingre (2012) asserts that conservative statistical corrections and constraints serve to reduce uncertainty around observed effects in the context of relatively smaller sample sizes. To that end, the current study has employed multiple statistical constraints targeted to reduce uncertainty around observed effects. Namely, the current study employs the use of a contiguous voxel cluster threshold, spatial smoothing, and a relatively conservative statistical threshold of $p < .001$. Additionally, the results of the current study are considered in context with the existing literature and ongoing research in our lab. In addition to limitations around sample size, the discrete nature of these component processes, and the areas that support them, warrant further research to establish the precise context and order in which these activations occur. While the design of the current study was unsuitable to address this question, future research employing an event-related design might. Further, the nature of these areas and their purported responsibilities set out clear predictions for the impact of disruption to

these regions in the context of updating a mental model. Future research using transcranial magnetic stimulation (TMS) or transcranial direct current stimulation (tDCS) could potentially isolate more specifically the roles played by each region, and therefore corroborate their importance in the updating process.

Chapter 3: Working memory correlates of updating².

3.1. Introduction

A great deal of research has examined the ubiquity with which humans model and adapt to regularities in their environment (Saffran et al., 1996, 1999, 2003; Turk-Browne et al, 2005, 2008; Abta & Okanoya, 2009; Green et al., 2010; Tenenbaum et al., 2011; Griffiths & Tenenbaum, 2006; Jueptner et al., 1997; Mayr, 1996; de Gardelle & Summerfield, 2011; Albrecht, Scholl, & Chun, 2012; Fiser & Aslin, 2001; 2002). Accurately modelling such regularities has been shown to be important for a broad range of functions from visual and auditory learning, to language acquisition, motor control, and predictive decision making (Turk-Browne et al, 2005; Green et al., 2010; Jueptner et al., 1997; Saffran et al., 1996; Danckert et al., 2012; Tenenbaum et al., 2011). Critically, in order to construct accurate models of variable and incomplete data sets, we must be capable of integrating information across time (Brown & Steyvers, 2009). Results of Chapter 2 indicate a network of brain regions that are involved in the ability to build and update mental representations that include frontal, parietal, and subcortical areas including the anterior insula and cingulate cortex.

While our prior work in the context of a zero-sum game shows that healthy individuals can indeed adapt to unannounced changes in an opponent's play strategy (Chapter 2; Danckert et al., 2012; Stöttinger et al., 2014a, b), what we don't know from this work, is how other cognitive mechanisms impact this ability. At minimum, playing 'rock, paper, scissors' involves attending to the current play (i.e., attentional resources) and its outcome (i.e., reward signals) over the course of at least some subset of prior plays (i.e., working memory resources; Baddeley, 1992;

² A version of this chapter has been published as Valadao, D. F., Anderson, B., & Danckert, J. (2015). Examining the influence of working memory on updating mental models. *The Quarterly Journal of Experimental Psychology*, 68(7), 1442-1456. It is reproduced here with permission.

Baddeley, 2003). With respect to the latter, in order to successfully exploit an opponent's bias in 'rock, paper, scissors' it would be imperative to maintain some subset of prior plays in working memory (WM) in order to detect the bias in the first instance, and, importantly for our current purposes, to detect any *change* in that bias over time.

There is ample evidence to suggest that WM resources are involved in learning through persistent exposure to spatial regularities; as demonstrated by the poorer performance in contextual cuing tasks when a WM load is imposed (Chun & Jiang, 1998; Annac, Manginelli, Pollmann, Shi, Müller & Geyer, 2013; Manginelli, Langer, Klose & Pollmann, 2013; Travis, Mattingley & Dux, 2013). Furthermore, paradigms exploring the ability to detect abrupt, salient changes are commonly understood to tax WM resources (Rouder et al, 2011). It is not well known, however, what impact a WM load would have on the ability to represent *changes* to information that occur over a longer time scale (i.e., slowly evolving changes to environmental regularities). In addition, while previous studies have investigated the impact of WM loads on learning, further research is needed in order to determine the relationship between the nature of the WM load (e.g., spatial, featural) and the resulting impact on our ability to detect changes in both congruent and incongruent environmental regularities.

The current study explored the effects of WM on our ability to exploit regularities in stimulus properties and to adapt to changes in those regularities – in other words, to build and update mental models. We employed a dual task where participants had to predict either the upcoming location or shape of a target while simultaneously completing a WM task. Target locations and shapes (in separate experiments) were presented with above chance regularity along the dimension of interest (see Methods). For example, target locations were initially drawn from one quadrant of the screen for 20 trials before changing to a different quadrant. Working

memory was then manipulated focusing on either the location, colour or shape of targets using versions of the N-Back task. In this way we were able to manipulate the *relevance* of the contents of WM (spatial vs. featural; Luck & Vogel, 1997) to the prediction task. Changing the underlying distribution of events on the primary prediction task without notice allowed us to examine the interaction of WM and the ability to update representations of regularities. Specifically, we hypothesized that WM load would differentially impact performance on the prediction task dependent upon the congruence (or lack thereof) of the WM task with the demands of the prediction task. Directing WM resources towards the stimulus property that also contained some level of probabilistic regularity (e.g., doing a location based n-back task while predicting impending target locations), ought to improve the ability to detect that regularity and adapt to any changes.

3.2. Experiment 1: Method

Participants

A total of 97 University of Waterloo undergraduate students (33 Male; mean age 19.9, $SD \pm 2.7$ years) completed the first experiment. All participants had normal or corrected to normal vision and were right-handed by self-report. Participants were instructed to use their right hand throughout the experiment. Informed consent was obtained prior to commencing the experiment and all procedures were approved by the University of Waterloo, Office of Research Ethics.

Apparatus

Stimuli were presented on a 19" NEC AccuSync 90 monitor with a resolution of 1024 by 768 at a refresh rate of 60Hz using a GeForce 4 Mx440 graphics processing unit. Button responses were captured via a Cedrus RB-530 Response Pad. Pointing behaviour was captured using Elo Intellitouch touch screen capture technology. Stimuli were created using Adobe Photoshop CS5 and the experimental protocol was created and administered using E-Prime 1.1 on Windows XP.

Procedure

Participants were seated approximately 57cm from a touch screen display. The experiment began with a touch screen calibration task, both to ensure that the touch screen was calibrated to the participant's responses, and to habituate participants to the use of the touch screen.

Each experimental trial consisted of two parts – a spatial prediction component and a WM component. To generate the spatial prediction component of the trial, participants were first shown a target stimulus. After 1250-1500 msec the target disappeared and the screen changed colour. This signaled to the participant that they were to predict – by pointing on the touch screen – the location they thought the *next* target would appear (Figure 3.1). Of course, in early trials such predictions amount to guesswork. Unbeknown to the participant, target locations were drawn from a specific spatial distribution centered around one quadrant of the screen ($SD_{x,y} = 60$ pixels) allowing participants to develop a representation over time that would more accurately predict upcoming spatial locations. After 20 trials the distribution shifted to a different quadrant

of the screen, requiring participants to detect the change in probabilistic regularity of target locations and update their representation in order to optimize predictions. Participants were not informed of the probabilistic regularity of the target locations or the shift in those regularities.

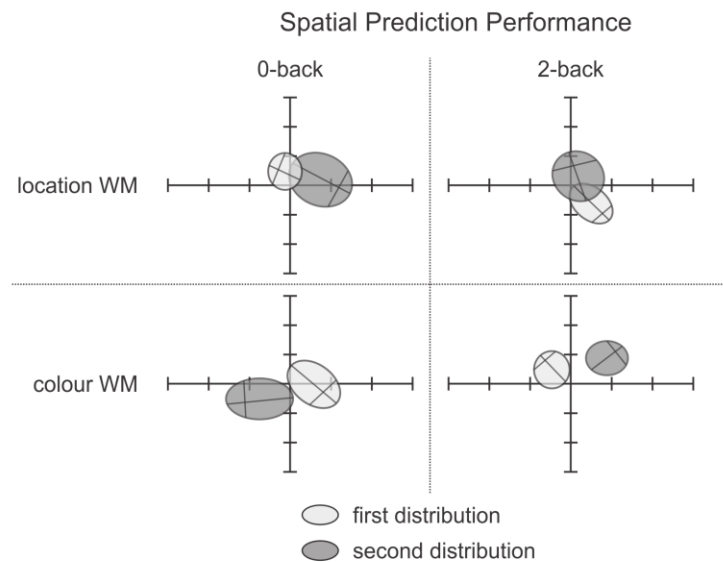
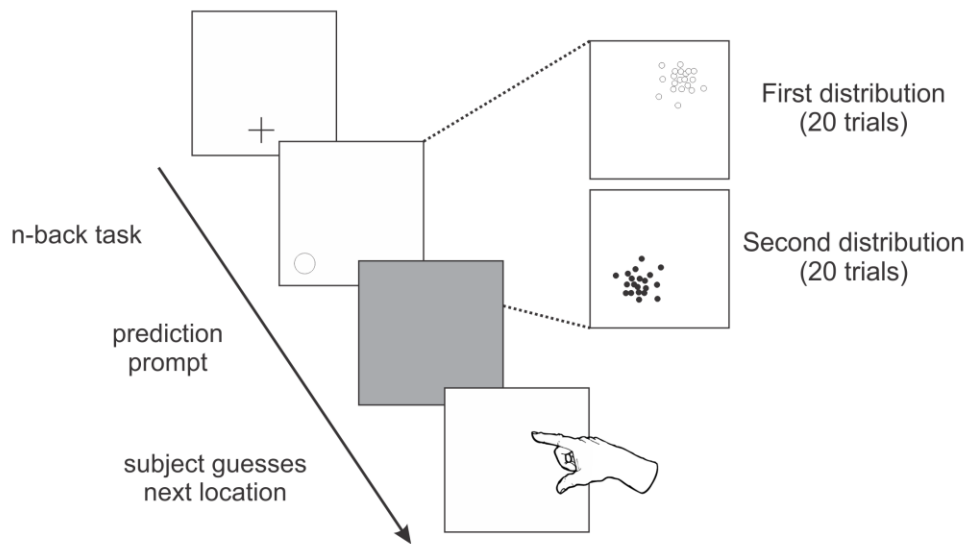


Figure 3.1. Working Memory Prediction Stimuli. Top depicts a schematic of the spatial prediction task in Experiment 1. After a brief fixation cross, an object would appear somewhere on screen. Participants made a button press response relative to the specific WM load (i.e., 0-back or 2-back) they were currently performing. The object would then disappear and the screen would change color to inform them of the accuracy of their button press response. This also acted as a cue for participants to touch the screen indicating their prediction for where the *next* object was most likely to appear.

While the first component of each trial examined how efficiently participants could predict (and update following a shift) a representation of the spatial locations of targets, the second component explored what impact, if any, WM had on their ability to complete this task. On each trial, participants completed a WM task with one of two load levels (2-back vs. 0-back), focusing on one of three specific stimulus properties: colour (red, green, blue, or yellow), shape (circle, square, star, or triangle), or location (the four quadrants of the screen; Figure 3.2). Participants completed six blocks of trials with each block consisting of one level of WM load (0 or 2-back) for each target property (location, colour, shape). For this component, while the target stimulus was on screen, participants responded by pressing a button on the touch pad if the stimulus met a specific criterion. If the stimulus did not meet the criterion, they were to withhold their response and wait for the target to disappear. For the 0-back trials, participants responded if the object was a particular colour, shape, or appeared in a specified quadrant of the screen; the particular criterion varied across different blocks. For the 2-back WM trials, participants responded by pressing a button when the current stimulus matched the stimulus two trials ago for either shape, colour, or screen quadrant. Consistent with both levels of the WM load manipulation, the target criterion (i.e., shape, colour, location) was varied across blocks of trials, but remained consistent within each block. In designing the experiment, it was possible to include different levels of the N-back task to explore the differential impact of WM difficulty per se on detecting changes in regularity. However, our primary goal was to explore the effect of WM *modality* and its congruence with the prediction task on participants' ability to adapt to changes in probabilistic regularity. We chose the 2-back for the WM load conditions as it provided an optimal balance of difficulty when combined with the spatial prediction task. That

is, we felt that a 1-back WM load would not be sufficiently difficult to expose differential influences of WM modality on detecting changes in regularities, whereas a 3-back load would be too difficult, leading all conditions to exhibit a floor effect. The 0-back detection task was included as a dual-task control with no WM load. While it was possible to include additional levels of the n-back task, we chose to limit the number of conditions in the study to limit the length of the task to a manageable level to minimize the possibility that fatigue and low motivation could account for our results.

