

Modelling Movement as an Ongoing Decision

by

Nathan James Wispinski

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

Department of Psychology

University of Alberta

© Nathan James Wispinski, 2017

Abstract

How do we make decisions, like which drink to reach and grab out of a fridge, or which path to take on a hike? Decades of research on decision making has led to powerful models of how options compete within our brains to be selected. However, these models, and the field of decision making in general, is split in two halves: the competition between options before a movement begins, and the competition between options after a movement has begun.

Here, I propose a computational model of decision making that aims to bridge these two halves. Specifically, we model decision making as one process that determines when to initiate a movement, and another that acts to move toward an option proportional to its ongoing desirability. In Chapter 2, this model is compared to data collected while people were asked to reach and touch which of two snack foods they preferred. Despite its relative simplicity, we find the model can account for a diverse range of human behaviour during decision making, and is highly generalizable to a variety of experimental paradigms.

In Chapter 3, I present neural data collected during a different decision making task - one where people were asked to decide which of two circles was brighter. These data do not align with many models focused on only one half of decision making, and instead provide support for the neural processes predicted by our model in Chapter 2.

Together, these results support a new formal way of thinking about decision making - as a single, continuous competition between options, only seemingly split in half by the question about when to move.

Preface

This thesis is an original work by Nathan James Wispinski. The research project, of which this thesis is a part, received research ethics approval from the University of Alberta Research Ethics Board, Project Name “ACE 1”, Pro00038718, 04/18/2013.

Some of the research conducted for this thesis forms part of a research collaboration, with the experiment in Chapter 2 designed by Dr. James T. Enns, Grace Truong, and Dr. Craig S. Chapman during Dr. Chapman’s postdoctoral fellowship at the University of British Columbia Department of Psychology. Behavioural data for Chapter 2 were collected by Tina S.T. Huang in the lab of Dr. James T. Enns with research ethics approval from the University of British Columbia Behavioural Research Ethics Board (BREB) project file # H11-00946, 09/17/2012.

Acknowledgements

I have way too many acknowledgements to make. But here's a start:

I am extremely grateful for everyone who has helped me in my journey through science the past seven years, and I am deeply appreciative that you have shaped my thinking in your own unique way. Thank you to my supervisors, supervisory committee, and examiners: Craig Chapman, Anthony Singhal, Kyle Mathewson, and Kelvin Jones. Your feedback and guidance have been invaluable in producing this thesis. Most importantly, thank you to Craig Chapman. You have treated me like a colleague rather than a student from day one, and I could not have asked for a better mentor. Thank you for your trust and invaluable insight - you encouraged me to set aside EEG and pursue the model in Chapter 2 despite neither of us knowing anything about modelling. I view this thesis as just me writing down your exciting theories that you have shared with me (plus some weird equations we made up). Finally, thank you for your generous and unending support. You always fix bugs in my code, willingly give up hours of your time to whiteboard ideas, and always have your door open when we in the lab are lost. Even when you are on sabbatical out of the province, your door is always open (I counted 40 full hours of Skype calls for running model simulations and saying "I think we're almost done this time" after every single call).

Thank you to my lab mates - Jenn Bertrand, Ewen Lavoie, Jeff Sawalha, and John Luu - who are stellar colleagues and great friends. You all brighten a windowless room with great discussions and even better jokes. An extra shout out to Jenn for fighting together with me in the trenches of learning EEG. In addition, thank you to everyone outside of the ACELab who took time to help Jenn and I pick up EEG during our first year, including Anthony, Kyle, Sarah Elke, Sayeed Devraj-Kizuk, Jeremy Caplan, and John Lind. Thank you to all the research assistants in the ACELab who have helped collect wonderful data despite complex and sometimes unreliable experimental setups. Special thanks to 2016 summer students Mirna Matta and Alice Xue for your heavy lifting with EEG collection. Further, I'd like to sincerely thank all who participate in our experiments for their time and patience.

Thank you to the (sometimes faceless) network of other researchers, who answer questions, provide toolboxes, guides, and answers. Your (often unrewarded) efforts help tremendously. Thank you to the Summer School for Computational Sensory-Motor Neuroscience (CoSMo) and all associated faculty for their amazing work. The majority of this thesis was only possible because of the intensive tutorials and guidance on computational modelling during those two weeks. Thank you to the 2016 class of CoSMo for the great discussions and even better times afterward. Thank you to everyone at the Department of Psychology, and the rest of the University of Alberta, who have made the U of A such a welcoming and exciting place to do research. This includes the Psychology General Office, members of the Collins, BLINC, RLAI, Singhal, and Mathewson labs.

Thank you to the Department of Psychology, the University of Alberta, the Alberta Gambling Research Institute, and the National Sciences and Engineering Research Council of Canada for funding and resources, without which this work would not have been possible.

Finally, thank you to all the amazing people in my life outside of research. Your love, support, encouragement, and friendship have helped tremendously the past two years. Thank you to all the boys on video games for all the laughs. Thank you to my family for the unending support. And thank you to Miranda for the unwavering encouragement. You are all answers on days when it feels like research has none.

Table of Contents

Abstract	ii
Preface	iii
Acknowledgements	iv
List of Tables	vi
List of Figures	vii
1 - Two Halves of a Whole Decision	1
1.1 - Between a Stimulus and a Response	1
1.2 - Between a “Response” and a Response	10
1.3 - Objectives and Predictions	18
2 - Modelling Movement as an Ongoing Decision	21
2.1 - Methods	22
2.1.1 - Behavioural experiment	22
2.1.2 - Data analysis	24
2.1.3 - Primer on drift diffusion parameters	25
2.1.4 - Modelling of reaction times	28
2.1.5 - Modelling of reach trajectories	31
2.2 - Results	37
2.2.1 - Behavioural results	37
2.2.2 - Model results	38
2.2.3 - Additional model analyses	43
2.3 - Discussion	46
3 - Neural Signals of Decision Making	49
3.1 - Methods	51
3.1.1 - Experiment	51
3.1.2 - EEG collection and preprocessing	55
3.1.2 - Model	56
3.2 - Results	58
3.2.1 - Behavioural results	58
3.2.2 - EEG results	60
3.2.3 - Model results	66
3.3 - Discussion	69
4 - General Discussion	73
References	79
Appendix	89
A.1 - Collapsing across rank difference	89
A.2 - Initial choices and biomechanical bias	90
A.3 - Rejecting bad drift diffusion fits	92
A.4 - Drift diffusion vs. race models	93
A.5 - Tables of model parameters	96

List of Tables

Appendix A.5 - Tables of model parameters

Table A.1: Fitted parameters for Chapter 2 participants.

Table A.2: Fitted parameters for Chapter 3 keypress participants.

List of Figures or Illustrations

Chapter 1: Two Halves of a Whole Decision

Figure 1.1: Mockup of a typical experimental figure in visual psychophysics.

Figure 1.2: Random dot motion task stimuli. Reprinted from Zhang, J. (2012). The effects of evidence bounds on decision-making: theoretical and empirical developments. *Frontiers in Psychology*, 3. Copyright 2012, Zhang.

Figure 1.3: Schematic of a race model (left panels) and a drift diffusion model (right panels) of decision making during a random dot motion task (top row).

Figure 1.4: Architecture of different decision models as depicted in Chen & Stuphorn (2015). Reprinted from Chen, X., & Stuphorn, V. (2015). Sequential selection of economic good and action in medial frontal cortex of macaques during value-based decisions. *Elife*, 4, e09418. Copyright 2015, Chen et al.

Chapter 2: Modelling Movement as an Ongoing Decision

Figure 2.1: Depiction of experiment setup for Experiment 1.

Figure 2.2: Schematic of a five-parameter drift diffusion model.

Figure 2.3: Model schematic with an example from one trial with a change of mind.

Figure 2.4: Comparison of behavioural data and model results for easy (green), medium (blue), and hard (red) conditions. Reaction times z-scored within individuals and collapsed across individuals for both model (dashed lines) and data (solid lines).

Figure 2.5: Comparison of group averaged behavioural data (filled circles) and model (black line) results for easy (green), medium (blue), and hard (red) conditions. A) Group averaged reaction time. B) Group averaged choice accuracy for initial choices (light colors) and final choices (dark colors). C) Group average proportion of changes of mind for corrective changes (dark colors) and erroneous changes (light colors).

Figure 2.6: Average trajectories from correct trials for behavioural data and model simulations.

Figure 2.7: Individual reach trajectories by condition for three selected subjects.

Figure 2.8: A) Average correct reach trajectory data from Chapman et al. (2010) Experiment 1, where participants were presented with one or two possible targets. B) Simulated group data from parameters fit to our 25 participants. Figure adapted from Christopoulos, V., & Schrater, P. R. (2015). Dynamic integration of value information into a common probability currency as a theory for flexible decision making. *PLoS Computational Biology*, 11(9), e1004402. Copyright 2015 Christopoulos, Schrater. Original data from Chapman, C. S., Gallivan, J. P., Wood, D. K., Milne, J. L., Culham, J. C., & Goodale, M. A. (2010). Reaching for the unknown: multiple target encoding and real-time decision-making in a rapid reach task. *Cognition*, 116(2), 168-176. Copyright 2010 Elsevier.

Figure 2.9: Group average reaching trajectories from model simulations. A) Obstacle value simulations. B) Obstacle location simulations.

Chapter 3: Neural signals of decision making

Figure 3.1: A) Experiment setup. B) Diagram of experimental stimuli.

Figure 3.2: A) Average correct reach trajectories for double difference (blue) and single difference (red) conditions. B) Difference plot over 100 normalized reach distance points between double and single difference conditions.

Figure 3.3: RT-locked event-related potentials (ERPs) from midline parietal electrode.

Figure 3.4: Luminance change-locked event-related potentials (ERPs) from midline parietal electrode.

Figure 3.5: RT-locked event-related potentials (ERPs) for keypress participants. A) CPP. B) Lateralized readiness potential (LRP), which is argued to index motor preparedness.

Figure 3.6: Raster plot of CPP voltage locked to the time when stimuli began to change in luminance.

Figure 3.7: A) Reaction times for keyboard responses in our EEG task (filled circles), along with model RTs from a drift diffusion model fit to these data (black lines). B) Accuracy for keyboard responses, and predicted accuracy from model simulations.

Figure 3.8: CPP data from keyboard participants (solid lines with S.E.M.) and drift diffusion traces from model simulations (dotted lines). A) Response-locked data. B) Luminance change-locked data.

Figure 3.9: Raster plot of A) CPP from keyboard response group, and B) drift diffusion traces from model simulations.

Appendix

Figure A.1: Group mean reaction times and accuracy for all six unique conditions, which were subsequently collapsed to three critical conditions based on rank difference.

Figure A.2: All raw trajectories from participant 4.

Figure A.3: Plot of changes of mind over time for all 32 participants.

Figure A.4: Example RT and accuracy fits for three individuals from Chapter 2.

Figure A.5: Grand mean reaction times from behavioural data (black), and theorized race (red) and drift diffusion (blue) models.

1 - Two Halves of a Whole Decision

Would you like a chocolate bar or an apple? Should you rent or buy a house? Do you want to read this thesis right now, or would you rather be doing something else?

Decision making is not only ubiquitous, but has significant consequences. Our decisions directly influence our physical health, financial security, personal well-being, and much more. In this respect, it is important to study decision making - not only to understand why we choose what we do at a fundamental level, but because if we can understand our decisions, we may be able to influence them to make our lives better.

But how exactly do we make decisions? The unfortunate answer is that we don't know. But we have some strong leads.

1.1 - Between a Stimulus and a Response

For almost two centuries, the field of psychophysics has been presenting people with different intensities of a physical stimulus and measuring their responses. For instance, an experimenter would show a participant a light of a certain brightness, and ask whether they thought it was brighter than the last light they saw (not unlike a modern-day eye exam). Most relevant to this thesis, psychologists had been instructing subjects to make decisions regarding difficult to discern stimuli (e.g., whether an arrow is pointing left or right). Time and time again, experiments showed that (unsurprisingly) people took time to make a decision. Further, people take more time to respond when the decision is more difficult (McCarthy & Donchin, 1981). While this result should come as a surprise to no one who has spent time around any human being, it underlies two critical points about ourselves. First, there is a temporal delay between external stimuli and the brain - even automatic reflexes take time (Enns, 2004; Evars & Tanji, 1976). Second, even after the delay between external information and the brain, processing

within the brain takes time. This scaling of response time with decision difficulty has long been thought to indicate that harder decisions take more processing time to resolve an answer (Hyman, 1953). In sum, whatever our brains are computing, it is not instantaneous.

A second, and more surprising result of these experiments is that people consistently make different responses to the same physical stimuli (McNicol, 2005). For example, when asked multiple times if an arrow tilted very slightly to the left was pointing left or right, a person might say “left” 80% of the time, but report “right” the other 20% of the time (see Figure 1.1). This result may seem benign, but it contradicts very influential ideas about human behaviour. For example, neoclassical economics assumes that people have stable and transitive preferences, and make decisions using full information (Henry, 2012). Further, mirroring the effects of decision difficulty on response time, researchers found that people are less accurate when the decision they are faced with is more difficult (McNicol, 2005). Of relevance to this thesis, human inconsistency despite external environmental consistency suggests the human brain is noisy. Noise can mean many things in neuroscience and psychology - predictable fluctuations tangential to what is directly relevant (Howell, 2011), degraded relevant signals (Shadlen & Newsome, 1994), or even random fluctuations (Ando & Graziani, 2012). Regardless of the source of this noise, the fact remains that humans are consistently inconsistent. While these results may seem intuitive, they hint at the architecture behind the fundamental cognitive process of decision making.

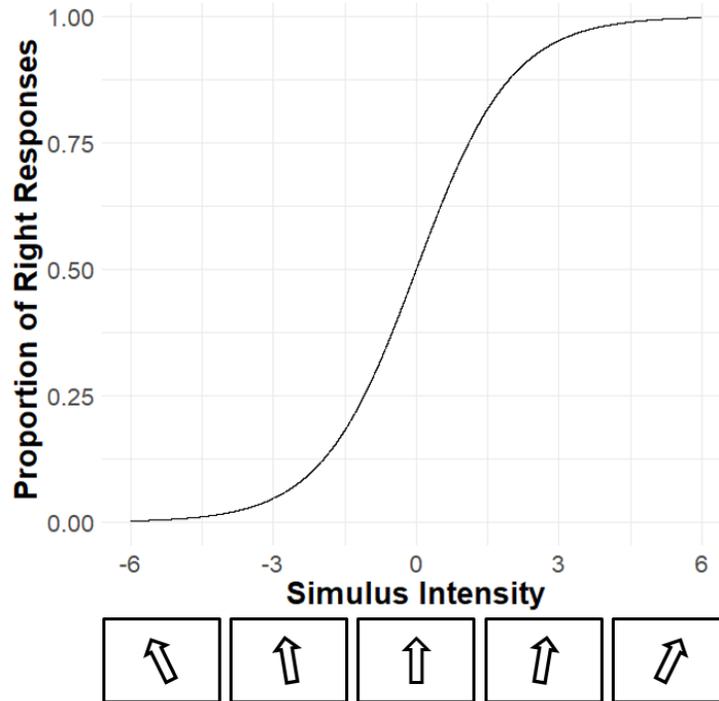


Figure 1.1: Mockup of a typical experimental figure in visual psychophysics. For a given stimulus intensity (e.g., how tilted an arrow is to the left or right, x-axis), when the participant needs to make one of two responses (two-alternative forced choice), there is a corresponding proportion the participant has responded “right” (y-axis). For this mockup, a psychometric function is used instead of data points, which is a widely applied inferential model that captures much of the data in these tasks (Gescheider, 2013). At medium stimulus intensities (e.g., an arrow pointing perfectly in the middle) participants tend to report “left” and “right” equally. The point at which subjects respond with each option equally is named the point of subjective equality, and can be different from true stimulus equality (and reveals that the subject has a bias to respond with one of the options). Of particular note, when the stimulus is slightly deviated from the middle stimulus intensity, participants still report the incorrect response, but less frequently.

Any model of decision making must therefore account for the findings that (1) information to the brain is delayed, (2) cognitive processes take time, and (3) neural processing is noisy. Given these constraints, models must predict that when decisions are more difficult, people tend to (1) be less correct, and (2) take more time to report their decision.