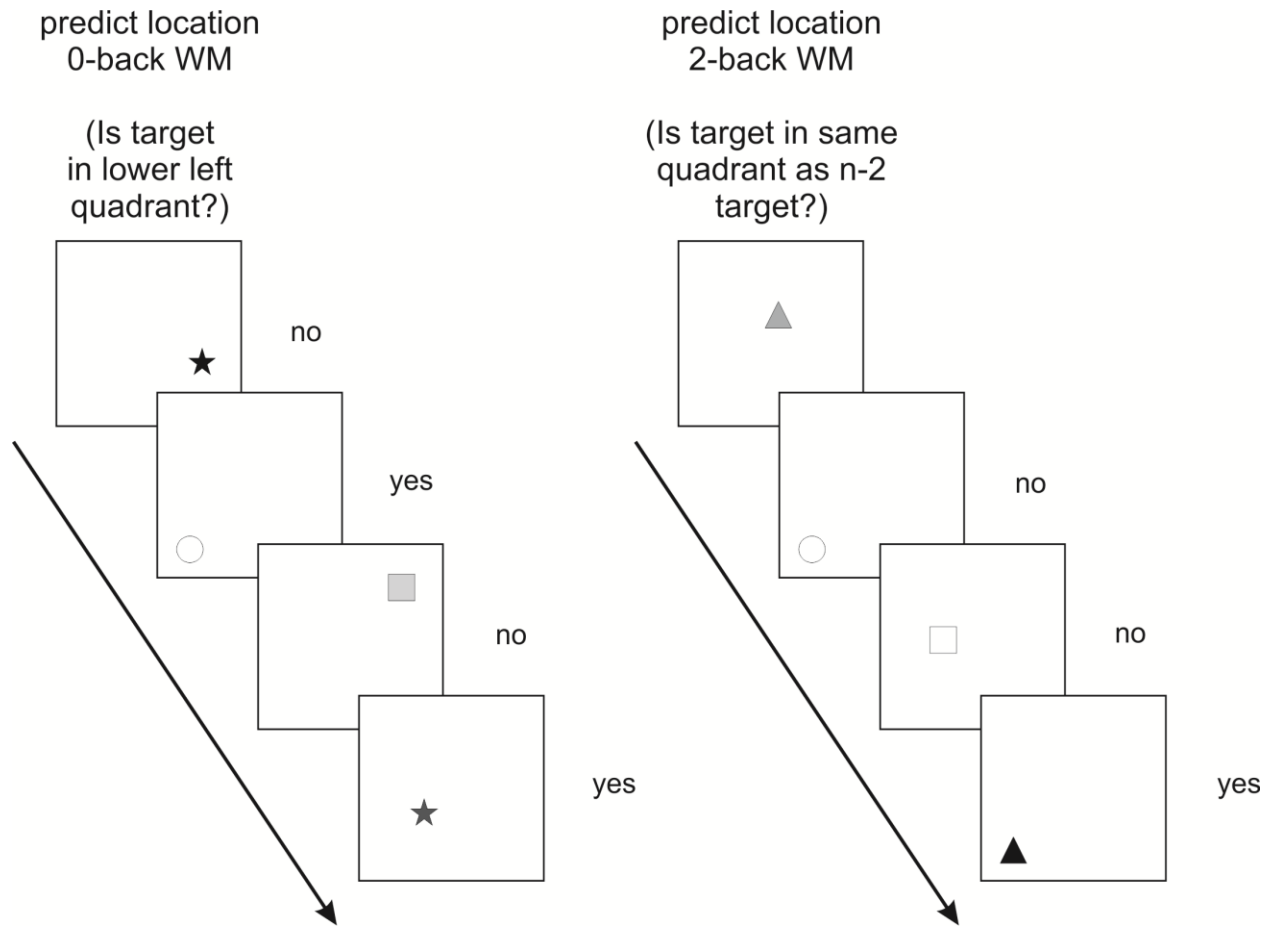


Figure 3.2. 2-Back Task. Depicts a schematic of the different working memory tasks. The tasks could focus participants on one of three features of the object (i.e., its location as depicted, or its color or shape) in each block of trials. Sample button press responses adjacent to the objects indicate the correct response for that object in that block of the experiment (the arrow indicates temporal order of target presentation).

After an initial practice phase consisting of 20 trials with both levels of the WM task, participants moved on to the experimental trials. At the beginning of each experimental block, participants were prompted as to which task they would be completing (e.g., 2-back; same shape). If participants responded correctly, the screen would turn green to indicate their button response was correct; the screen turned red for incorrect responses. Examination of WM performance across the current study indicates that participants were capable of completing the 2-back task, with a mean accuracy of 76%. Notably, 2-back accuracy was somewhat lower in the location condition relative to colour and shape conditions for Experiments 1, $F(2,288) = 103.2$, $MSE = .02$, $p < .001$, and 2, $F(2,69) = 5.05$, $MSE = .02$, $p < .01$, but not Experiment 3, $F(2,72) = 2.29$, $MSE = .01$, $p = .11$ (Figure 3.3). Irrespective of their performance on the WM component of the task, participants then had to predict the *location* of the next object by touching the screen to indicate their prediction. In other words, the spatial prediction task was constant across all variants of WM load and WM target type.




Figure 3.3. N-Back Task Accuracy. Depicts dual-task N-Back accuracy as a function of WM load across all experiments.

The experiment was performed within-participants and the order of blocks was counterbalanced across participants. We predicted that participants would perform better at the spatial prediction task when simultaneously completing a WM task that focused their attention on the information that was relevant to the statistical learning task (i.e., the location of the objects), as opposed to information that was irrelevant to the task (i.e., the stimulus colour or shape).

Data Analysis

Accuracy on the spatial prediction component of the task in Experiment 1 was computed by calculating the Euclidian distance between the participant's prediction on each trial and the known mean of the distribution from which the objects were drawn (i.e., $\sqrt{error_x^2 + error_y^2}$), also known as radial error). For the remaining two experiments, accuracy for shape prediction was characterized as a proportion of how often the participant correctly predicted the object most likely to appear on subsequent trials (optimal prediction).

We first explored spatial prediction performance in the first 20 trials in each experimental block (i.e., prior to any change in the distribution of target locations) to establish a behavioural baseline from which it would be possible to evaluate how well participants were able to detect a *change* in the distribution of target locations. To do this we calculated the individual accuracy of each participant for the first 20 trials of each initial target distribution (Figure 3.1). Non-linear curves were then fitted to the data to investigate whether a significant improvement in prediction error was evident for the first distribution. A significant curve fit would demonstrate that participants had become more accurate in their predictions demonstrating some degree of learning of the spatial distribution of targets of the first distribution (i.e., over the first 20 trials).

Analysis of these data revealed that participants' clearly altered their predictions to reflect the nature of the first distribution of targets in both the 0-back, $b = 0.106$, $F = 21.92$, $p < .001$, and 2-back, $b = 0.084$, $F = 13.67$, $p < .001$, conditions (Figure 3.4). Having shown that participants accurately learned the first distribution we next explored the improvement in participants' ability to detect the *change* in the distribution of target locations using the same measure of accuracy as described for the first distribution. This allowed us to quantify how well participants' were able to update their representation of object locations as a function of different WM load conditions. All data were analyzed using SPSS Statistics version 20.

3.3. Experiment 1: Results

Analysis of prediction errors after the switch in target distributions indicated that individuals were capable of effectively detecting the change in regularity of target locations when no WM load was present (i.e., 0-back conditions; Figure 3.4). Inverse curve estimations showed significant learning in the 0-Back conditions when participants focused on the object's location, $b = 0.201$, $F = 81.70$, $p < .001$, color, $b = 0.072$, $F = 10.05$, $p < .01$, and shape, $b = 0.065$, $F = 8.10$, $p < .01$ (Figure 3.4).

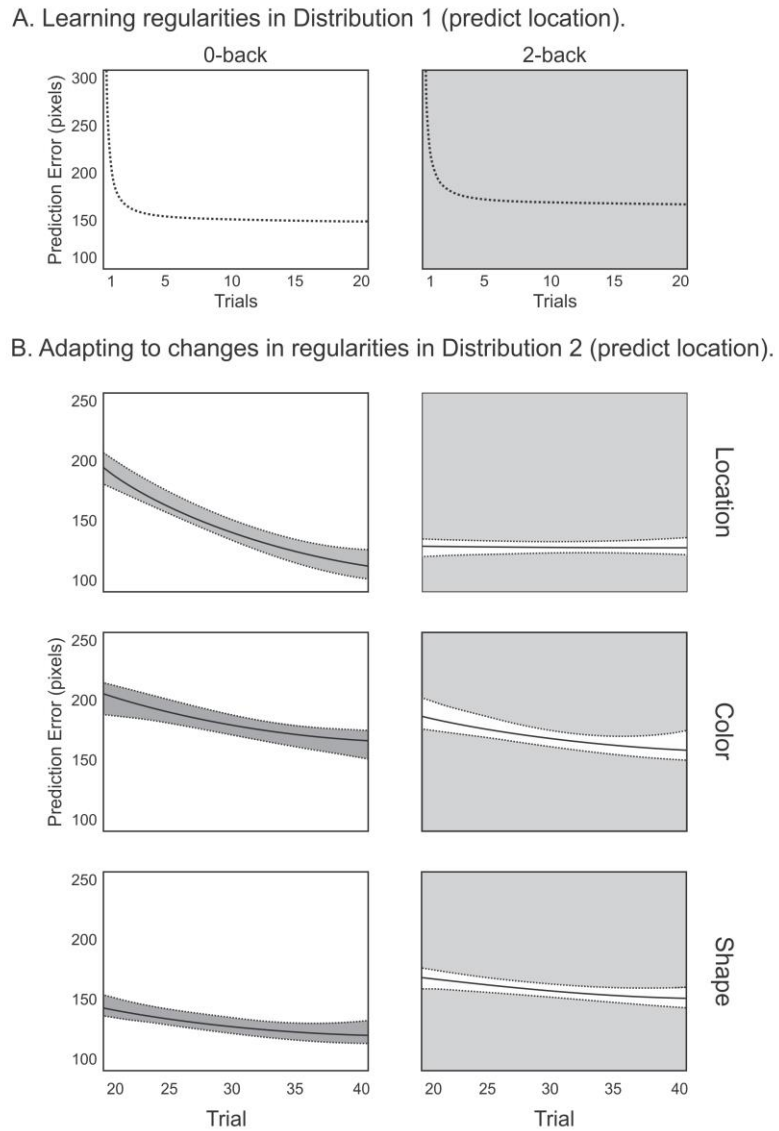


Figure 3.4. Experiment 1 Location Prediction Data.

A. Depicts improvement in radial prediction error (relative to known distribution means) for the first 20 trials prior to a switch in object location distribution as a function of concurrent WM load (i.e., 0-back in white, and 2-back in gray). B. Depicts best fitting curve estimations for change in radial prediction error after the switch of object location distribution occurred across all experimental conditions (i.e., trials 21-40).

Under the different WM loads (i.e., the 2-back conditions) participants only showed improvement in learning for the colour, $b = 0.056$, $F = 6.02$, $p < .05$, and shape, $b = 0.053$, $F = 5.60$, $p < .05$, conditions (Figure 3.4). However, when the 2-back WM task shared the same property as the prediction task (i.e., target location) participants showed no significant improvement in prediction, $b = 0.003$, $F = 0.01$, $p = .91$, (Figure 3.4). Directly contrasting improvements in prediction across the 0-back and 2-back WM conditions showed a significant difference for the location task only, $t(96) = 6.97$, $SE = 436.57$, $p < .001$. That is, when focusing on location there was a strong learning trend in the 0-back condition that was rendered essentially flat in the 2-back condition (Figure 3.4). A graph depicting the raw prediction behavior of our sample from a subset of conditions in our dataset is shown in Figure 3.1. Visually, this shows a larger degree of overlap in prediction behaviour when the WM task targeted the object location, further suggesting some difficulty adapting to changes in distributions under congruent WM loads. Taken together, the above results suggest a significant effect of congruent WM load on learning the second distribution.

Table 3.1. Inverse Learning Curve Estimation Statistics.

	<i>Distribution 1</i>	b	Fit	F _{sig}	<i>Distribution 2</i>	b	Fit	F _{sig}
Exp. 1	0-Back				0-Back			
	Combined	.106	21.92	<.001	Location	.201	81.70	<.001
					Color	.072	10.05	<.01
					Shape	.065	8.10	<.01
	2-Back				2-Back			
	Combined	.085	13.67	<.001	Location	.003	0.01	.91
				Color	.056	6.02	<.05	
				Shape	.053	5.60	<.05	
Exp. 2	0-Back				0-Back			
	Combined	.103 [†]	5.12	<.05	Location	-.109	5.75	<.05
					Color	-.097	4.55	<.05
					Shape	-.127	7.87	<.01
	2-Back				2-Back			
	Combined	-.099	4.75	<.05	Location	-.072	2.52	.11
				Color	-.094	4.27	<.05	
				Shape	.035	0.56	.44	
Exp. 3	0-Back				0-Back			
	Combined	-.158	12.79	<.001	Location	-.161	13.30	<.001
					Color	-.129	8.40	<.01
					Shape	-.229	27.47	<.001
	2-Back				2-Back			
	Combined	.09 [†]	4.09	<.05	Location	-.101	5.16	<.05
				Color	.023	0.26	.61	
				Shape	-.049	1.18	.28	

[†]Where indicated, inverse curve estimation did not provide good fit for these data. Therefore, linear estimation statistics, which provided a much better fit in these conditions, are provided.

3.4. Experiment 2: Introduction

The first experiment demonstrated that a spatial WM load *impairs* the ability to update a representation of the spatial distribution of objects when that distribution changes. However, the nature of the distributions used was somewhat problematic. Specifically, to avoid having the 2-back task be overly simple, we presented targets near the border of screen quadrants (Figure 3.1). That is, we felt that distributions clearly contained within a given quadrant would be too easy to represent. This significantly limited our ability to switch the distribution of object locations (i.e., the new distribution also had to have targets appearing close to quadrant borders). Therefore, while we did see a reduction in learning rate during the location-based WM load condition relative to its no WM load counterparts, it is possible that a ceiling effect accounted for this difference. It was therefore important to replicate our findings using a different distribution of object features that could be varied without being constrained by the nature of the WM task. For the remaining experiments, we used shape as the feature participants had to predict on each trial. Since objects can be represented in WM both in terms of spatial and featural components (Luck & Vogel, 1997), we expected to replicate our findings when we now explored featural processing. That is, we expected prediction performance for the second distribution of shapes to be worst when WM load also involved shape processing.

3.5. Experiment 2: Method

Participants

20 University of Waterloo undergraduate students participated in the second experiment (7 Male; mean age 20.6, SD +/- 2.1 years).


Procedure

The objectives, structure, and stimuli for Experiment 2 were identical to Experiment 1 save for one key difference. Instead of predicting future target locations, participants were asked to predict the *shape* of upcoming targets and indicate their prediction on the touch pad. As with Experiment 1, in the first 20 trials, one shape was more likely to appear than the other shapes (e.g., squares appeared on 70% of trials, with every other shape appearing 10% of the time). After these first 20 trials, the most likely target object was changed unbeknown to participants (e.g., circles appeared on 70% of trials, with every other shape appearing 10% of the time). Accuracy was computed as a proportion of trials in which participants correctly identified the most likely future shape to be the one that was actually most likely to appear. Trial-by-trial accuracy for each participant was submitted to non-linear curve estimation to examine the relative improvement in prediction accuracy as a function of WM condition. Data on participants learning improvement as a function of working memory load were analyzed in the same manner as Experiment 1.

3.6. Experiment 2: Results

Analysis of learning the first distribution (see Experiment 1) revealed significantly improved prediction performance for targets in the first distribution over the first 20 trials in both the 0-back and 2-back conditions (Table 3.1). As with Experiment 1, under no WM load, improvement in shape prediction accuracy occurred for location, ($b = -0.109$, $F = 5.75$, $p < .05$), color, $b = -0.097$, $F = 4.55$, $p < .05$, and shape conditions, $b = -0.127$, $F = 7.87$, $p < .01$, (Figure 3.5).