Several formal models have been proposed to explain decision making in humans. The most successful class of models in this regard are sequential sampling models of decision making

(Gold & Shadlen, 2007). Sequential sampling models posit that the human brain takes repeated samples of information relevant to a decision (i.e., evidence) over a period of time, and commits to a decision when there is enough information to support a particular action. The most common task used by modern experiments using sequential sampling models is the random dot motion (RDM; see Figure 1.2) task (Britten, Shadlen, Newsome, & Movshon, 1992). A circle filled with moving dots is presented to a participant, and they are asked to determine if in general the dots are moving left or right. The decision is made easier by manipulating the level of motion coherence (i.e., how many “coherent” dots are moving in the same direction). Additionally, the stimulus is noisy, as the motion of the remaining “non-coherent” dots are in a randomly determined direction. Sequential sampling models state that humans performing this task take a small sample of motion, process this information to determine whether, and how much, the motion favors responding left or right, and then either decide to continue sampling the motion or make a response.

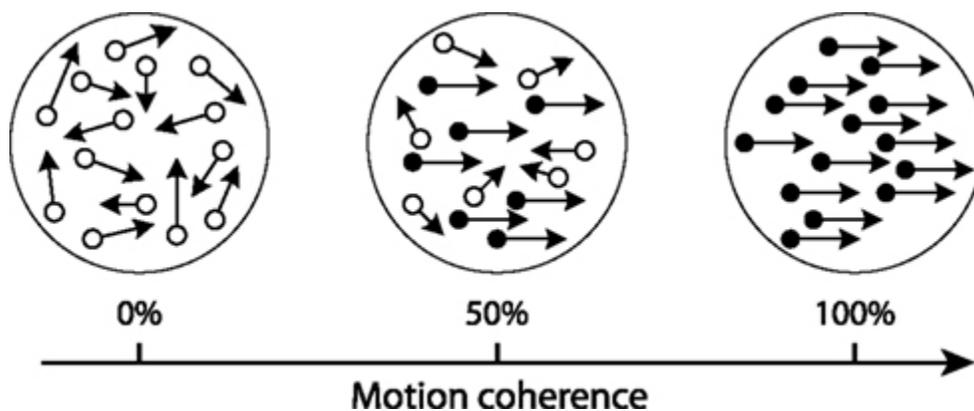


Figure 1.2: Random dot motion task stimuli. Several dots in a circle move in a specified direction and are periodically replaced by new dots. Trials typically have a set level of motion coherence, where a subset of dots move in the same direction (usually directly left or right; depicted as black dots), and the remainder of the dots move in a randomly determined direction (depicted as white dots in this mockup - all dots appear the same in the task). Participants are asked to determine whether the dots are moving to the left or to the right. More motion

coherence makes the left/right motion discrimination decision easier to participants. Reprinted from Zhang, J. (2012). The effects of evidence bounds on decision-making: Theoretical and empirical developments. *Frontiers in Psychology*, 3. Copyright 2012, Zhang.

Two of the most successful and widely used sequential sampling models are race models (Smith & Vickers, 1988) and drift diffusion models (Ratcliff & Rouder, 1998; see Figure 3). Race models generally state that humans sample relevant information and after some neural delay, partition this information into evidence favoring each choice option. To illustrate, in the RDM task, this would mean that after a neural delay, each sequential sample of the world provides some noisy evidence for responding left, and some noisy evidence for responding right. Two “accumulators”, one for left motion and one for right motion, add up all of the past samples favoring their respective direction. When one accumulator reaches some specified threshold value, that accumulator wins the race and its respective action is executed. This accumulation process has several desirable qualities in a decision making model. Namely, it includes some delay between the external stimulus and processing, its processing necessarily takes time, and it specifies that the values of evidence favoring left and right motion are subject to internal noise. In this respect, it considers all three constraints of a foundational decision model. Further, it provides a mechanism for why people make inconsistent choices - internal processing noise. If many dots are moving to the left and very few to the right, but the evidence values favoring right motion over a period of time are high due to random internal noise, then on that trial a right response is more likely even though this does not align with the motion of the stimulus. Race models of decision making not only consider the three constraints previously mentioned, but have been shown to accurately predict the pattern of accuracy and reaction times in decision tasks as well (Bogacz, 2007).

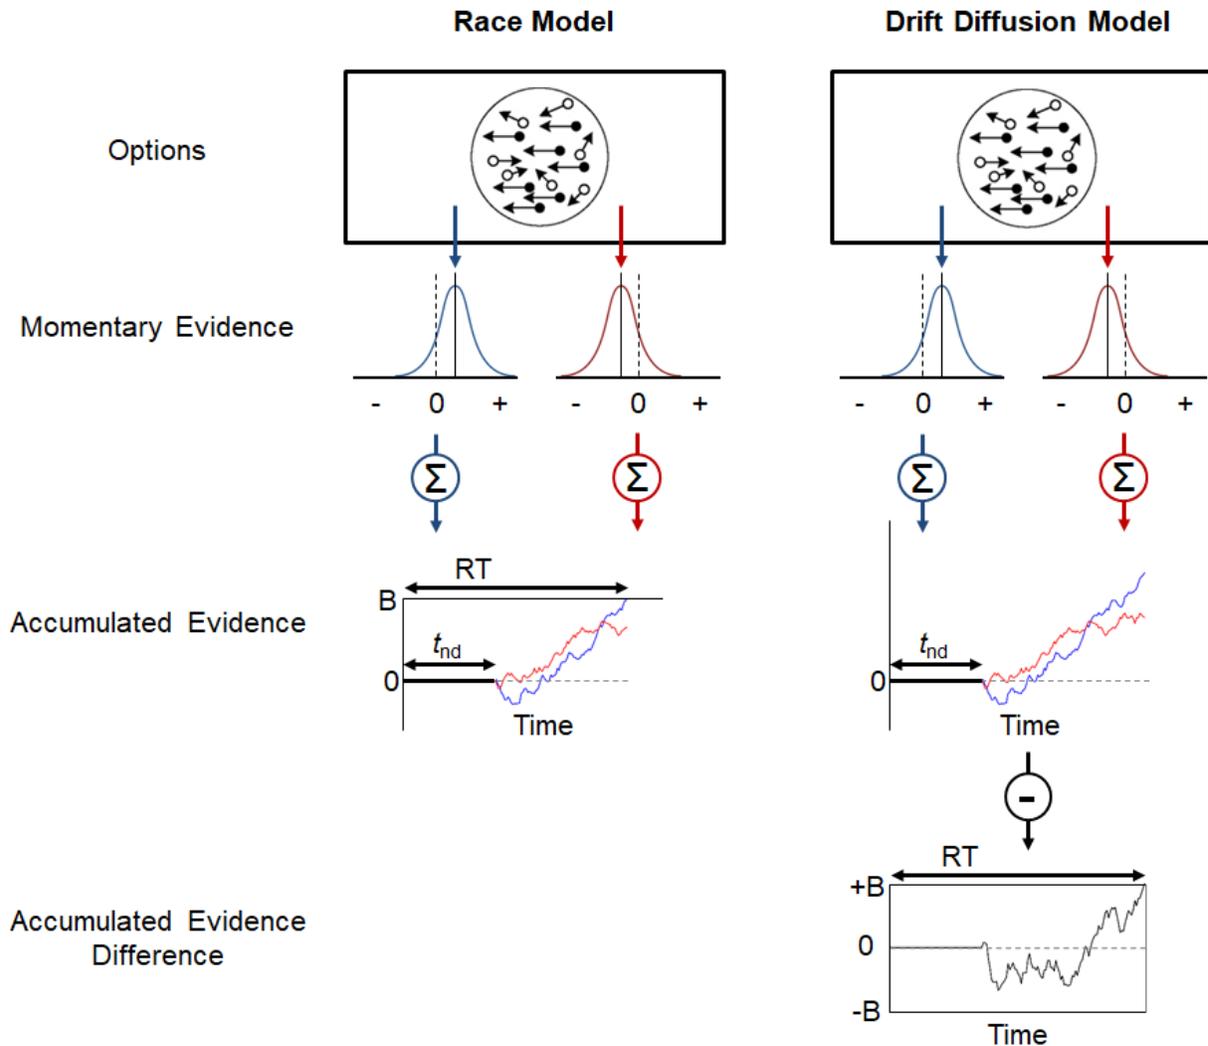


Figure 1.3: Schematic of a race model (left panels) and a drift diffusion model (right panels) of decision making during a random dot motion task (top row). In both models, response options (e.g., “left” in blue, and “right” in red) are represented internally. Evidence for these response options is sampled from the stimulus, which leads to values for that moment favoring responding “left” and favoring responding “right” (second row). Here, the mean evidence for a left decision is higher than evidence for a right decision, as there are on average more dots moving coherently to the left. These momentary samples of evidence are summed up into independent accumulators (third row, here represented on the same plot for simplicity). Evidence arrives after some non-decision time (t_{nd}) because of neural and processing delays. In a race model, evidence is accumulated until one accumulator reaches some threshold (B). In a drift diffusion model, the *difference* in accumulated evidence is taken (fourth row), and a decision is made when the *difference* crosses a threshold ($+/-B$). On this example trial, both response options have accumulated positive evidence at similar rates, and so a race model with an absolute threshold is faster at arriving to a decision than a diffusion model based on relative evidence.

Race models make accurate predictions about human behaviour, and provide an explanation (at a level of abstraction) regarding how decision making may be actually implemented in the human brain. Not only have race models been tested with behavioural data from different behavioural tasks (Smith & Ratcliff, 2004), but neural recordings have shown patterns of activity consistent with the model (Platt & Glimcher, 1999) - a slow buildup of activity until a response is executed (i.e., the accumulation of evidence over time). However, race models fail to account for a specific but important behavioural result. In random dot motion tasks, evidence for one response (e.g., left moving dots) comes at the expense of evidence for the alternative response. In other words, having more left moving dots necessarily means fewer right moving dots. Therefore, both left and right options cannot have high levels of evidence at the same time. This proves to be a problem, as in real life sometimes we need to make choices between two very desirable options (e.g., watching the rest of a playoff hockey game vs. sleeping). While race models predict that high evidence for both of two options would lead to very fast decisions (both accumulators reach bound very quickly), this prediction is not borne out in the data (Bogacz, 2007).

An alternative to the race model that solves this particular issue is the drift diffusion model (Ratcliff & Rouder, 1998). The drift diffusion model has all of the processes and qualities of a race model, but instead uses the *difference* in accumulated evidence as determining reaction times and choices (see “Accumulated Evidence Difference” in Figure 1.3). Independent accumulators for each option are typically still present, but an online difference between these accumulators is calculated. When the difference reaches a positive threshold or a negative threshold, that respective action is selected. Because a difference in evidence is taken, it solves the above problem. Namely, when two highly desirable responses compete, it takes more time

for the difference to reach threshold (and for a response to be reported), and therefore makes this kind of decision a difficult one rather than an easy choice. Drift diffusion models have thus been applied widely and with great success on a range of tasks (Ratcliff, Smith, Brown, & McKoon, 2016).

Some of the strongest support for drift diffusion models, and evidence accumulation models in general, comes from studies of neural recordings. Single cell recordings, predominantly in monkey lateral intraparietal area (LIP), have shown firing rates that match general evidence accumulation models - namely a signal buildup for several hundred milliseconds that drops sharply just before a response (typically a saccade; Gold & Shadlen, 2007; Platt & Glimcher, 1999; Shadlen & Newsome, 2001). Further, studies in humans using MEG (Donner, Siegel, Fries, & Engel, 2009), fMRI (Krueger et al., 2017), and EEG (O'Connell, Dockree, & Kelly, 2012) have shown general parietal signals that match evidence accumulation model predictions, despite the non-invasive nature and noisy signal resolution of these methods.

Of particular note, activity at central parietal electrodes using EEG is consistent with a bounded evidence accumulation model. Termed the centro-parietal positivity (CPP; O'Connell et al., 2012), the timing and build-up rate of this signal are modulated by the strength of decision evidence. Importantly, this signal has several desirable properties for candidate decision signals. First, it is argued to be distinct from observed sensory encoding (steady state visually evoked potentials; SSVEP) and motor preparation (lateralized readiness potential; LRP) signals in EEG. Second, the CPP is observed when making a decision without an overt action (e.g., when reporting decisions at the end of a block of trials). Third, the CPP is observable for several kinds of decision evidence including changes in visual contrast, auditory volume, auditory frequency, and motion using the random dot motion task (Kelly & O'Connell, 2013; O'Connell et al.,

2012). Further, in a Go/No-Go task, parameters of a one-choice bounded accumulation model were related to CPP build-up rate and peak latency (Murphy, Robertson, Harty, & O'Connell, 2015). Finally, the CPP is broadly consistent with single cell recordings from monkey area LIP, as topographies strongly indicate a parietal source for this signal. For this reason, studies have recently argued that the CPP is synonymous with the classic P300 component in EEG research (Twomey et al., 2015). Together, these results suggest the brain may implement a process very much like a drift diffusion model.

While a drift diffusion model of decision making, and evidence accumulation models in general, are extremely appealing, they are not without criticisms and limitations. Notably, diffusion models are unable to account for a pattern of consistent and very early responses (explored in depth by LATER models of decision making; Noorani & Carpenter, 2016). Additionally, evidence accumulation models in general break down when trying to explain choices that take minutes, hours, or even days (such as buying a house). Finally, even today new studies are being published which show decision making is more complex than past theories and models can account for. For example, behavioural data and neural recordings point to the presence of an urgency mechanism (Thura, Beauregard-Racine, Fradet, & Cisek, 2012; Thura & Cisek, 2016). Sampling information takes neural resources, and not executing an action can sometimes lead to a missed opportunity. Researchers have proposed different mechanisms, such as an urgency signal multiplicative to evidence (Thura et al., 2012), or a dynamic threshold that collapses over time (Drugowitsch, Moreno-Bote, Churchland, Shadlen, & Pouget, 2012; Hawkins, Forstmann, Wagenmakers, Ratcliff, & Brown, 2015), in order to push people toward action. On the other hand, longer decisions have shown results inconsistent with a “pure” diffusion model, which have led some to propose that accumulators are leaky in nature (i.e., they

slowly return to baseline over time) and drive individuals to refrain from making a decision (Usher & McClelland, 2001). In sum, while drift diffusion models are arguably our best candidate for how we make decisions, many questions still remain.

1.2 - Between a “Response” and a Response

While cognitive models of decision making inspired by psychophysics are extremely attractive and useful for behaviour in psychology labs, they have until recently failed to consider an entire second half of decision making taking place outside of the lab - moving. When we make choices in our world, most require the execution of a movement that takes at least as long as the theorized decision (which only predicts how long it takes to react, or start moving). For example, whether reaching to grab a jar of creamy or crunchy peanut butter off a grocery shelf, or biking to a restaurant for the evening, both involve a period of time between action onset and the completion of a decision. The vast majority of psychology, even today, views movement as an output to a function, or an arrow protruding from the mysterious black box of cognition.

These approaches largely assume that any movement required to enact a choice (e.g., reaching to the crunchy over the creamy peanut butter) occurs after, and is isolated from, any cognitive deliberation. Such models can be viewed as goods-based models of decision making (see Figure 1.4 A; Padoa-Schioppa, 2011). These include many traditional models of decision making discussed above (e.g., race, drift diffusion) which specify a competition between available options until a threshold determines which action to then plan and execute.