A. Learning regularities in Distribution 1 (predict shape).



B. Adapting to changes in regularities in Distribution 2 (predict shape).




Figure 3.5. Experiment 2 Shape Prediction Data.

A. Depicts improvement in shape prediction accuracy for the first 20 trials prior to the switch in shape distribution as a function of concurrent WM load (i.e., 0-back in white, and 2-back in gray). B. Depicts best fitting curve estimations for change in shape prediction accuracy after the switch of object shape distribution occurred across all experimental conditions (i.e., trials 21-40).

Under the context of a WM load, participants demonstrated a significant amount of learning in the color condition, $b = -0.094$, $F = 4.27$, $p < .05$. However, participants showed no improvement in prediction accuracy when the WM task focused on the object's location, $b = -0.072$, $F = 2.52$, $p = .11$ or its shape, $b = 0.035$, $F = 0.56$, $p = .44$ (Figure 3.5). Critically, the largest difference in learning curves was observed when comparing the 2-back and 0-back conditions that focused on the target feature *congruent* with the prediction task (i.e., shape). As with Experiment 1 (i.e., location), these data indicate a significantly impaired ability to detect changes in regularity of object shapes when a congruent WM load (i.e., shape) was applied, $t(19) = 2.53$, $SE = 4.63$, $p < .05$.

3.7. Experiment 3: Introduction

To keep Experiment 2 procedurally similar to Experiment 1 we continued to vary the spatial location of targets as we had in Experiment 1 (Figure 3.1). That is, not only did the task include regularity in the shape information of successive targets, there was also redundant regularity in the spatial information conveyed by successive targets (i.e., not only were targets more predictably of one shape or another, they were also more likely to appear in one quadrant or another). Given that we have shown previously that redundant spatial regularities can aid updating of mental models (Filipowicz et al., 2014), we felt it was important to examine shape prediction performance with spatially redundant regularities removed from the task. This was also considered important in light of prior research showing that WM load is decreased by redundancy and regularity (Brady et al., 2009). Therefore, we conducted an experiment in which the most likely shape appeared 70% of the time as in Experiment 2, but with locations of target objects drawn at random from all possible locations.

3.8. Experiment 3: Method


Participants

20 University of Waterloo undergraduate students participated in the third experiment (2 Male; mean age 19.8, SD +/- 1.1 years).

3.9. Experiment 3: Results

Analysis of learning the first distribution (see Experiment 1) showed significant improvement in the first 20 trials (Table 3.1), indicating that participants improved their prediction performance for targets in the first distribution over the first 20 trials. Consistent with the first two experiments, participants demonstrated the ability to effectively detect changes in the regularity of object shapes after an unannounced shift in the most likely shape occurred. This was true regardless of whether the 0-back task focused participants on location, $b = -0.161$, $F = 13.30$, $p < .001$, color, $b = -0.129$, $F = 8.40$, $p < .01$, or shape, $b = -0.229$, $F = 27.47$, $p < .001$ (Figure 3.6).

A. Learning regularities in Distribution 1 (predict shape).



B. Adapting to changes in regularities in Distribution 2 (predict shape).




Figure 3.6. Experiment 3 Shape Prediction Data.

A. Depicts improvement in shape prediction accuracy for the first 20 trials prior to the switch in shape distribution as a function of concurrent WM load (i.e., 0-back in white, and 2-back in gray). B. Depicts best fitting curve estimations for change in shape prediction accuracy after the switch of object shape distribution occurred across all experimental conditions (i.e., trials 21-40).

Under a modest WM load, we see learning demonstrated only in the location condition, $b = -0.101$, $F = 5.16$, $p < .05$. However, participants demonstrated no learning under the color, $b = 0.023$, $F = 0.26$, $p = .61$ and shape WM conditions, $b = -0.049$, $F = 1.18$, $p = .28$ (Figure 3.6). Furthermore, consistent with the first two experiments, we again demonstrate a significant impairment in shape prediction improvement when resources are focused on a *congruent* WM task (i.e., 2-Back shape) relative to its no-load control condition, $t(19) = 2.87$, $SE = 4.54$, $p < .01$.

As with previous experiments, we see a marked decrease in learning the new most likely shape after the shift when WM resources were focused on the object's location, $b = -0.101$, $F = 5.16$, $p < .05$, color, $b = 0.023$, $F = 0.26$, $p = .61$, or shape, $b = -0.049$, $F = 1.18$, $p = .28$ (Figure 3.6), relative to the 0-back conditions that placed no WM load on participants. Indeed, comparing learning curves directly indicate that participants are significantly poorer at improving prediction accuracy when WM was focused on features *congruent* with the prediction task (i.e., shape), $t(19) = 2.87$, $SE = 4.54$, $p < .01$.

3.10. Discussion

Results of the current study indicate that WM moderates the efficiency with which participants were able to detect and exploit a change in the regularity of their environment – in other words, updating a mental model. The reduced capacity for detecting change under a moderate WM load was modality specific. All experiments showed a consistent interaction between WM content and prediction performance such that a WM load that was *congruent* with the content of the prediction task was most detrimental to the ability to represent changes to regularities in the prediction task.

The purpose of mental models is to optimize predictive behavioural control given incomplete and noisy inputs (Howhy, 2012). The success of any given mental model depends on the predictive accuracy of the model. When a model's predictions cease to lead to optimal behavioural outcomes we need to either fine tune the model given new information, or abandon it altogether in favour of a new model. The determination of the accuracy of a mental model depends on a process of comparing current observations with expected, or predicted outcomes generated by the model. What the current data suggest is that this comparative process to determine the efficacy of a mental model requires free WM resources. While we expected that focusing WM resources on the same property in which a change in regularities was to be detected would be beneficial to updating, the data showed the opposite. It was as if participants could not see the forest for the trees. Focusing on single instances for the purposes of the WM task obscured the ability to detect trends evolving over longer time scales (i.e., 20 trials; Figure 3.4).

The current results concur with other research showing that WM loads interfere with our ability to attend to regularities (Annac et al., 2013; Manginelli et al., 2013; Travis et al., 2013) and detect salient changes (Rouder et al., 2011). Indeed, prior research has shown that individual WM capacity correlates with the capacity to learn sequences (Unsworth & Engle, 2005). Here we have extended these findings in three important ways. First, as discussed above, WM load interferes with the ability to detect slowly evolving changes (i.e., prior work examined only the influence of WM on detecting singleton events; Rouder et al., 2011) or repeated exposures to identical spatial configurations of stimuli (Annac et al., 2013; Manginelli et al., 2013; Travis et al., 2013). Second, the current results show that the influence of WM on detecting changes to regularities was modality specific (Figure 3.4). Specifically, prediction performance was only

disrupted when the WM task shared the same feature as the prediction task (e.g., remembering shape while predicting future shapes). Prior work has focused primarily on the impediment of spatial/configural learning due to spatial (but not featural) WM loads. Finally, the n-back nature of the WM manipulation in the current study resulted in a persistent and maintained use of WM resources across learning trials, whereas the WM loads of previous tasks did not have this maintenance built in to their manipulation. Interestingly, prediction performance while completing the two-back and the zero-back memory control tasks was similar when the features of the objects participants were required to remember differed from the target feature being predicted. This result may also suggest that the regularity in stimuli may have eased WM demands under certain circumstances, corroborating the notion that probabilistic regularity allows for compression of WM resources (Brady et al., 2009).

In summary, the current study examined the relationship between working memory and the ability to detect changes in the environment in order to update an existing mental model. Results indicated that free WM resources are needed to effectively appreciate statistical regularities (i.e., increased efficiency under no-load conditions) and that modality specific WM tasks hindered the ability to detect *changes* in statistical regularities. These results establish at least one cognitive constraint acting on our ability to accurately represent regularities in the environment and changes to those regularities.

Chapter 4: Attentional Influences on Building and Updating Mental Representations

4.1. Introduction

Chapter 2, examining the brain regions involved in building and updating mental models, highlighted an important role for parietal cortex during instances where individuals successfully modeled and exploited biases in a computer's play strategy. Notably, parietal cortex is typically considered important in governing the allocation of attentional resources (Wojciulik & Kanwisher, 1999; Hopfinger et al., 2000; Ferber & Danckert, 2006; Heilman & Valenstein, 2011) and decision making processes (Straube et al., 2011, Paulus et al, 2004). Despite these findings, the role attention plays in facilitating the model building process is relatively unknown.

Indeed, attention is a natural candidate for further exploration on its own merits. Attention is commonly understood as an important mechanism for facilitating information processing by selecting or orienting to relevant or salient information in the environment (Corbetta & Shulman, 2002). Attention can be oriented both volitionally towards goal-relevant stimuli in accordance with expectations (i.e., endogenous cueing; Posner, Snyder & Davidson, 1980) or can be captured by salient events in the environment (i.e., exogenous cues; Müller & Rabbitt, 1989; Yantis & Jonides, 1990). However, additional research has begun to demonstrate that attention can be oriented not just by the presence of visually salient stimuli, but by visual and temporal regularities in the environment.

Indeed, there are a number of studies that demonstrate a relationship between visual or statistical regularities and attentional biases. Zhao and colleagues (2013) demonstrated that the presence of statistical regularities in visual stimuli biased the allocation of attention such that participants were quicker to correctly identify search targets interleaved within structured vs. random arrays of objects (temporal regularities within streams of objects provided the

‘structure’; Zhao, Al-Aidroos, and Turk-Browne, 2013). This is consistent with other research which demonstrates that probabilistic regularity in the visual environment enhances the allocation of attention in visual search tasks such that frequently repeated search scenes lead to much faster identification and classification of target orientations (e.g., contextual cueing; Chun & Jiang, 1998; Geng & Behrmann, 2002; Rausei, Makovski, and Jiang, 2007; Jiang, Swallow, and Rosenbaum, 2012; Jiang, Swallow, Won, Cistera, and Rosenbaum, 2015). Additionally, regularity in the spatial or temporal presentations of visual stimuli has been shown to facilitate perceptual classification tasks (Druker & Anderson, 2010; Shaqiri & Anderson, 2012) and associative prediction tasks (Turk-Browne, Scholl, Johnson, and Chun, 2010).

Taken together, this research clearly demonstrates that the presence of probabilistic visual regularities in the environment affects the allocation of attention. However, it is unclear whether the relationship works in the opposite direction. Namely, whether the explicit allocation of attention can improve one’s ability to build a mental representation of the regularities they observe. For example, while some suggest that attentional resources are a critical component of learning processes such as sequence learning (Nissen & Bullemer, 1987; Kabata, Yokoyama, Noguchi, and Kita, 2014), more recent research has demonstrated that sequences can be learned equally well in the context of impoverished attention resources (through the use of a dual task) as they can be in circumstances allowing full allocation of attentional resources (i.e., no dual task; Stadler, 1995). The influence of attention on statistical learning is equally controversial. Classically defined as the ability to passively observe probabilistic regularities between events, statistical learning is seen as a critical component of early lexical segmentation in the development of human language (Aslin, Saffran, and Newport, 1998; Saffran, Johnson, Aslin, and Newport, 1999). While Saffran and colleagues argue that such learning can occur without

explicit awareness of the translational probabilities, such as when they had individuals listen to streams of speech while completing a distractor task (Saffran, Newport, Aslin, Tunick, & Barrueco, 1997), others argue that such statistical learning is impossible without freeing up attentional resources. For example, Toro and colleagues had participants either listen to speech streams with predictable translational probabilities in isolation or, in one instance, complete a detection task while listening to the same speech stream (i.e., responding to repetitions; Toro, Sinnet, and Soto-Faraco, 2005). In all cases, participants performed worse at the statistical learning task under the context of depleted attentional resources. However, in a similar study, Turk-Browne and colleagues presented participants with streams of visual shapes that formed regular triplets, with two groups corresponding to distinct colours. Participants viewed these groups of shapes interleaved by colour and were asked to detect repetitions of green shapes (thereby forcing participants to focus on only one colour of shapes). Results showed significantly increased learning of the triplet organization of attended relative to unattended shapes, again indicating that attention affects statistical learning (Turk-Browne, Jungé, and Scholl, 2005).

What the above discussion highlights is that there remains some controversy concerning the influence of attention on our ability to represent regularities in the environment (i.e., build mental models). Importantly, none of these studies explored the influence of attention on updating mental representations when regularities change. While some argue that mere exposure to statistical regularity in the environment is sufficient, others demonstrate that directing attentional resources can have *either* a facilitative (Turk-Brown et al., 2005) or deleterious (Toro et al., 2005) effect on the learning process. Furthermore, to the extent that attention is an important component of the ability to build mental representations, it is unknown whether volitionally allocating attention will improve the model building process or whether attention

serves as a gating mechanism such that attentional constraints impair the process (Mangun, 1995; Fischer & Whitney, 2012). The purpose of the current study was to explore the role of attentional resource allocation on the ability to build a mental representation of visual regularity in one's environment.

4.2. Experiment 1: Introduction

In order to examine the influence of attention on model building, it was first necessary to establish that participants could indeed represent two conflicting data constructs simultaneously. Therefore, the first experiment served as a behavioural baseline and proof of concept to establish our experimental protocol and explore the possibility that participants could later benefit above a known baseline of performance when they are eventually asked to allocate attention to one aspect of the environment over another.

4.3. Experiment 1: Method

Participants

A total of 26 undergraduate students ($M = 22.2$ years, $SD = 3.6$ years) from the University of Waterloo participated in the current experiment. All participants reported normal or corrected-to-normal vision and gave written consent to participation. As remuneration for their involvement, participants received credit towards one of their Psychology courses that term. The experiment and all procedures were approved by the University of Waterloo, Office of Research Ethics.

Apparatus

Stimuli were presented on a 19” monitor with a resolution of 1024 by 768 at 60Hz. Participants used a mouse to input their responses throughout the task. The task itself was created using Psychopy, a superset of the python programming language developed for the purpose of creating experimental stimuli and was administered on a computer running XFCE Linux (Peirce, 2007).