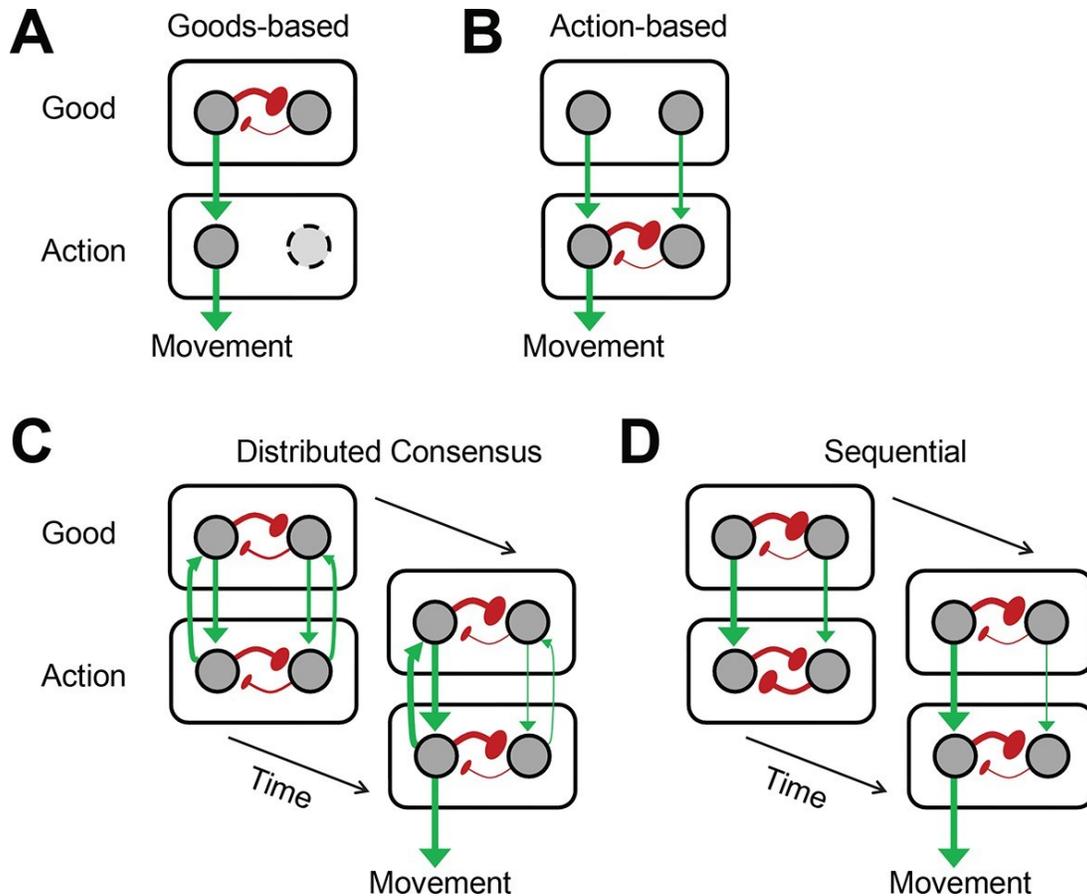


Figure 1.4: Architecture of different decision models as depicted in Chen & Stuphorn (2015). Classic decision models outlined in the previous section fall into A) Goods-based decision making, where a decision is made, and then the respective action is planned and executed. B) Action-based models of decision making state that evidence for each option is input to motor representations for each option, and that it is motor representations that compete for selection. C) A distributed consensus model of decision making states that coordinated competition occurs in many specialized areas, and a decision in one area is often enough to send a cascade through areas in order to come to a decision by consensus. D) Sequential models of decision making state that competition occurs at both the representation of option value, and at motor representations. Information about the ongoing competition regarding the value of available options continuously influences the competition of motor representations (and not the other way around). Reprinted from Chen, X., & Stuphorn, V. (2015). Sequential selection of economic good and action in medial frontal cortex of macaques during value-based decisions. *Elife*, 4, e09418. Copyright 2015, Chen et al.

Goods-based models of decision making have recently undergone a significant change because of the recent surge of interest in “changes of mind” (Albantakis & Deco, 2011;

Kaufman, Churchland, Ryu, & Shenoy, 2015; Kiani, Cueva, Reppas, & Newsome, 2014; Resulaj, Kiani, Wolpert, & Shadlen, 2009; Selen, Shadlen, & Wolpert, 2012; van den Berg et al., 2016) – the infrequent (e.g., 5%) but reliable observation that individuals will often initiate an action toward one choice option, but switch to another choice option before the action is complete. This action switching behaviour mid-movement supports the idea that decision making continues even after an action is initiated. From this observation, Resulaj and colleagues (2009) put forth an extremely successful model which expands a drift diffusion model of decision making into movement. Specifically, when a threshold is reached, an action toward the winning option is initiated, but the drift diffusion process does not stop. If the continuing drift diffusion process crosses a new post-initiation threshold before a specified time into the movement, the participant then switches motor plans to the alternative option. This model not only includes all of the attractiveness of a drift diffusion model, but is also able to explain the frequency of changes of mind (Resulaj et al., 2009), and is able to predict confidence ratings (van den Berg et al., 2016). Further, it holds an intuitive value, as all of us have experienced a moment where our hand may reach toward one option in our environment, only to be pulled toward the alternative by some last second insight.

The changes of mind model (Resulaj et al., 2009) only necessitates a slight change, if any, to the goods-based architecture. Instead of a goods-based competition selecting a single action when a threshold is reached, the goods-based competition continues to compete and sends a switching signal to the action system when a decision has changed. Importantly, this architecture makes an implicit prediction about what neural signals should look like. It would be unnecessary for continuous information to be relayed from the goods-based system to the action-based system, as all the action system cares about is which singular action to take. Therefore, we

should observe a single burst of information from the goods system to the action system whenever an action is to be taken or switched, along with the classic accumulation-to-bound signal in decision systems (e.g., Platt & Glimcher, 1999).

However, studies of neural signals in motor cortex instead support an action-based model of decision making (Figure 1.4 B), where options compete at the level of motor representations. The most prominent theory of this class of models is the affordance competition hypothesis (Cisek, 2007). This framework states that instead of a serial process of decision making, multiple potential actions are represented and compete in parallel. In other words, instead of a decision making process designating a winner to communicate to the motor system, the motor system already represents potential actions to all available options, which themselves compete. One of the benefits of such an architecture is that it affords an adaptive behavioural advantage for changing decisions quickly if our environment changes. While a changes of mind model affords the ability to change actions, an action-based model with both actions already planned is much quicker at this adaptation. Most notably, the affordance competition hypothesis is supported by neural recordings from monkey dorsal premotor cortex (Cisek & Kalaska, 2005). When monkeys are presented with two possible targets to reach toward on a screen, dorsal premotor neural population activity increased in the directions of both the potential targets. Additionally, this increased activity was sustained throughout a period of time when the potential options were removed from the screen, suggesting the activity reflects target representation rather than mere stimuli input. Finally, when told the direction they were to reach toward (through a central color cue), population activity in the correct direction increased further while activity in the incorrect direction was suppressed. This pattern of activity - an increase in activity associated with one option coincident with suppression of the alternative - suggests competition at the level of

premotor cortex representations. Interestingly, downstream activity in primary motor cortex only showed increased activity to the correct direction, suggesting that any representation and competition of parallel motor plans are resolved by primary motor cortex into a single coherent motor command for execution.

The affordance competition hypothesis, and action-based models of decision making in general, are further supported by a wealth of experimental data. If forced to move before participants are cued as to which of two options are correct (termed go-before-you-know tasks), reach trajectories reveal spatial averaging, as they initially move toward the midpoint of options before deviating toward one (Chapman et al., 2010). Additionally, participants show averaging of feedback gains, as participants apply an intermediate corrective force in response to a perturbation of their hand while moving (Gallivan, Logan, Wolpert, & Flanagan, 2016). These averaging effects suggest that if competition between parallel motor plans is not resolved, an average plan is taken as both options are still likely to be the ultimate winner. Further, trajectories have revealed that motor control and movement costs are integrated with decision value (Chapman, Gallivan, & Enns, 2014; Trommershäuser, Maloney, & Landy, 2008; Wolpert & Landy, 2012), suggesting that the motor system has a say in which option is selected. In addition, a recent approach in decision science has used the measurement of continuous movement behaviour to capture decision-making dynamics, revealing how physical movements read out “hidden” cognitive states (Song & Nakayama, 2009). Several recent reviews on this specific topic (Freeman, Dale, & Farmer, 2011; Gallivan & Chapman, 2014; Song & Nakayama, 2009; Spivey & Dale, 2006; Wolpert & Landy, 2012) document the reliable effect that conflict between response options results in similarly conflicted movements (e.g., curved hand or mouse trajectories). For instance, when asked to move a computer mouse to click on one of two images

that matches a read aloud word (e.g., “picture”), mouse trajectories are more direct when choosing between two dissimilar options (e.g., picture and jacket), relative to images that are phonetically similar (e.g., picture and pickle; Song & Nakayama, 2009). These characteristic curved trajectories are unlikely to come from a goods-based model of decision making where movements are binary (i.e., straight toward one option or the alternative), and instead lend support to competition at some level of motor representation.