Procedure

The experimental stimuli consisted of a visual analogue of the “Plinko” game featured on the commonly known daytime TV game show “The Price is Right.” Participants would observe a coloured ball fall from the top of the screen through an array of pegs before falling into one of the 40 slots at the bottom of the screen (Filipowicz, Valadao, Anderson, & Danckert, 2014; Figure 4.1). While the nature of the experimental protocol imparted a sense of randomness to the eventual destination of each ball, the eventual landing spot of each ball was pre-ordained to conform to a canonical distribution that was created prior to the experimental run. These distributions had a known mean and variance and created an array of eventual landing slots for the experiment. In the original iteration of the Plinko task, participants would first indicate their “prior” by drawing what they thought the distribution would look like before observing ball drops. Afterwards, participants would observe a single distribution of ball drops while making estimations on every trial as to what their internal model of the true distribution looked like (Figure 4.1). For the purpose of this chapter, the original task was extended to include two distributions. Participants always saw ball drops conforming to two source distributions grouped by colour (red and green balls). In the current experiment, ball drops were interleaved such that

only one ball (and hence only one colour) was dropped on each trial. Participants were asked to pay attention to both colours of balls equally as they fell so as not to impart any advantageous use of attention at this time.




Figure 4.1. Original Plinko Task. Participants observed balls dropping from the top of the screen through an array of virtual pegs into one of 40 landing slots. Participants were free to estimate their representation of the ball drops on each trial (but could also choose to leave their estimate unchanged from trial to trial).

The primary objective of each experimental run was for participants to observe these ball drops and, at 35 trial intervals, input their estimations of what the true distributions looked like. This led to a total of 4 estimates for each distribution in each condition and marked a change from prior iterations of this task where participants were free to adjust their estimations on every trial. The change was deemed necessary to ensure exposure to many more ball drops in a shorter time span than would be possible if they were stopped to input data every trial. Participants completed this task by allocating a height to each of the 40 bars underneath the landing spots on the Plinko stimuli such that the respective heights of each bar would form their internal model of the distribution. That is, a taller bar represented a higher probability of balls falling there, a smaller bar represented a lower expected probability and no bar at all indicated that participants thought balls would never fall in that slot. As in all experiments in the current chapter, participants were asked to estimate both the red and green distributions (Figure 4.2). Participants would be prompted to input their estimate for the appropriate colour distribution with a coloured text prompt at the bottom of the screen. Additionally, the bars themselves were coloured so that participants knew which distribution they were estimating at all times (Figure 4.2). Whenever the program paused to prompt participants for data, they were given prompts for both the red and green distribution, with the order in which those prompts were given randomized.




Figure 4.2. Experiment 1 Plinko Task. After observing alternating ball drops according to a red and a green distribution for 35 trials, participants would draw what they thought both red and green distributions looked like when prompted.

Each run of the experiment consisted of four pairs of distributions (i.e., red and green in each pair). The distributions themselves, from which ball drop landing positions were drawn, were pre-generated during the design phase and were split into two conditions. The first condition consisted of distribution pairs in which the means for each distribution were separated

from each other by 0.5SD, 1.0SD, and 1.5SDs, including a final pair that shared the same mean but different variance (i.e., wide v narrow Gaussians). The second condition consisted of distribution pairs that overlapped with each other by 22%, 44%, and 66%, followed by a final pair that shared the same mean but difference variance (i.e., the same wide v narrow Gaussian pair as the other condition). The purpose of these two conditions was to determine whether participants' capacity to represent each individual distribution differed as a function of the degree of similarity and whether any such differences could be addressed by allocation of attentional resources in subsequent experiments. The two distinct conditions were administered between subjects, while distribution pairs for each condition were counter-balanced within-subjects.

Data Analysis

Participants' data were first pre-processed to normalize the height of their bar estimations in each given slot to a percentage of the sum total of their bar heights across all slots. This was necessary to ensure that the relative heights of the bars were equal across participants, preventing individuals who had a tendency to give all bars a stereotyped baseline height a disproportionate amount of bias on the data relative to an individual who left bar heights relatively low. This same normalization process was done for the source distributions such that each of the 40 slots had a percentage for the participants estimate ($P.Slot_i$) and the true percentage of ball drops coming from the source distribution ($D.Slot_i$).

From there, participant's accuracy for estimating the source distribution was computed according to the following formula which sums the lower of the two percentages of each slot.

$$\sum_{i=1}^{40} \text{MIN}(P.\text{Slot}_i | D.\text{Slot}_i)$$

This creates an accuracy value that penalizes relative under- and over-estimation of the source distribution across each slot, such that a value of 0 would be returned if the participant made absolutely no slot estimates in a region where the distribution existed and would approach a value of 1 as a participant's slot estimates began to match the exact distribution slot percentages (Figure 4.3).

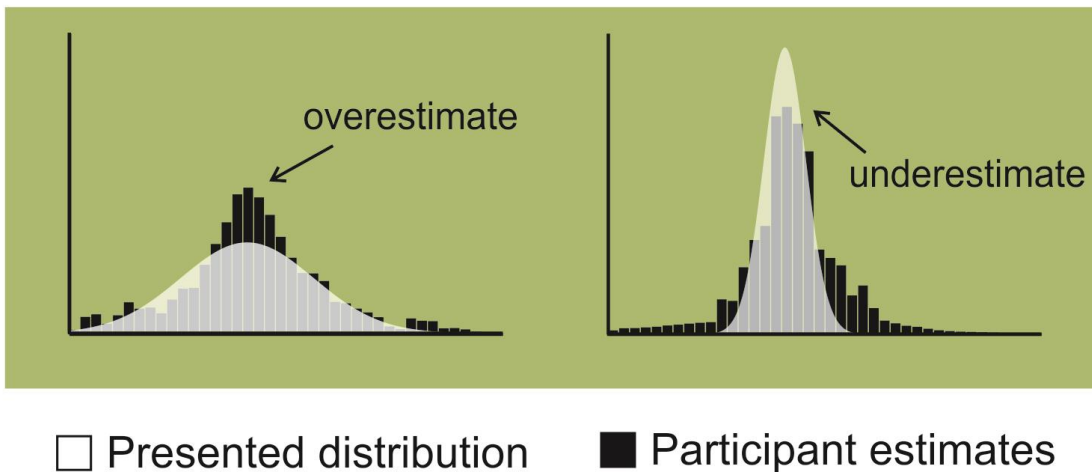


Figure 4.3. Accuracy Estimation Example. Participant's estimates are depicted by solid bars while source distribution is depicted by a transparent white curve. Accuracy is calculated according to the minimum overlap between the two curves and then converted into a percentage of the source distribution's surface area. A maximum accuracy score is only achievable by precisely estimating both the mean and variance of the source distribution. Left depicts an example of overestimation that would reduce overall accuracy while the right depicts an underestimation error.

Participant's accuracy for each of the eight distributions (i.e., four pairs of red and green distributions) were averaged across all four trials where they made their respective estimates (i.e., participants made 4 estimates approximately every 35 trials for each of the red and green distributions under each condition). These accuracy scores were then entered into a General Linear Model to examine the effect of distribution pairs on how well individuals were able to accurately estimate the source distributions. All data were analyzed using SPSS Statistics version 20.

4.4. Experiment 1: Results

Analysis of participant's overall accuracy revealed a main effect of distribution, $F(2.2, 53.24) = 7.17$, $MSE = .02$, $p < .01$, colour, $F(1, 24) = 5.29$, $MSE = .01$, $p < .05$, a distribution by colour interaction, $F(3, 72) = 8.70$, $MSE = .004$, $p < .001$, and a distribution by colour by condition interaction, $F(3, 72) = 4.69$, $MSE = .004$, $p < .01$ (Figure 4.4). However, follow-up analyses demonstrated that much of these effects were driven by the fourth distribution pair that both groups of participants observed. Namely, the distribution that features a wide Gaussian (green) overlapping with a narrow Gaussian (red) sharing the same mean.




Figure 4.4. Experiment 1 Estimation Accuracy. Depicts distribution estimation accuracy for each distribution pair. Left panel shows performance for the distance condition pairs while the right panel shows performance in the overlap condition. Lines are coloured according to the distribution that was estimated by the participant. For this experiment the green distribution in the wide/narrow pairing was always wide while the red distribution was always narrow.

Examination of accuracy estimates for this distribution pair revealed an overall reduction in accuracy for this final distribution pair relative to the three pairs that did not share the same mean, $F(1,24) = 14.04$, $MSE = .02$, $p < .001$. Furthermore, comparing the two distributions directly revealed that participants tended to perform worse at estimating the narrow (red) distribution relative to the wide (green) distribution, $t(24) = 3.83$, $SE = .02$, $p < .01$ (Figure 4.4), mistakenly estimating it as being much wider than it actually was. This suggests that participants tended to over-generalize their representation of variability to match the widest distribution observed at any given time. This is also consistent with other work in our lab demonstrating that participants have particular trouble transitioning from representing wide distributions to narrower distributions (Filipowicz, Valadao, Anderson, & Danckert, under review).




Figure 4.5. Experiment 1 Estimation Accuracy by Condition. Depicts distribution estimation accuracy as a function of distribution pair condition. Distance condition had distribution pairs at 0.5SD, 1.0SD, and 1.5SD apart for distributions 1, 2, and 3, respectively. Overlap condition had distributions overlapping at 26%, 44%, and 66% for distributions 1, 2, and 3, respectively.

Indeed, when excluding this distribution pair, both main effects of colour, ($F < 1$) and distribution, $F(1.5, 35.5) = 2.26$, $MSE = .01$, $p = .13$, disappear, indicating that the uniquely difficult fourth distribution pair was driving these effects. In terms of between subject comparisons, while participants tended to estimate distributions more poorly when they were closer together (the overlapping condition, Figure 4.5), this difference did not reach statistical significance in the current experiment, $F(1, 24) = 2.28$, $MSE = .06$, $p = .14$.

4.5. Experiment 1: Discussion

In summary, results of Experiment 1 demonstrate that, without biasing the allocation of attention towards either of the two distributions, participants were able to estimate both with equal accuracy. While there was some indication that participants did more poorly overall when distributions overlapped with each other, this effect did not reach statistical significance.

However, participants did particularly poorly at estimating a narrow distribution when presented in tandem with an overlapping wider distribution, a result that is consistent with prior work in our lab demonstrating difficulty in representing narrow distributions after first being presented with a wider distribution (i.e., this work showed each distribution sequentially as opposed to the interwoven presentation used here).

4.6. Experiment 2: Introduction

Given that the prior experiment demonstrated individuals were able to estimate two distributions simultaneously reasonably well, the central question of the current study was to explore whether or not the allocation of attention would improve an individual's ability to model

their environment. The second experiment set out to test this hypothesis by examining performance in estimating two distributions while informally allocating attention preferentially towards one of them.

4.7. Experiment 2: Method

Participants

A total of 30 undergraduate students ($M = 19.0$ years, $SD = 1.3$ years, 12 male) from the University of Waterloo participated in the current experiment. All participants reported normal or corrected-to-normal vision and gave written consent to participation. As remuneration for their involvement, participants received bonus grade credit towards one of their Psychology courses that term. The experiment and all procedures were approved by the University of Waterloo, Office of Research Ethics.

Procedure

In contrast to the first experiment in which participants were asked to focus on both colours of ball drops, participants were now asked to attend only to the green ball drops, despite the fact that they'd later be asked to estimate distributions of both colours, a fact participants were informed of at the outset. This would allow for the exploration of the impact of selective attention on how well individuals represent regularities.

Despite our request to participants that they focus only on green ball drops, we did not implement a mechanism in the current experiment to strictly *enforce* allocation of attention. However, one critical change made to the current experiment was in relation to how the balls fell into the landing slots. Given that the prior experiment displayed ball drops one at a time, it was

possible that, under the current instructions, participants would be able to shift their focus between distributions as interleaved trials of red and green ball drops were presented. Therefore, the current experiment displayed both red and green ball drops at approximately the same time. To accomplish this, a modification to the Plinko task rendered two ball drops, one of each colour, falling through the pegs at roughly the same time. A temporal jitter was implemented such that pairs of balls never landed at the same time (Figure 4.6). Participants observed the same distribution pairs as in Experiment 1. Each trial consisted of the presentation of one green and red ball, with data being collected at trials 17, 35, 53, and 70.

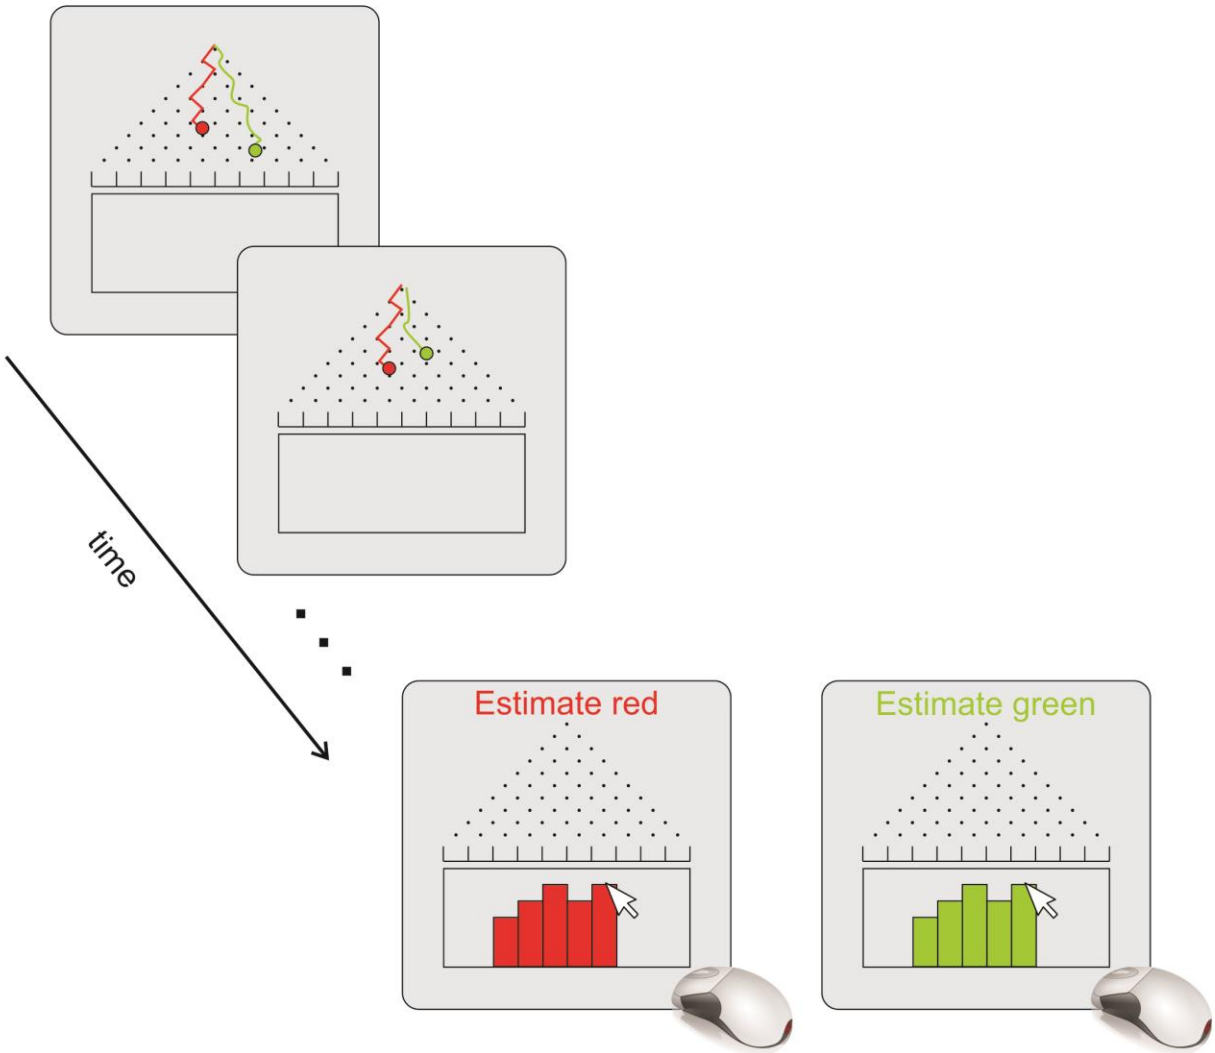


Figure 4.6. Simultaneous Plinko Task. Depicts a version of the Plinko task where participants would observe two simultaneous ball drops, one of each colour, rather than one colour at a time (as in Experiment 1). As with Experiment 1, participants also estimated both green and red distributions when prompted.

4.8. Experiment 2: Results

Analysis of participants' estimation accuracy reveals a main effect of attended colour, such that participants performed significantly better when estimating the attended as opposed to the unattended distribution, $F(1,28) = 8.91$, $MSE = .01$, $p < .01$. There was no main effect of distribution, $F < 1$, nor were there any significant interactions (Figure 4.7).

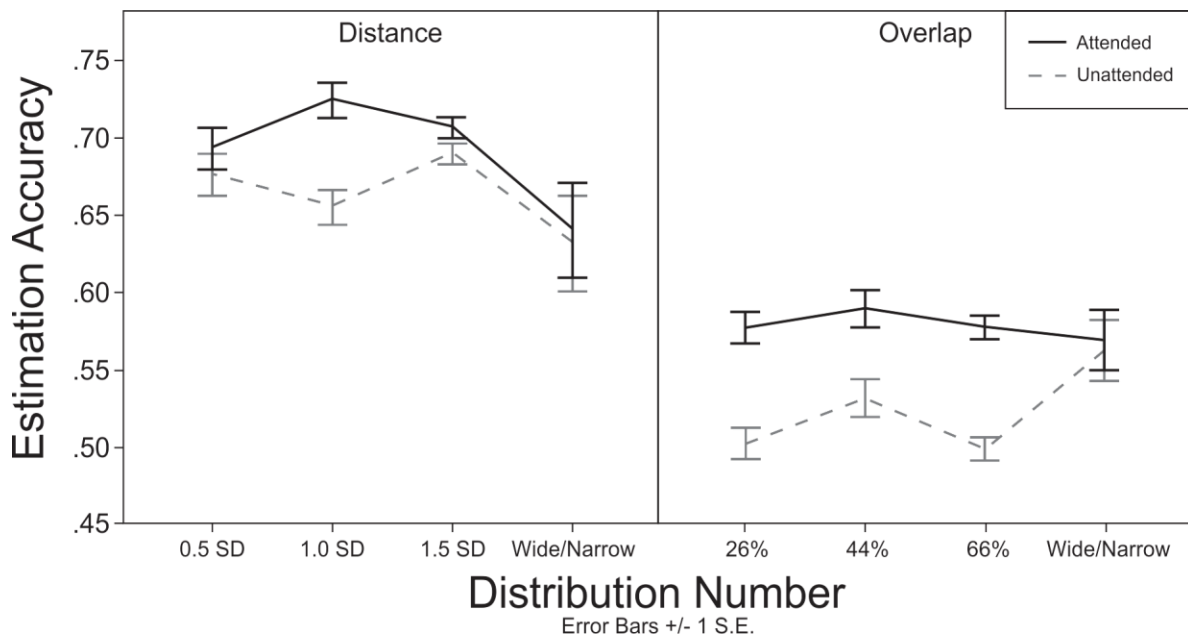


Figure 4.7. Experiment 2 Estimation Accuracy. Depicts distribution estimation accuracy for all distribution pairs across both distance and overlap conditions as a function of attention. While no manipulation was used to sustain participant's attention on the target distribution, they were asked to focus one distribution over another, despite knowing they would eventually have to estimate both.

In contrast to Experiment 1, there was a significant effect of condition in this experiment, $F(1,28) = 7.05$, $MSE = .03$, $p < .05$. Specifically, participants tended to do worse overall when

estimating distributions that overlapped with each other relative to conditions in which the distribution differed in the distance between means (Figure 4.8). This was even true for distributions that participants were asked to attend to. Although the same effect was not significant in Experiment 1, it was numerically in the same direction – better performance for distributions differing in distance vs. overlap conditions (compare Figures 4.5 and 4.8).

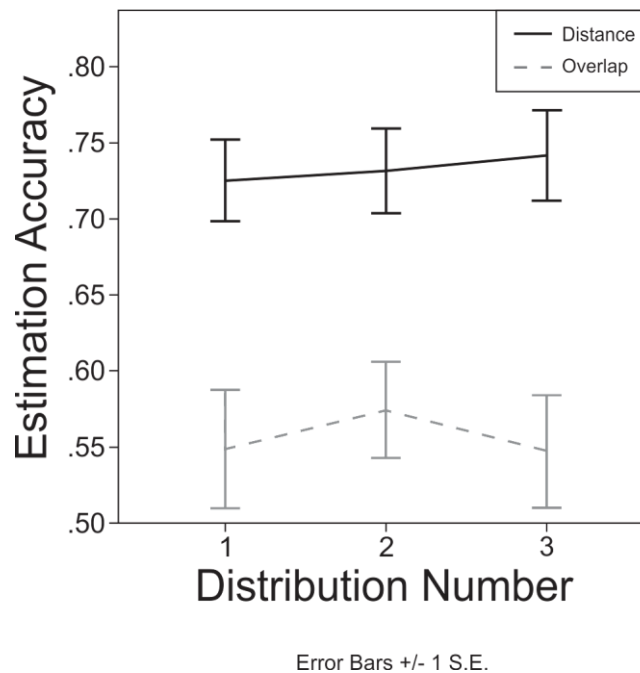


Figure 4.8. Experiment 2 Accuracy by Condition. Depicts distribution estimation accuracy for Experiment 2 as a function of distribution pair condition. Distance condition had distributions pairs at 0.5SD, 1.0SD, and 1.5SD apart for distributions 1, 2, and 3, respectively. Overlap condition had distributions overlapping at 26%, 44%, and 66% for distributions 1, 2, and 3, respectively.

Interestingly, participants still tended to have difficulty estimating a narrow distribution that overlapped with a wide one. Counterintuitively, that was true particularly when participants were specifically asked to pay more attention to the narrow distribution, $F(1,28) = 6.52$, $MSE = .01$, $p < .05$ (Figure 4.9). Despite the improvement in accuracy when attending to a target

distribution across all other distribution pairs, here they continued to estimate the attended distribution as being much wider than it actually was.




Figure 4.9. Experiment 2 Overlapping Distribution Accuracy. Depicts distribution estimation accuracy for Experiment 2 as a function of attended distribution demonstrating a significant reduction in narrow vs. wide accuracy despite attentional focus.

4.9. Experiment 2: Discussion

The results of Experiment 2, where we had participants allocate their attention towards one of the two distributions in each pair, demonstrated an improvement in estimation accuracy for the attended vs. the unattended distribution. In addition, while Experiment 1 showed performance trending downwards for distribution pairs that overlapped with each other, that effect was significant in Experiment 2. Finally, as with Experiment 1, participants had

considerable difficulty estimating the narrow distribution when it was presented in context with a wider distribution, despite the fact that they were asked to attend to it.

4.10. Experiment 3: Introduction

Despite the fact that results of Experiment 2 showed a benefit of attention on representational accuracy, there was no direct measure to determine whether participants did indeed allocate their attention to the target distribution. Furthermore, the effect of attentional allocation was subtle, with the largest effects being the impact of placing the distributions in an overlapping configuration (i.e., the overlap condition and the wide/narrow distribution pair). Therefore, the purpose of the current experiment was to replicate the findings of Experiment 2 while introducing an added element to the Plinko task that both enforced and measured the allocation of attention.

4.11. Experiment 3: Method

Participants

A total of 39 undergraduate students ($M = 20.5$, $SD = 3.2$, 6 male) from the University of Waterloo participated in the current experiment. All participants reported normal or corrected-to-normal vision and gave written consent to participation. As remuneration for their involvement, participants received bonus grade credit towards one of their Psychology courses that term. The experiment and all procedures were approved by the University of Waterloo, Office of Research Ethics.

Procedure

Participants completed a similar Plinko task to that of Experiment 2 with simultaneous ball drops of both red and green distributions. The same distribution pairs used in Experiments 1 and 2 were also used here. The colour of ball drops that participants were asked to attend to was counterbalanced between participants. While the prior experiment presented simultaneous ball drops as a means to prevent participants from shifting their attention to the unattended distribution on trials where the unattended balls fell, there wasn't a mechanism to encourage participants to sustain their attention on the target colour. Theoretically, participants could still attend to the other colour of ball drops despite instructions not to.

To ensure participants sustained attention on the colour instructed, participants were given a dual Sustained Attention to Response Task (SART; Robertson et al., 1997) whilst they observed the ball drops (Figure 4.10). In the original SART task, participants are presented with a rapid sequence of numbers and are told to press a button to every number except one. Commission errors – a failure to withhold a response to the specified number – represent failures of sustained attention (Robertson et al., 1997). The total proportion of go stimuli was 67%, a proportion associated with minimal errors and limited influence of speed accuracy trade-offs (Wilson, Finkbeiner, de Joux, Russell, & Helton, 2016).




Figure 4.10. Plinko SART Task. Participants would observe ball drops from the top of the screen into one of 40 landing slots. As the balls fell between the 16th and 21st row of pegs, both red and green balls would change to an array of possible shapes or remain as circles. Participants were asked to press the mouse button if the colour ball they were attending to changed to any shape *except* for a star.

In the current experiment, the SART task was modified such that participants had to pay attention to one of two coloured distributions of balls. As the balls fell, both the attended and unattended balls independently changed to an array of possible shapes mid-fall. Participants were instructed to click the left mouse button when they saw the attended colour circle change to a square, diamond, triangle, or semi-circle. Participants were asked to withhold their response if they saw the attended ball change to a star or on trials in which the ball did not change. The change point on each trial occurred in the middle half of its travel time with its change point determined randomly within that range to prevent a perseverative response style. The middle range was chosen to avoid early or late stages of a ball's trajectory. Early changes might mean the participant could ignore the ball's continuing trajectory for that trial and late changes had the possibility of inducing a kind of attentional blink potentially impairing participant's ability to accurately encode the shape's final landing position (Dux & Marois, 2009).

4.12. Experiment 3: Results

Evaluation of participant's performance on the SART task indicates that participants completed this component of the experiment according to instructions, with mean response accuracy of 89% (SD = .07%; Figure 4.11). As with the previous experiment, distribution estimation results showed a main effect of attended distribution on estimation accuracy, $F(1,37) = 5.0$, $MSE = .01$, $p < .05$ (Figure 4.12), demonstrating that participants performed better when estimating ball drops for the attended distribution (i.e., the same distribution for which they completed the SART task).

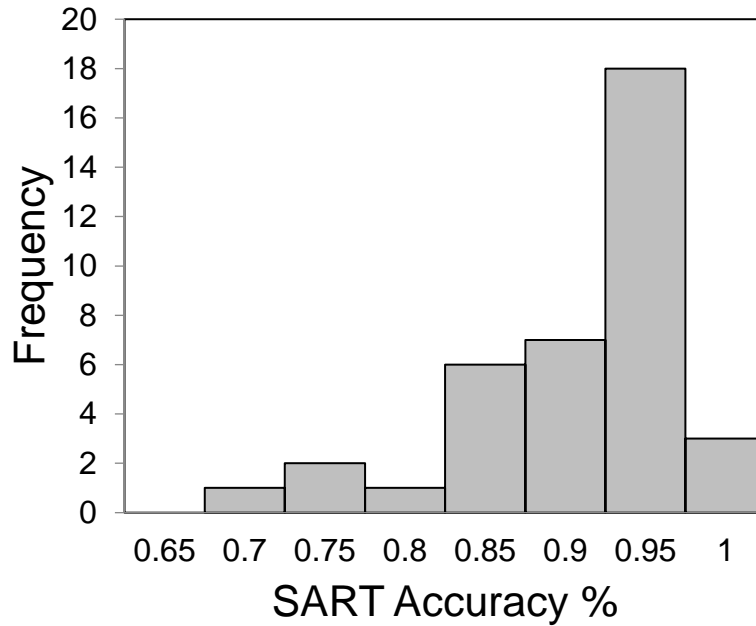


Figure 4.11. Experiment 3 SART Task Accuracy. Depicts histogram of SART accuracy for all participants in Experiment 3.

There was also a main effect of condition such that, as with prior experiments, participants were less accurate estimating distributions that were overlapping with each other, $F(1,37) = 7.82$, $MSE = .04$, $p < .01$ (Figure 4.13). However, there were no other significant main effects or interactions. While Experiment 2 demonstrated an overall reduction in performance when participants had to estimate wide and narrow distributions that completely overlapped, this effect was not present in the current experiment, $F < 1$. However, as with all prior experiments, participants did a worse job overall at estimating the narrow distribution when it overlapped with a wide one, $t(38) = -6.8$, $SE = .02$, $p < .001$ (Figure 4.14).




Figure 4.12. Experiment 3 Estimation Accuracy. Depicts distribution estimation accuracy as a function of attention. Attended distributions were ones where participants completed the SART task whilst observing ball drops.