The most successful computational model within an action-based decision making architecture is able to successfully reproduce movement competition dynamics (Christopoulos, Bonaiuto, & Andersen, 2015), spatial averaging of trajectories in go-before-you-know tasks, and can additionally account for observed neural patterns during reaching tasks (Christopoulos & Schrater, 2015). This “biologically plausible computational theory” integrates spatial sensory input, expected reward of options, and task context into a dynamic neural field, which simulates the activity of hundreds of neural populations each tuned to a different direction in space. These directionally-tuned neuronal populations compete (in this case, laterally excite and inhibit each other), and once a population reaches a specified activity threshold, an optimal control policy (read: motor plan; Todorov & Jordan, 2002) for that respective direction is activated. A weighted average of active policies then determines how the hand moves in space, before the process is updated by a new state of the hand in space.

While this formal model aims to closely mirror the complexity of biological systems, and is able to explain a wealth of behavioural and neural data, it suffers the opposite problem from that of many goods-based models (e.g., the race and diffusion models detailed above). Namely, the model has not yet been shown to predict reaction times and accuracy in decision making tasks. In addition, the complexity of its architecture and the number of parameters (over 30)

make this formalization difficult to validate in order to explain human decision making behaviour and individual differences between decision makers.

Although an action-based model of decision making is an appealing solution to the problem of decision-modulated movements, two additional architectures of decision making models are worth mentioning when discussing the link between the pre- and post-movement initiation halves of decision making. One additional architecture put forth by Cisek is the distributed consensus model of decision making (Figure 1.4 C; Cisek, 2012). This idea states that competition between available options takes place at many loci in the brain at once. A distributed consensus model of decision making is appealing because information can compete in areas that are already specialized for subsets of information (e.g., motion information competes in brain area MT/V5, whereas motor cost information competes in premotor areas). Additionally, this model predicts that decision signals can come from any of these specialized regions, which offers an explanation to the elusiveness of a decision making area in the brain (Katz, Yates, Pillow, & Huk, 2016; Yates, Park, Katz, Pillow, & Huk, 2017). However, there is little current evidence which supports a distributed consensus model of decision making, and studies have shown a lack of coordinated and recurrent connections between value and action information in specific brain regions as predicted by a distributed consensus model (Chen & Stuphorn, 2015).

Finally, Chen and Stuphorn (2015) suggest a fourth kind of decision architecture termed sequential models of decision making (Figure 1.4 D; not to be confused with sequential *sampling* models of decision making). In this architecture, value information competes in the goods-based system, is communicated continuously to the action-based system with some neural delay (50-100 ms), and additionally competes within the action-based system. All potential actions are represented and compete in the action system, much like action-based models of decision

making, but receive inputs from the separate and competitive goods-based system.

The sequential model of decision making is supported by recordings of directionally-tuned neurons in the supplementary eye field (SEF), where monkeys made decisions between different values presented at four possible locations for a saccade (Chen & Stuphorn, 2015). Because target value and physical location were disentangled, the researchers were able to analyze the distinct contribution of value and action on firing rates in the SEF. Classification accuracy analysis showed a ~100 ms delay between value and action information (Chen & Stuphorn, 2015). In other words, neural signals in monkey SEF carry information about what option is going to be selected ~100 ms *before* information about the upcoming saccade direction can be decoded. Further, the authors argue that competition additionally occurs within the action system. Initially, activity in neurons tuned for both potential saccade directions increases, and later this activity is gradually increased in the to-be-selected direction and suppressed in the direction associated with the non-chosen option. However, while this two process model of value and action in decision making is promising, support for this recent model is still relatively limited.

In summary, outside of traditional psychology labs, movement plays an integral part in decision making. Not only do we commonly execute lengthy movements after a period of deliberation, but studies have shown that decision making continues while we move, and even influences how we physically move in space. Why then is it important to identify the exact system with which decision making and movement are linked? Comparing processes and architectures of decision making is important not only because they can explain how our brains are organized, they also predict what we choose, when we choose it, and how we behave in between. For instance, any model without competition between action plans has a very difficult

time explaining the wealth of behavioural effects seen in reach trajectories. Further, any model without goods-based competition implies that there is no “pure” value signal in the brain not modulated by the value of an action in physical space. In sum, although the differences between models and architectures may seem relatively small, the implications of one model over another have vast consequences for human behaviour.

1.3 - Objectives and Predictions

In this thesis I propose a new formal model of decision making, and provide support for the model with both behavioural and neural data. The proposed model is aimed at bridging the gap between the two halves of decision making: pre- and post-movement initiation. In short, this formal model blends a drift diffusion model with competition between motor representations.

In order for this formalization to be a useful model of decision making, it must predict accuracy and reaction times for decisions of varying difficulties. Additionally, for the model to bridge pre- and post-movement initiation, it must also predict movements. Specifically, the model should explain key findings in the literature: changes of mind, curved movements for more difficult decisions, and motor averaging. Further, the model should be flexible enough to explain differences between individuals, and behaviour in a wide range of tasks.

For this model to be appealing, it must also address limitations of competing models of decision making. Competing models (specifically, that of Resulaj et al., 2009) cannot account for predictable fluctuations in reach trajectories when changes of mind are removed, cannot account for effects such as motor averaging behaviour, and are fit to all of the data they aim to predict. The new formalization should be able to address all of these shortcomings. Therefore, in the first part of this thesis (Chapter 2), I will spell out this new model in detail and validate the model by applying it to a behavioural data set from a reach decision task.

In the second part of this thesis (Chapter 3), I will further support this new formal model with neural data. Specifically, models seeking to explain foundational human behaviour do well to constrain possible solutions with plausible neural mechanisms and observed neural patterns. Many solutions to a problem can provide useful predictions, but only a subset of those solutions may possibly be implemented in neural architecture. For example, template matching provides a solution to the problem of letter recognition, but is much less resource efficient than other models and therefore less likely to be implemented in the human brain (Grainger, Rey, & Dufau, 2008). Other models, like those of dopamine and reward processing may have several appealing models for a function of the neurotransmitter (e.g., value coding vs. reward prediction error), but observed firing rates of neurons push one model to be rightly favored over the alternative (Montague, Hyman, & Cohen, 2004). Therefore, we can use neural data to help constrain models of decision making.

While our model is formalized around abstractions of several neural mechanisms, it is important to test key predictions of our model against observed neural data. Our model makes three key predictions about neural signals observed during a reaching task. (1) Our model predicts an accumulation-to-bound signal (specifically a diffusion process) representing a decision to begin movement. (2) The model predicts signals observed in sequential models of decision making (Chen & Stuphorn, 2015), where goods-based information is communicated to the action-based system with some delay. (3) Our model predicts that motor plans compete, and that this competition should be reflected in ongoing reach trajectories. By testing these predictions against observed neural patterns, we should be able to tell what parts of our model are supported by neural data, and more importantly what our model cannot account for. Additionally, applying the model to a new set of behavioural data provides a valuable

opportunity to test model flexibility.

“All models are wrong, but some are useful” (Burnham & Anderson, 2007). While the proposed model is aimed at capturing important aspects of the process of decision making as it is actually implemented in the human brain, it will of course be wrong. The human brain is much more rich and complicated. However, this model is aimed at building upon the important work outlined (and missed) in the above introduction, by fusing concepts to address current limitations in the field. Ultimately, the goal of this thesis is to define decision making as a single, continuous process in which deciding and moving are intimately linked.

2 - Modelling Movement as an Ongoing Decision

Here, we put forth our own formal computational model of how decision making and movement are linked. This model is inspired by action-based and sequential models of decision making (Chen & Stuphorn, 2015), and aims to act as a bridge between models largely focused on only one half of decision making (e.g., pre-movement initiation, Resulaj et al., 2009; post-movement initiation, Christopoulos et al., 2015).

In this chapter, we model decision making as two competitive processes arising from the same evidence: one representing the decision to begin a movement, and one representing the ongoing desirability of potential actions. We adapt a drift diffusion model to explain reaction times of participants in a reaching task, as in Resulaj et al. (2009). The diffusion process acts as a goods-based competition, as options compete based on the accumulating differences in subjective value. The result of the diffusion process is a “go” signal to the action-based system to begin moving when the goods-based system is confident enough in the selection of an option. In addition, we implement a parallel accumulation process in the action system, which begins at the same time as the goods system but has no bound. The action system is continuously fed information from the goods system, and integrates this information into motor representations. These motor representations compete, and once a participant is moving, act to shift the participant's hand toward options proportional to their desirability. In this way, we aim to explain when people begin to move, how they move in physical space, and ultimately what they choose.