Figure 4.13. Estimation Accuracy by Condition. Depicts distribution estimation accuracy as a function of condition. As with Experiments 1 and 2, the distance condition placed distribution pairs 0.5SD, 1.0SD, and 1.5SD apart while the overlap condition had distribution pairs overlapping by 24%, 44%, and 66%.

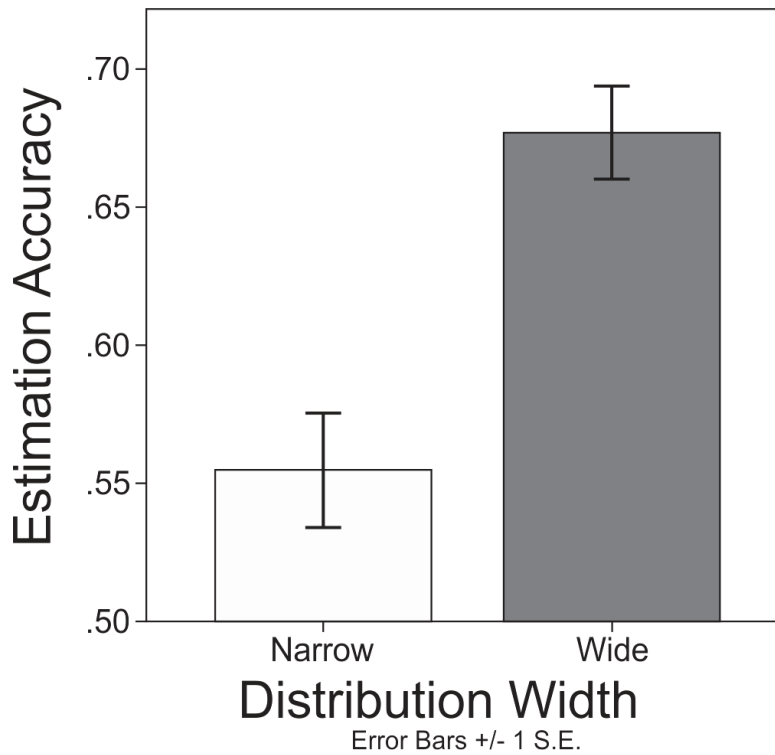


Figure 4.14. Estimation Accuracy for Wide vs. Narrow Distribution. Depicts distribution estimation accuracy for the fourth and final distribution pair common to both groups. Specifically, the distribution pair that had a narrow distribution overlapping a wider distribution with the same mean.

4.13. Experiment 3: Discussion

Consistent with the results from Experiment 2, the results of Experiment 3 demonstrated an advantage in distribution estimation accuracy when participants completed the SART task on those ball drops. Additionally, participants continued to have relatively more difficulty with distributions that were overlapping. Notably, participants again had the most difficulty with estimating the narrow distribution when it overlapped with the wider one, even when asked to attend to it through the use of the SART task.

4.14. Experiment 4: Introduction

While the prior two experiments demonstrated that the allocation of attention improved participants' ability to represent distributions, these effects did not extend to instances where the distribution pairs consisted of a wide and narrow distribution with overlapping means. Although the results consistently demonstrated lower overall performance as the distributions began to overlap significantly, the marked reduction in performance in the most difficult of distribution pairs warranted further exploration. Given that both the wide and narrow distributions always shared the same mean, it was important to explore whether the difficulty observed with this distribution pair was merely a result of that fact. That is, if participants' past performance with these distributions was a true reflection of difficulty with narrow distributions in the context of wider ones, then separating the means of the distributions should still produce the same pattern of performance. However, in the event that the results reflect the fact that both distributions shared the same mean, then separating the means should modulate the effect. Additionally, while prior experiments demonstrated a performance improvement for attended vs unattended distributions, performance was relatively poor for narrow distributions that overlapped wide

distributions even though their attention was allocated to ball drops from the narrower distribution. Consequently, Experiment 4 set out to investigate whether attention affected the ability to represent these narrow distributions relative to when they were unattended in the context of overlapping with a wider distribution.

4.15. Experiment 4: Method

Participants

A total of 20 undergraduate students ($M = 20.4$, $SD = 1.7$ years, 6 male) from the University of Waterloo participated in the current experiment. All participants reported normal or corrected-to-normal vision and gave written consent to participation. As remuneration for their involvement, participants received bonus grade credit towards one of their Psychology courses that term. The experiment and all procedures were approved by the University of Waterloo, Office of Research Ethics.

Procedure

The current study contained distribution pairs consisting only of a narrow distribution overlapping with a wide distribution. There were a total of 4 distribution pairs, with the distance between means of the narrow and wide distributions differing slightly in each pair (4, 6, 8, and 10 slots apart for distribution pairs 1-4, respectively), as opposed to previous experiments where the means of the wide and narrow distributions were identical. This allowed us to explore whether distance between distribution means would ameliorate the performance deficit observed in prior experiments where the narrow and wide distributions shared the same means.

As in Experiment 2, balls from each pair of distributions fell from the top of the pegs into one of 40 landing slots. Participants were asked either to attend to only green balls or red balls. The colour of the aforementioned distributions varied such that, regardless of which color participants were asked to attend to, there would always be two narrow and two wide distributions of that colour (Figure 4.15). This means that, taking the narrow distributions for example, they would always have two narrow distributions they *did* attend to and two narrow distributions where they *didn't* regardless of whether they were asked to attend to green or red ball drops. For example, for participants asked to attend to green balls, they would be attending to the wide distribution in pairs 1 and 4 and the narrow distribution in pairs 2 and 3. This allowed us to directly compare the effect of attention on representation accuracy of narrow and wide distributions within-subjects as opposed to prior experiments where participants were either attending to the wide *or* narrow distribution between-subjects as had been the case for all prior experiments in this chapter. Additionally, a third group of participants were given no instructions regarding the allocation of attention to serve as a control condition. The order of distribution pairs was counterbalanced across participants.




Figure 4.15. Experiment 4 Distribution Pairs. Top panel depicts distribution pairs used for Experiment 4. All distribution pairs consisted of a narrow distribution overlapping with a wider one, comparable to the fourth distribution pair in Experiments 1-3. Middle and bottom panels demonstrate how participants always attended to two narrow and two wide distributions when given instructions to attend to one ball colour. Solid lines denote attended distributions while dashed lines denote unattended distributions.

4.16. Experiment 4: Results

To examine the hypothesis that prior results for one specific pair of distributions was due to the fact that the means were identical (i.e., the narrow and wide Gaussians), an analysis was run looking at estimation accuracy as a function of distribution (with the means of both distributions being further apart in each successive pair) and width (wide vs. narrow). As with prior experiments that demonstrated poorer estimation performance for narrow distributions when it overlapped with a wider one, there was a significant main effect of distribution width such that participants were less accurate at estimating the narrower distributions, $F(1,18) = 35.84$, $MSE = .04$, $p < .001$ (Figure 4.16). There were no other main effects or interactions, suggesting that participants did not benefit in their estimations as a function of how far apart the wide and narrow distributions were placed relative to each other, all $F_s \leq 1$. These results suggest that the difficulty noted in previous experiments with respect to poor performance in narrow vs wide distributions was not an artefact of the distribution means being the same.

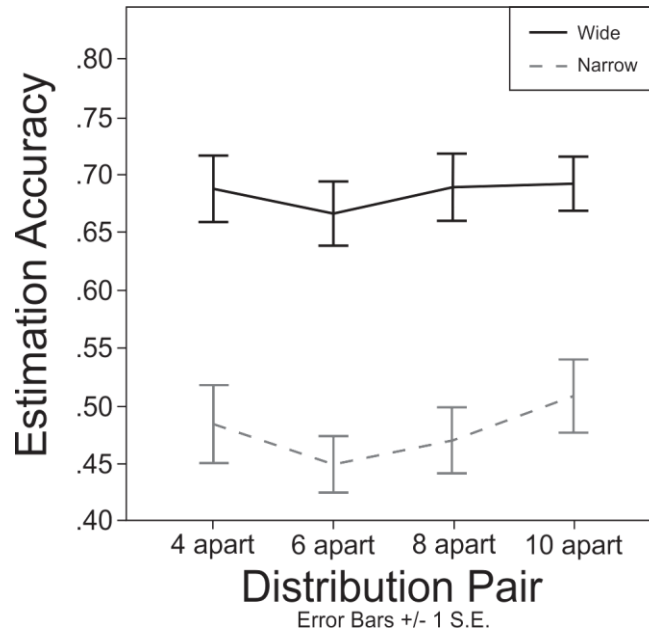


Figure 4.16. Experiment 4 Estimation Accuracy by Distribution Width. Depicts estimation accuracy for the wide distributions in solid lines and narrow distributions in dashed lines. Distance between distribution means within pairs increased from left to right.

Given the only factor affecting performance in the distribution estimation task appeared to be the width of the distribution, focus shifted to the effect of attention on how well individuals estimated both wide and narrow distributions. As previously mentioned, individuals attending to either color of ball drops would end up attending to a combination of both narrow and wide distributions over the course of the experiment (Figure 4.16). Analysis of within-subjects effects revealed a main effect of distribution width, $F(1,18) = 35.84$, $MSE = .02$, $p < .001$, attention, $F(1,18) = 4.6$, $MSE = .004$, $p < .05$, and an attention by distribution width interaction, $F(1,18) = 5.12$, $MSE = .004$, $p < .05$ (Figure 4.17).

Further exploration of the attention by distribution width interaction revealed that the primary improvement of estimation accuracy occurred when participants were attending to the narrow distribution. That is, under conditions where participants performed at their worst in prior experiments (i.e., the narrow distribution), performance was significantly improved when they were asked to focus on *those* ball drops relative to the condition in which they attended to the wider distributions (Figure 4.17).

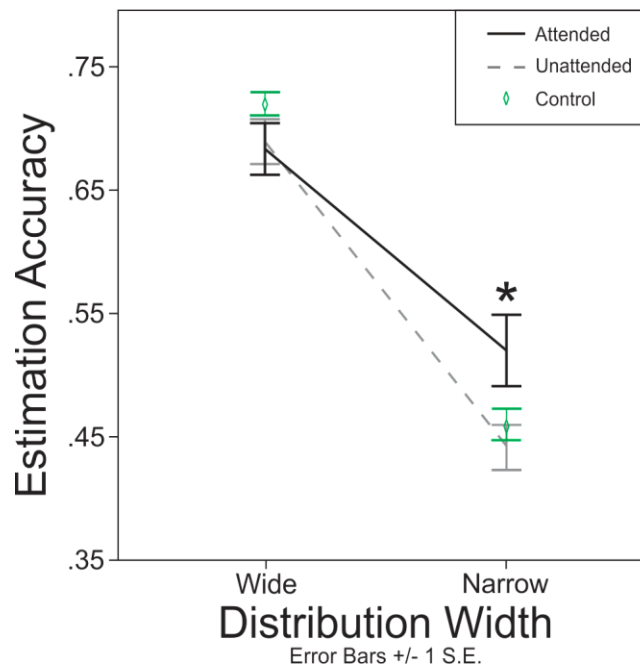


Figure 4.17. Experiment 4 Wide vs. Narrow Estimation Accuracy by Attention. Depicts distribution estimation accuracy for attended and unattended distributions as a function of distribution width. Green markers indicate no-instruction control performance. Asterisk denotes significant difference between attended and unattended narrow distribution estimation accuracy.

With respect to whether these results represent a facilitative effect of *focusing* attention above and beyond passive observation or whether attention represents a gating mechanism, performance was also examined as a function of whether participants received a specific distribution colour to focus on or given no attention instructions at all. Overall estimation accuracy was comparable between the groups with no main effect of attention instruction, $F < 1$. While performance in the attention condition appears to be higher when participants attended to the narrow condition, this difference failed to reach statistical significance, $F(1,18) = 2.26$, $MSE = .02$, $p = .15$.

4.17. Experiment 4: Discussion

These results suggest that participants were able to represent both distributions quite well regardless of whether or not they were instructed to attend to either at the exclusion of the other. However, when asked to allocate attention to one of the distributions in each pair, participants were more accurate when they were attending to the narrow distribution than when they weren't. Taken together, the results seem to suggest that attention gates the process of representing regularities, such that impairing the ability to freely allocate attention to the environment may impair performance.

4.18. Discussion

The results of the current study appear to demonstrate the robust and reliable role that attention plays in facilitating the building of mental representations. Through multiple experiments, our results demonstrated that individuals consistently did better at estimating the mean and variance of an unknown distribution of visual events when attentional resources were

allocated to those events. This was true regardless of whether participants were merely instructed to voluntarily attend to one of two distributions (Experiments 2 and 4) or whether there was a secondary task designed to enforce sustained attention on one distribution at the expense of the other (Experiment 3). These results are generally consistent with literature that shows attention is an important component of many visual learning processes including visual statistical learning (Baker, Olson, & Behrmann, 2004; Turk-Browne, Junge, & Scholl, 2005), probability learning (Kabata et al., 2014), sequence learning (Kimura, Widmann, & Schröger, 2010), and implicit learning (Jiménez & Mendéz, 1999).

However, when viewed in a broader context, these results indicate that individuals are quite capable of observing and representing multiple distributions simultaneously. This was particularly evident in the first experiment where no attentional constraints were placed on participants. That said, the results of the current study did demonstrate that individuals find it much more difficult to build a representation in the context of other competing information that are similar to the one being represented. Throughout all experiments, participants found it much harder to estimate distributions that overlapped with each other as opposed to when the boundaries of each respective distribution were separated from each other. This difficulty was exacerbated when the distributions shared the same mean (i.e., narrow over wide). This is consistent with literature demonstrating that similarity between targets and distractors require more attentional scrutiny (Rausei, Makovski, & Jiang, 2007), possibly explaining the overall lower performance under these circumstances. This is also consistent with work in our lab demonstrating that individuals have considerable difficulty shifting from representing a wide distribution towards a narrow distribution even when those distributions are presented sequentially (Filipowicz et al., under consideration).

In terms of the mechanism by which attention affects the model building process, the results of Experiment 1 strongly suggested that individuals were able to effectively represent two distributions at once without any attentional manipulations. While later manipulations did result in improved performance for the attended vs. unattended distributions, it appeared as though much of that difference resulted from *a reduction* in performance for the unattended distributions relative to Experiment 1 levels rather than a heightening of accuracy for the attended distributions (Figure 4.18). This is also consistent with results from Experiment 4 where we manipulated the target of participants' focus as an independent variable while asking them to perform what was the most difficult estimation task from all three prior experiments. While participants did benefit from focusing on the narrow distributions relative to when their attention was on the wider ones, participants performed quite well when they were simply free to attend to either distribution relative to when their focus was constrained by task instructions. This is consistent with the results from Chapter 2 showing that participants were best able to represent regularity in their environment when working memory resources were freed up rather than exploited to hold recent events in mind, even when those events were relevant to the mental representation they needed to build in the first place.

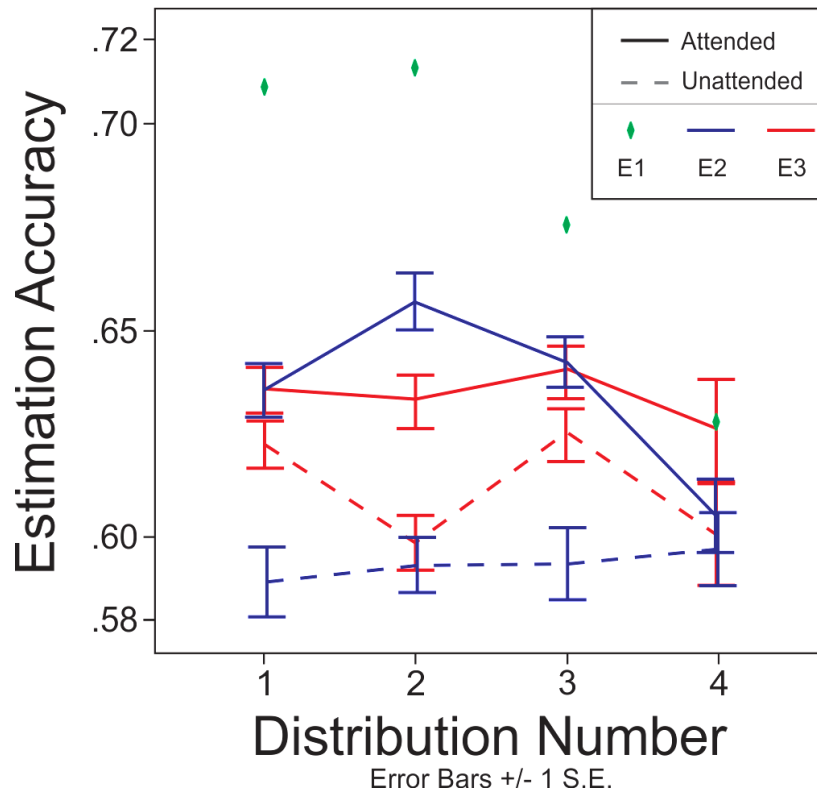


Figure 4.18. Experiments 1-3 Combined Attended vs. Unattended vs. Control. Depicts relative performance across Experiments 1-3. Dashed lines indicate unattended distributions in each respective pair while solid lines indicate performance when attention was focused on the target distribution. Both Experiments 2 and 3 had significant main effects of attention.

In summary, it appears that attentional resources act as a gating mechanism and are most helpful for the building of mental representations when they are freed up to passively observe visual regularity in the environment. Through multiple experiments, participants were consistently better at determining the underlying structure of a series of visually presented ball drops when they were focused on those events relative to a second distribution to which they were not. However, the overall pattern of performance under these circumstances were comparable to conditions where individuals were free to self-regulate the allocation of their attentional resources. To the extent that attentional resources act as a mechanism for building and

updating mental representation, it appears as though it serves as a gating mechanism in the information processing system as individuals seek to infer the structure of their environment through limited and efficient observation (Gershman & Niv, 2010; Gottlieb, 2012).

Chapter 5: General Discussion

The ability to build and update mental models of the environment represents an important component to effective and efficient decision making as it allows individuals to predict the consequences of action choices without over-sampling the environment (Griffiths & Tenenbaum, 2006; Tenenbaum et al., 2011, Johnson-Laird, 2013). The current thesis explored mental model building and updating from the point of view of the neural correlates, and the ancillary cognitive abilities that facilitate the process.

Chapter 2 explored the brain regions putatively involved in facilitating mental model building and updating. Building on existing work in the literature with brain-damaged individuals highlighting a likely role of the right-hemisphere in this process (Danckert et al., 2012), Chapter 2 investigated the neural regions that played an important role in the ability of healthy individuals to build and update mental models. Neurologically healthy participants played a serial zero-sum game with a computer that utilized a variety of potentially exploitable strategies. Results indicated that individuals were reasonably skilled at modeling and exploiting the bias of the computer opponent as demonstrated by steadily rising rates of optimal play choice over the course of each relevant epoch. Imaging data captured during these instances where individuals were able to model and exploit their computer opponent demonstrated significant activation in a variety of areas. Most notably, a network of parietal cortex, cingulate cortex, striate cortex, and prefrontal cortex.

The network of regions activated likely contributes in distinct ways to model building and updating, with further research needed to elucidate the separable contributions of each region. Given research findings from other domains it is reasonable to speculate that the parietal cortex, classically characterized as an area important for the allocation of visual attention

(Mesulam, 1981; Hopfinger et al., 2000; Corbetta & Shulman, 2002; Ferber & Danckert, 2006), coupled with prefrontal cortex, typically implicated in executive control (Hare et al., 2009; Geng, 2013), are involved in the exploration and accumulation of evidence in pursuit of a mental model (d'Acremont et al., 2013a; 2013b). Once built, such a model needs to be evaluated in the context of expected and observed outcomes. Relevant activations in both the anterior cingulate cortex, typically implicated in error detection and belief-based decision making tasks (Zhu et al., 2012), and striate cortex, typically characterized as representing or encoding relative risk and reward signals (Levin et al., 2012) may co-ordinate this process. These results are consistent with other work in our lab demonstrating their importance in the model building and updating process (Danckert et al., 2012, Stöttinger, Filipowicz, Valadao, Culham, Goodale, Anderson, & Danckert, 2015) in addition to patient studies that implicate a subset of these areas in patients who demonstrate difficulty with this ability (Danckert et al., 2012; Stöttinger et al., 2014). It is worth noting some limitations to this work that suggest future directions of research. First, the nature of the block-design coupled with a relatively smaller sample size limits the ability to detect areas of activation that come online *the moment* a model is built or updated and instead infers that these processes are underway throughout the block of trials. Nevertheless, our results are consistent with similar work employing more event-related designs (Stöttinger et al., 2015). A potentially fruitful avenue for future research would involve replication of these findings while exploring the temporal dynamics of how this network operates. It would be important to corroborate the purported roles that the above areas play in the updating process with research that isolates their role in context. In addition to utilizing event-related designs in future research, this investigation could involve the use of alternative imaging techniques with greater temporal resolution than fMRI (e.g., EEG/ERP).

Chapter 3 consisted of a series of behavioural experiments examining the role of Working Memory in model building and updating. Participants completed a serial prediction task asking them to estimate the mean of a distribution whilst completing a dual task that either constrained WM resources through the use of a 2-back task or a 0-back dual-task control that wouldn't constrain working memory resources. While it was initially hypothesized that expending WM resources to focus on elements of the environment that contained information relevant to the model to be built (e.g., recent examples of the distribution), those were the conditions in which performance was at its worst. Specifically, participants were slowest to update their representations of the central tendency of a new distribution when they were asked to hold in WM information about the stimuli that were relevant or congruent with the model they needed up update (i.e., Location in E1, Shape in E2/3). These results are consistent with prior research demonstrating constrained WM resources impair the ability to learn and benefit from regularity in the environment and detect salient changes (Rouder et al., 2011; Annac et al., 2013; Travis et al., 2013). Taken together, Chapter 2 suggested that free WM resources represent an important gating mechanism to effective and efficient updating of mental representations such that constrained or reduced capacity may impair the process. However, these results are somewhat limited by the relatively constrained response choices made by participants, coupled with the fact that the presence of said models were inferred by the tendency to choose one of these limited options. This could be addressed in future research by increasing the range and complexity of possible response choices (as was done in the "Plinko task"). Additionally, while the results point to the importance of WM resources, further research could investigate the relationship between WM *capacity* or WM rehabilitation and the ability to build and update mental models.

Chapter 4 examined the role of attention in the building of mental representations. Using a novel task (Filipowicz, Valadao, Anderson, & Danckert, 2014) designed to allow participants to input data that more closely reflected their internal representation, participants observed ball drops reflecting two distinct distributions of events. Across four experiments, a variety of manipulations were applied to alter the allocation of attentional resources to one distribution over the other. Results generally showed that, without any attentional constraints, participants were able to represent and report both distributions with comparable accuracy. However, when asked to attend to one of them, participants were significantly better at representing the attended vs. the unattended distribution. However, when taken together with data where participants were not told where or how to attend to the distributions (i.e., Experiment 4), performance in the attended conditions were not significantly different from controls. These results are consistent with the interpretation that attention, as with WM, act as a gating mechanism for the ability to build and update mental representations. Despite this conclusion, only a limited subset of the dataset directly contrasted attended vs. unattended in relation to a control condition without attentional constraints (i.e., Experiment 4). While Experiment 1 did demonstrate better performance for the distributions pairs without any attentional instruction, it's important to note that the interleaved ball drops arguably made the task of representing each distribution more difficult than when both distributions displayed ball drops simultaneously. Additional research would seek to replicate the findings of Experiment 4 using distributions without equal means (as in Experiments 1-3) with simultaneous presentation of ball drops and comparable controls (i.e., no attentional instructions).

In summary, the work presented in this thesis contributes to a more comprehensive understanding of how individuals build and update mental representations of their environment.

The present research helps establish a working knowledge of the neural network that co-ordinates this process including parietal cortex, anterior cingulate and anterior insula – regions implicated in both patient work and related fMRI studies (Karnath, 2004; Danckert et al., 2012, Stöttinger et al., 2014). These areas appear to play a co-ordinated role in the integration of evidence and the evaluation of predicted vs. observed outcomes. Additional data in this thesis outline the cognitive abilities that play a facilitative role as a gateway for incoming information, namely attention and working memory.

References

- Abla, D., & Okanoya, K. (2009). Visual statistical learning of shape sequences: An ERP study. *Neuroscience Research*, *64*(2), 185-190.
- Albrecht, A. R., Scholl, B. J., & Chun, M. M. (2012). Perceptual averaging by eye and ear: Computing summary statistics from multimodal stimuli. *Attention, Perception, & Psychophysics*, *74*(5), 810-815.
- Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, *9*(4), 321-324.
- Baddeley, A. (1992). Working memory. *Science*, *255*(5044), 556.
- Baddeley, A. (2003). Working memory: looking back and looking forward. *Nature Reviews Neuroscience*, *4*(10), 829-839.
- Baker, C. I., Behrmann, M., & Olson, C. R. (2002). Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. *Nature Neuroscience*, *5*(11), 1210-1216.
- Balleine, B. W., Delgado, M. R., & Hikosaka, O. (2007). The role of the dorsal striatum in reward and decision-making. *The Journal of Neuroscience*, *27*(31), 8161-8165.
- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214-1221.
- Berniker, M., Voss, M., & Kording, K. (2010). Learning priors for Bayesian computations in the nervous system. *PloS one*, *5*(9), e12686.
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, *138*(4), 487.
- Brown, S. D., & Steyvers, M. (2009). Detecting and predicting changes. *Cognitive*

Psychology, 58(1), 49-67.

- Buchsbaum, B. R., Greer, S., Chang, W. L., & Berman, K. F. (2005). Meta-analysis of neuroimaging studies of the Wisconsin Card-Sorting task and component processes. *Human Brain Mapping*, 25(1), 35-45.
- Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition*, 121(1), 127-132.
- Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, 280(5364), 747-749.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36(1), 28-71.
- Cohen, J. D., McClure, S. M., & Angela, J. Y. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1481), 933-942.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3), 201-215.
- Culham, J. C., Danckert, S. L., De Souza, J. F., Gati, J. S., Menon, R. S., & Goodale, M. A. (2003). Visually guided grasping produces fMRI activation in dorsal but not ventral stream brain areas. *Experimental Brain Research*, 153(2), 180-189.
- Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, 7(9), 415-423.
- d'Acremont, M., Fornari, E., & Bossaerts, P. (2013). Activity in inferior parietal and medial prefrontal cortex signals the accumulation of evidence in a probability learning task. *PLoS*

Comput Biol, 9(1), e1002895.

d'Acremont, M., Lu, Z. L., Li, X., Van der Linden, M., & Bechara, A. (2009). Neural correlates of risk prediction error during reinforcement learning in humans. *Neuroimage*, 47(4), 1929-1939.

d'Acremont, M., Schultz, W., & Bossaerts, P. (2013a). The human brain encodes event frequencies while forming subjective beliefs. *The Journal of Neuroscience*, 33(26), 10887-10897.

d'Acremont, M., Fornari, E., & Bossaerts, P. (2013). Activity in inferior parietal and medial prefrontal cortex signals the accumulation of evidence in a probability learning task. *PLoS Comput Biol*, 9(1), e1002895.

Danckert, J., Stöttinger, E., Quehl, N., & Anderson, B. (2012). Right hemisphere brain damage impairs strategy updating. *Cerebral Cortex*, 22(12), 2745-2760.

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876-879.

De Gardelle, V., & Summerfield, C. (2011). Robust averaging during perceptual judgment. *Proceedings of the National Academy of Sciences*, 108(32), 13341-13346.

Decety, J., & Lamm, C. (2007). The role of the right temporoparietal junction in social interaction: how low-level computational processes contribute to meta-cognition. *The Neuroscientist*, 13(6):580-593.

Dinstein, I., Hasson, U., Rubin, N., & Heeger, D. J. (2007). Brain areas selective for both observed and executed movements. *Journal of Neurophysiology*, 98(3), 1415-1427.

Domenech, P., & Koehlin, E. (2015). Executive control and decision-making in the prefrontal cortex. *Current Opinion in Behavioral Sciences*, 1, 101-106.