Here in Experiment 1, participants were presented with a choice between two food items and were simply asked to reach and touch the one they prefer. While we argue that any decision making mechanism would remain almost entirely the same between this value-based decision making task and other more traditional decision making tasks (e.g., random dot motion task;

commonly termed “perceptual decision making”; Gold & Heekeren, 2014), there are several differences of note. While most other experiments ask participants to make a judgement about a tightly controlled physical stimulus, here participants are instead asked to choose between familiar snack items. Participants in this experiment all have unique preferences for each of the items, which we argue better reflects natural tasks. Further, while other experiments state that evidence for a decision is derived from sampling an external stimulus, the evidence in our experiment is derived from internally generated samples from memory (Shadlen & Shohamy, 2016). While internally generated evidence is argued to follow the same general mechanism, the pathway with which evidence arrives (e.g., visual pathway vs. memory retrieval) is worth keeping in mind. Importantly, many previous tasks have used a noisy stimulus (e.g., random dot motion), and past studies and models have shown that hand trajectories are sensitive to random fluctuations in physical sensory evidence (Resulaj et al., 2009). Here, we use static stimuli, and so any fluctuations in decision making are due entirely to noisy internal processes. Finally, our task involves two spatially separate stimuli as opposed to a single central stimulus (e.g., random dots), which we argue again better reflects natural task requirements and may provide clearer insight into spatial mechanisms for movement.

2.1 - Methods

2.1.1 - Behavioural experiment. Thirty-two right-handed undergraduates made right-hand reaches to indicate which of two projected images of food items they preferred on every trial (see Figure 2.1). There were four unique food items (Dairy Milk, Oh Henry, Doritos, and Lays Original chips) presented in all six non-matching pairs. Each pair was presented 60 times (side counterbalanced) for a total of 360 test trials. On each trial participants started with their finger near a home position and waited 1-2 seconds (random and uniformly distributed) for the

images to appear (and a coincident beep). Images always appeared 40 cm forward from the start position, and were positioned 20 cm to the left and right of the midline. Black outlined squares (± 7.5 cm) appeared around the images to indicate the area where the participant should touch to select that option. Once the choice options were available, the participants had 2.5 s to initiate their movement. Once movement had been initiated, participants had a further 2.5 s to complete their choice movement by touching the option they preferred. Consistent with our intent to have participants comfortably self-initiate their movements and choices, reaction and movement times did not approach these time constraints (RT $M = 442$ ms, range: 52-1527 ms; MT $M = 600$ ms, range: 216-2366 ms). To motivate consistent and veridical preferences, at the end of the test session, two trials were selected at random and the treats selected for those particular trials were given to the participant (consistent with past value-based decision making studies; e.g., Krajbich, Armel, & Rangel, 2010). Finally, at the end of the reaching trials, participants rated how much they liked the treats on a 9-point scale, from -4 (dislike very much) to +4 (like very much), and rank ordered the 4 treat items from most to least preferred. The rankings were used to collapse trials into three decision difficulties: easy (rank difference of 3; i.e., 1 vs. 4), medium (rank difference of 2; i.e., 1 vs. 3, 2 vs. 4), and hard (rank difference 1; i.e., 1 vs. 2, 2 vs. 3, 3 vs. 4). Collapsing conditions into easy, medium, and hard was additionally supported by an analysis of accuracy and reaction times (see Appendix A.1). In this task, correct choices were defined as the trials where participants chose the more preferred option (according to their self-report at the end of the experiment). This is a common definition in value-based decision making research (Krajbich & Rangel, 2011), but nonetheless assumes that participants' preferences are stable across the experiment, and that their self-report matches their true preferences. Finally, unlike many other reaching experiments, this task is intuitive and simple for participants, as they are

asked to simply “touch what you like when it appears.”

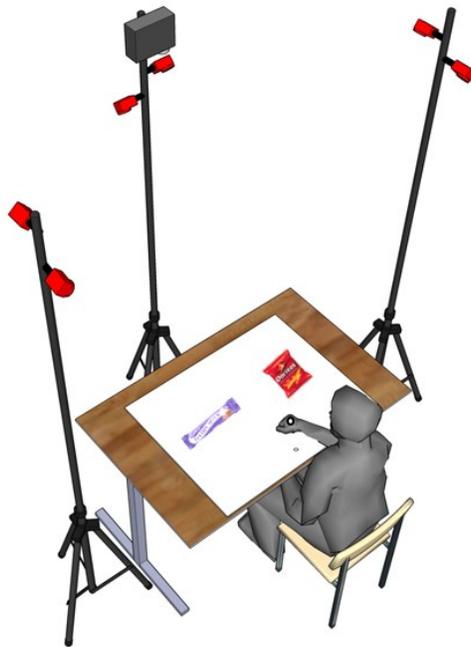


Figure 2.1: Depiction of experiment setup for Experiment 1. Six motion tracking cameras captured the three-dimensional position of a reflective marker on the participant’s right index finger. A projector displayed the start position and available options on each trial onto the table in front of the participant.

2.1.2 - Data analysis. Participants’ responses were recorded at 60 Hz using Optitrak motion tracking cameras via a reflective motion-tracking marker placed near the tip of their right index finger. Reaction time was defined as the time between choice option presentation and when a participant started moving, determined by velocity and distance criteria (Gallivan & Chapman, 2014). Movement time was defined as the time between movement onset and when the participant’s finger touched one of the two choice options. Consistent with past research, a change of mind was said to have occurred if the area between the midline and the hand trajectory on the side of space opposite to the final chosen side of space exceeded 0.1 cm^2 (Resulaj et al., 2009; van den Berg et al., 2016). Trials with a reaction time $< 100 \text{ ms}$ were rejected (2.4% of all

trials), as were trials where the participant initiated movement before the choice options appeared or were slower than the experimentally mandated 2.5 s reaction and movement time limits (6.6%). Initial choices were defined as the side of space of the hand at 20% through each reaching movement (~6 cm). This criterion was chosen to reduce biomechanical biases that significantly inflated estimated changes of mind (see Appendix A.2). As such, changes of mind were calculated exclusively after 20% of each reaching movement for both behavioural data and model simulated data.

2.1.3 - Primer on drift diffusion parameters. We model a drift diffusion process as determining when to begin a reaching movement (as in Resulaj et al., 2009). Our specific implementation relies on 5 model parameters. Before we explain our model in depth, it may be useful for some readers to familiarize themselves with what these parameters mean, and their impact upon the process and results.

A diffusion process typically begins with a randomly sampled non-decision time on each simulated trial (see Figure 2.2). This non-decision time acts to simulate the neural delay between external information and the decision process (and typically includes delay between deciding on an action and responding using a keypress or saccade as well). Non-decision time is sampled from a normal distribution with mean t_{nd} and some variance (usually a fixed number). t_{nd} is the first parameter of the diffusion process, which can vary between individuals (i.e., some people are faster at processing information than others) and tasks (i.e., some tasks take less time for information to reach the decision system than others). All else held equal, longer non-decision times result in longer reaction times, and shorter non-decision times result in shorter reaction times.

The second parameter is sensitivity to evidence (k). In perceptual decision making, an

exact value of external information is given to the participant (e.g., luminance in candela per square metre). Remember from Chapter 1 that in evidence accumulation models, evidence at each time point is randomly sampled from a gaussian distribution with a mean proportional to external evidence (e.g., for a luminance of 400 cd/m^2 , evidence is sampled from a distribution with mean 400). However, participants each have their own sensitivity to this evidence, which acts as a multiplier on the external evidence to determine the internal impact of this information (e.g., $k * 400 \text{ cd/m}^2 = kC = u = \text{mean evidence}$). A greater sensitivity value results in more extreme evidence, while a smaller sensitivity value means less extreme impact of external information on evidence accumulation. When k is zero, external information has no impact on evidence accumulation, which results in random choices without regard for external information.

The third and final parameter in a “pure” drift diffusion model is the decision bound (B). This bound, or threshold, determines how much information is needed in order to commit to an action. When the difference in accumulated evidence crosses the upper or lower bound on a simulated trial, that respective action is taken. Here, $+B$ represents the left option, and $-B$ represents the right option. A higher decision threshold means that participants are more conservative, and need a greater difference in evidence between options to commit to a decision. This hesitation to respond means participants are more likely to be correct, but take more time to act. Conversely, a lower decision threshold means participants require little difference in evidence favoring one option to commit to it, which makes decisions quicker and more impulsive. In this way, a drift diffusion model can be analogous to signal detection theory (Pleskac & Busemeyer, 2010).