- Donoso, M., Collins, A. G., & Koechlin, E. (2014). Foundations of human reasoning in the prefrontal cortex. *Science*, *344*(6191), 1481-1486.
- Druker, M., & Anderson, B. (2010). Spatial probability aids visual stimulus discrimination. *Frontiers in Human Neuroscience*, *4*, 63.
- Dux, P. E., & Marois, R. (2009). The attentional blink: A review of data and theory. *Attention, Perception, & Psychophysics*, *71*(8), 1683-1700.
- Ferber, S., & Danckert, J. (2006). Lost in space—the fate of memory representations for non-neglected stimuli. *Neuropsychologia*, *44*(2), 320-325.
- Filipowicz, A., Anderson, B., & Danckert, J. (2014). Learning what from where: Effects of spatial regularity on nonspatial sequence learning and updating. *The Quarterly Journal of Experimental Psychology*, *67*(7), 1447-1456.
- Filipowicz, A., Valadao, D., Anderson, B., & Danckert, J. (2014). Measuring the influence of prior beliefs on probabilistic estimations. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 2198–2203). Austin, TX: Cognitive Science Society.
- Filipowicz, A., Valadao, D., & Anderson, B., & Danckert, J. Rejecting outliers: Surprising changes do not always improve belief updating. *Revise and resubmit to Decision*. October, 2016.
- Fischer, J., & Whitney, D. (2012). Attention gates visual coding in the human pulvinar. *Nature Communications*, *3*, 1051.
- Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, *12*(6), 499-504.
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from

- visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3), 458.
- Friston, K. (2012). Ten ironic rules for non-statistical reviewers. *Neuroimage*, 61(4), 1300-1310.
- Furl, N., & Averbeck, B. B. (2011). Parietal cortex and insula relate to evidence seeking relevant to reward-related decisions. *The Journal of Neuroscience*, 31(48), 17572-17582.
- Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, 109(3), 416-422.
- Geng, J. J., Soosman, S., Sun, Y., DiQuattro, N. E., Stankevitch, B., & Minzenberg, M. J. (2013). A match made by modafinil: Probability matching in choice decisions and spatial attention. *Journal of Cognitive Neuroscience*, 25(5), 657-669.
- Geng, J. J., & Vossel, S. (2013). Re-evaluating the role of TPJ in attentional control: contextual updating? *Neuroscience & Biobehavioral Reviews*, 37(10), 2608-2620.
- Geng, J. J., & Behrmann, M. (2002). Probability cuing of target location facilitates visual search implicitly in normal participants and patients with hemispatial neglect. *Psychological Science*, 13(6), 520-525.
- Gershman, S. J., & Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology*, 20(2), 251-256.
- Gläscher, J., Adolphs, R., Damasio, H., Bechara, A., Rudrauf, D., Calamia, M., & Tranel, D. (2012). Lesion mapping of cognitive control and value-based decision making in the prefrontal cortex. *Proceedings of the National Academy of Sciences*, 109(36), 14681-14686.
- Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron*, 76(2), 281-295.
- Green, C. S., Benson, C., Kersten, D., & Schrater, P. (2010). Alterations in choice behavior by manipulations of world model. *Proceedings of the National Academy of Sciences*, 107(37),

16401-16406.

- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science, 17*(9), 767-773.
- Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science, 324*(5927), 646-648.
- Heilman, K.M., & Valenstein, E. (2011). *Clinical Neuropsychology*. New York, NY: Oxford University Press.
- Hohwy, J. (2012). Attention and Conscious Perception in the Hypothesis Testing Brain. *Frontiers in Psychology, 3*, 96.
- Hopfinger, J. B., Buonocore, M. H., & Mangun, G. R. (2000). The neural mechanisms of top-down attentional control. *Nature Neuroscience, 3*(3), 284-291.
- Ingre, M. (2013). Why small low-powered studies are worse than large high-powered studies and how to protect against “trivial” findings in research: Comment on Friston (2012). *Neuroimage, 81*, 496-498.
- Jiang, Y. V., Swallow, K. M., & Rosenbaum, G. M. (2013). Guidance of spatial attention by incidental learning and endogenous cuing. *Journal of Experimental Psychology: Human Perception and Performance, 39*(1), 285.
- Jiang, Y. V., Swallow, K. M., Won, B. Y., Cistera, J. D., & Rosenbaum, G. M. (2015). Task specificity of attention training: the case of probability cuing. *Attention, Perception, & Psychophysics, 77*(1), 50-66.
- Jiménez, L., & Méndez, C. (1999). Which attention is needed for implicit sequence learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25*(1), 236.

- Johnson-Laird, P. N. (2013). Mental models and cognitive change. *Journal of Cognitive Psychology*, 25(2), 131-138.
- Johnson-Laird, P. N. (2010). Mental models and human reasoning. *Proceedings of the National Academy of Sciences*, 107(43), 18243-18250.
- Jueptner, M., Stephan, K. M., Frith, C. D., Brooks, D. J., Frackowiak, R. S., & Passingham, R. E. (1997). Anatomy of motor learning. I. Frontal cortex and attention to action. *Journal of Neurophysiology*, 77(3), 1313-1324.
- Kabata, T., Yokoyama, T., Noguchi, Y., & Kita, S. (2014). Location probability learning requires focal attention. *Perception*, 43(4), 344-350.
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, 10(12), 1625-1633.
- Kadota, H., Nakajima, Y., Miyazaki, M., Sekiguchi, H., Kohno, Y., & Kansaku, K. (2009). Anterior prefrontal cortex activities during the inhibition of stereotyped responses in a neuropsychological rock–paper–scissors task. *Neuroscience Letters*, 453(1), 1-5.
- Karnath, H. O., Berger, M. F., Küker, W., & Rorden, C. (2004). The anatomy of spatial neglect based on voxelwise statistical analysis: a study of 140 patients. *Cerebral Cortex*, 14(10), 1164-1172.
- Kimura, M., Widmann, A., & Schröger, E. (2010). Top-down attention affects sequential regularity representation in the human visual system. *International Journal of Psychophysiology*, 77(2), 126-134.
- Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition*, 38(6), 667-676.
- Kwon, O. S., & Knill, D. C. (2013). The brain uses adaptive internal models of scene statistics

- for sensorimotor estimation and planning. *Proceedings of the National Academy of Sciences*, 110(11), 1064-1073.
- Levin IP, Xue G, Weller JA, Reimann M, Lauriola M, Bechara A. 2012. A neuropsychological approach to understanding risk-taking for potential gains and losses. *Frontiers in Neuroscience*. 6, 1-11
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279-281.
- MacDonald, A. W., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science*, 288(5472), 1835-1838.
- Manginelli, A. A., Langer, N., Klose, D., & Pollmann, S. (2013). Contextual cueing under working memory load: Selective interference of visuospatial load with expression of learning. *Attention, Perception, & Psychophysics*, 75(6), 1103-1117.
- Mangun, G. R. (1995). Neural mechanisms of visual selective attention. *Psychophysiology*, 32(1), 4-18.
- Mayr, U. (1996). Spatial attention and implicit sequence learning: Evidence for independent learning of spatial and nonspatial sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(2), 350.
- McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, 84(4), 870-881.
- Mesulam, M. (1981). A cortical network for directed attention and unilateral neglect. *Annals of Neurology*, 10(4), 309-325.
- Miller, M. B., Valsangkar-Smyth, M., Newman, S., Dumont, H., & Wolford, G. (2005). Brain

- activations associated with probability matching. *Neuropsychologia*, 43(11), 1598-1608.
- Müller, H. J., & Rabbitt, P. M. (1989). Reflexive and voluntary orienting of visual attention: time course of activation and resistance to interruption. *Journal of Experimental Psychology: Human Perception and Performance*, 15(2), 315.
- Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *The Journal of Neuroscience*, 30(37), 12366-12378.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19(1), 1-32.
- Paulus, M. P., Feinstein, J. S., Tapert, S. F., & Liu, T. T. (2004). Trend detection via temporal difference model predicts inferior prefrontal cortex activation during acquisition of advantageous action selection. *Neuroimage*, 21(2), 733-743.
- Paulus, M. P., Feinstein, J. S., Leland, D., & Simmons, A. N. (2005). Superior temporal gyrus and insula provide response and outcome-dependent information during assessment and action selection in a decision-making situation. *Neuroimage*, 25(2), 607-615.
- Paulus, M. P., Hozack, N., Zauscher, B., McDowell, J. E., Frank, L., Brown, G. G., & Braff, D. L. (2001). Prefrontal, parietal, and temporal cortex networks underlie decision-making in the presence of uncertainty. *Neuroimage*, 13(1), 91-100.
- Pearson, J. M., Heilbronner, S. R., Barack, D. L., Hayden, B. Y., & Platt, M. L. (2011). Posterior cingulate cortex: adapting behavior to a changing world. *Trends in Cognitive Sciences*, 15(4), 143-151.
- Peirce, J. W. (2007). PsychoPy—psychophysics software in Python. *Journal of Neuroscience Methods*, 162(1), 8-13.

- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, *109*(2), 160.
- Rangel, A., & Hare, T. (2010). Neural computations associated with goal-directed choice. *Current Opinion in Neurobiology*, *20*(2), 262-270.
- Rausei, V., Makovski, T., & Jiang, Y. V. (2007). Attention dependency in implicit learning of repeated search context. *The Quarterly Journal of Experimental Psychology*, *60*(10), 1321-1328.
- Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T., & Yiend, J. (1997). Oops!': performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia*, *35*(6), 747-758.
- Rouder, J. N., Morey, R. D., Morey, C. C., & Cowan, N. (2011). How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin & Review*, *18*(2), 324-330.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, *70*(1), 27-52.
- Saffran, J. R., & Wilson, D. P. (2003). From Syllables to Syntax: Multilevel Statistical Learning by 12-Month-Old Infants. *Infancy*, *4*(2), 273-284.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*, 5294.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, *70*(1), 27-52.
- Saffran, J. R., Newport, E. L., Aslin, R. N., Tunick, R. A., & Barrueco, S. (1997). Incidental language learning: Listening (and learning) out of the corner of your ear. *Psychological*

- Science*, 8(2), 101-105.
- Shaqiri, A., & Anderson, B. (2012). Spatial probability cuing and right hemisphere damage. *Brain and Cognition*, 80(3), 352-360.
- Stadler, M. A. (1995). Role of attention in implicit learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(3), 674.
- Stöttinger, E., Filipowicz, A., Danckert, J., & Anderson, B. (2014). The effects of prior learned strategies on updating an opponent's strategy in the rock, paper, scissors game. *Cognitive Science*, 38(7), 1482-1492.
- Stöttinger, E., Filipowicz, A., Marandi, E., Quehl, N., Danckert, J., & Anderson, B. (2014). Statistical and perceptual updating: correlated impairments in right brain injury. *Experimental Brain Research*, 232(6), 1971-1987.
- Stöttinger, E., Filipowicz, A., Valadao, D., Culham, J. C., Goodale, M. A., Anderson, B., & Danckert, J. (2015). A cortical network that marks the moment when conscious representations are updated. *Neuropsychologia*, 79(A), 113-122.
- Straube, B., Wolk, D., & Chatterjee, A. (2011). The role of the right parietal lobe in the perception of causality: a tDCS study. *Experimental Brain Research*, 215(3-4), 315-325.
- Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain. 3-Dimensional proportional system: an approach to cerebral imaging*. Thieme, New York.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. *Science*, 331(6022), 1279–1285.
- Toro, J. M., Sinnett, S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition*, 97(2), 25-34.
- Travis, S. L., Mattingley, J. B., & Dux, P. E. (2013). On the Role of Working Memory in Spatial

- Contextual Cueing. *Journal of experimental psychology. Learning, Memory, and Cognition*, 39(1), 208-219.
- Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(2), 399–407.
- Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The Automaticity of Visual Statistical Learning. *Journal of Experimental Psychology: General*, 134(4), 552–564.
- Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *The Journal of Neuroscience*, 30(33), 11177-11187.
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*. 185(4157), 1124–1131.
- Uddin, L. (2014). Salience processing and insular cortical function and dysfunction. *Nature Reviews Neuroscience*, 16(1), 1-7.
- Unsworth, N., & Engle, R. W. (2005). Individual differences in working memory capacity and learning: Evidence from the serial reaction time task. *Memory & Cognition*, 33(2), 213-220.
- Valadao, D. F., Anderson, B., & Danckert, J. (2015). Examining the influence of working memory on updating mental models. *The Quarterly Journal of Experimental Psychology*, 68(7), 1442-1456.
- Vickery, T., Chun, M., & Lee, D. (2011). Ubiquity and Specificity of Reinforcement Signals throughout the Human Brain. *Neuron*, 72(1), 166-177.
- Walton, M., Crosson, P., Behrens, T., Kennerley, S., Rushworth, M.. (2007). Adaptive decision making and value in the anterior cingulate cortex. *Neuroimage*, 36, 142-154.

- Wilder, M., Jones, M., & Mozer, M. (2009). Sequential effects reflect parallel learning of multiple environmental regularities. *Advances in Neural Information Processing Systems*, 22, 2053–2061.
- Wilson, K., Finkbeiner, K., de Joux, N., Russell, P., & Helton, W. (2016). Go-stimuli proportion influences response strategy in a sustained attention to response task. *Experimental brain research*, 234(10), 2989-2998.
- Wolbers, T., Hegarty, M., Büchel, C., & Loomis, J. M. (2008). Spatial updating: how the brain keeps track of changing object locations during observer motion. *Nature Neuroscience*, 11(10), 1223-1230.
- Wolford, G., Miller, M.B., & Gazzaniga, M. (2000). The left hemisphere's role in hypothesis formation. *The Journal of Neuroscience*, 20(6), 1-4.
- Wojciulik, E., & Kanwisher, N. (1999). The generality of parietal involvement in visual attention. *Neuron*, 23(4), 747-764.
- Wolpert, D.M. (2007). Probabilistic models in human sensorimotor control. *Human Movement Science*, 26(4), 511–524.
- Yantis, S., & Jonides, J. (1990). Abrupt visual onsets and selective attention: voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception and Performance*, 16(1), 121-134.
- Yu, A., & Cohen, J. (2009). Sequential effects: superstition or rational behavior? *Advances in Neural Information Processing Systems*, 21, 1873-1880.
- Yu, A., & Dayan, P. (2005). Uncertainty, Neuromodulation, and Attention. *Neuron*, 46(4), 681–692.
- Yu, A. (2007). Adaptive Behaviour: Humans Act as Bayesian Learners. *Current Biology*, 17(22),

977-980.

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention is spontaneously biased toward regularities. *Psychological Science, 24*(5), 667-677.

Zhu, L., Mathewson, K., & Hsu, M. (2012). Dissociable neural representations of reinforcement and belief prediction errors underlie strategic learning. *Proceedings of the National Academy of Sciences of the United States of America, 109*(5), 1419-1424.