Additionally, we implement two more commonly used parameters which both act to explain different aspects of response bias. The initial drift offset (y_0) determines where the drift

diffusion process initially starts. If y_0 is zero, then there is no bias, and the model simply reduces to the three-parameter “pure” diffusion model. However, if y_0 is positive, then less of a difference in evidence is needed in order to reach $+B$, and relatively more evidence is needed to reach $-B$. In this way, the initial drift offset biases participants to choose the option on one side of space faster, and more often.

Finally, the second response bias parameter (and fifth and final parameter in our implementation) is drift bias (u_0). This bias parameter influences the mean of the momentary evidence distributions on each side of space. Remember that evidence at each time point is sampled from a gaussian distribution with mean $u = kC$, and some variance, where C is the amount of external evidence (e.g., 400 cd/m²), and k is the person’s sensitivity to that evidence. Drift bias acts to shift the mean of these distributions, such that the above equation is now $u = kC + u_0$ for the left option and $u = kC - u_0$ for the right option. For instance, if u_0 is positive, then the mean of evidence for the option on the left will be positively shifted, while the mean of evidence favoring the right option will be negatively shifted. In this way, drift bias acts to constantly pull the evidence for one option toward threshold, and push the alternative away. Like initial drift offset (y_0), drift bias (u_0) also acts so that one side of space is chosen more often than the other, and the reaction times for the biased side are faster. The critical difference between these two parameters is the nuanced way in which they shape reaction time distributions for left and right choices.

To summarize, we implement five parameters for our drift diffusion model: non-decision time (t_{nd}), sensitivity to evidence (k), movement initiation bound (B), initial drift offset (y_0), and drift bias (u_0). The first three parameters are part of a “pure” drift diffusion model (Ratcliff & Rouder, 1998), while the last two parameters give us flexibility to account for side of space

biases commonly seen in reaching and decision making tasks (Gallivan & Chapman, 2014).

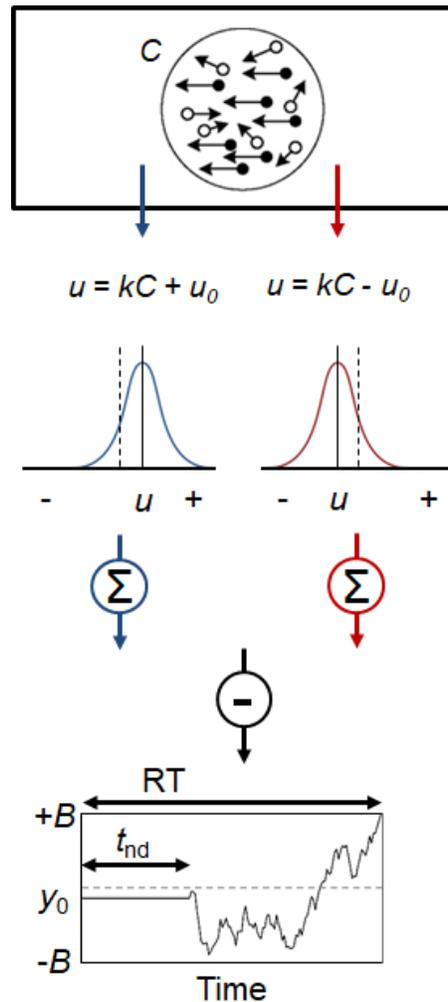


Figure 2.2: Schematic of a five-parameter drift diffusion model. An external stimulus (C) is sampled by the system, which has some sensitivity to this information (k). Along with a constant side of space bias (drift bias, u_0), this determines the mean of a random samples of evidence ($u = kC + u_0$). After some time between stimulus onset and this information reaching the decision system (non-decision time, t_{nd}), sequential sampling begins. At every time point, the stimulus is sampled, which results in a noisy value of evidence for choosing left and for choosing right. Evidence is accumulated, and a difference in accumulated evidence is calculated. When the difference in accumulated evidence crosses a decision threshold ($\pm B$), the corresponding response is made.

2.1.4 - Modelling of reaction times. Here, we model decision making and movement as effectively two accumulation processes arising from the same evidence: one representing the

decision to begin moving, and one representing the ongoing desirability of choice options in physical space for movement. A drift diffusion process determines reaction times (as in Resulaj et al., 2009), and another accumulation process based on the same evidence determines how much each motor representation influences reach trajectories.

We implemented a simplified drift diffusion processes (Gold & Shadlen, 2007; Palmer, Huk, & Shadlen, 2005; Resulaj et al., 2009; Smith & Ratcliff, 2004) that accumulated evidence for choosing the left option and choosing the right option on each trial. Starting at y_0 (initial drift offset) the difference between the two accumulators is computed, and the process ends when the difference between accumulators reached a fixed bound ($\pm B$). The time it takes for the difference between accumulators to reach bound (t_d), along with some sensory and motor delay (t_{nd}) determines movement initiation time after stimuli are presented (i.e., reaction time = $t_d + t_{nd}$).

Evidence at each sequential sample for each accumulator was drawn from individual normal distributions with mean $u = kC + u_0$, and a standard deviation of σ , where C is the subjective value of the option at that location (empirically, from the rank-order questionnaire data, so values of C could range from 0.25, the lowest ranked snack food, to 1, the best ranked snack food), k is sensitivity to evidence, u_0 is the constant bias to choose an option on a particular side of space (drift bias), and σ is evidence noise per unit time. This is a common implementation of a drift diffusion model, and has been shown to explain initial choices (which side of space the hand begins reaching to initially), final choices (which option the participant ultimately selects), reaction times using key presses, eye movements, and reaching movements, as well as neural recordings (Smith & Ratcliff, 2004; Palmer, Huk, Shadlen, 2005; Gold & Shadlen, 2007; Resulaj et al., 2009).

We fit five parameters for each participant: sensitivity to evidence (k), movement initiation bound (B), non-decision time (t_{nd}), drift bias (u_0), and initial drift offset (y_0). Models with similar parameters have been previously found to give acceptable fits for initial choice behaviour in a similar reaching task (Resulaj et al., 2009; van den Berg et al., 2016). Fits were obtained by minimizing the negative log likelihood of initial choice accuracy using a binomial distribution, and mean reaction times using a Gaussian distribution, for each of the three conditions (easy, medium, hard), for each side of space (left preferred, right preferred), and for each of the 32 participants. Consistent with past models, we only fitted reaction times from correct trials (van den Berg et al., 2016) because of the exclusion of a collapsing bound (e.g., Drugowitsch et al., 2012) in our formalization. Fits were computed using `fminsearchbnd` in MATLAB. One set of five drift diffusion parameters was used as an initial guess, and then `fminsearchbnd` attempted to find the best fitting parameters starting from these initial guess parameters. This process was repeated with 32 different sets of initial parameters, which results in 32 different solutions for each person. The solution with the best fit was used for that participant. 32 runs were used to increase the probability the final parameters for each person were at global minimum.

To reduce the number of parameters that needed to be fit and to remain consistent with past accumulation models of reaching decision making, we fixed evidence noise per timestep (σ) to 1 SD per second, and set the standard deviation of non-decision times (σ_{nd}) to 60 ms (Burk et al., 2014; van den Berg et al., 2016).

Fits of drift diffusion model parameters were obtained by simulating 10,000 trials at 500 Hz for each condition (easy, medium and hard), for each side (more preferred option on the left and more preferred option on the right). This large number of simulated trials was used to assure

stability in the estimates of initial choices and reaction times, as the process is noisy. Seven participants (21.9%) were removed from model analysis because of the inability of their drift diffusion model fits to predict reaction time, or initial choice accuracy. Prediction ability was quantified using the proportion of variance that the model could account for in the behavioural data (i.e., coefficient of determination; R^2). R^2 was calculated by dividing the sum of squared residuals (sum of squared deviations between easy, medium, and hard conditions for model and data) by the total sum of squares (sum of squared deviations between easy, medium, and hard conditions in the data about their own grand mean). Any participants where the R^2 of both reaction time and initial choices was less than zero was excluded (see Appendix A.3). Poor fits are not unexpected given our relatively small trials per participant. Without a reasonable drift diffusion fit for these participants, modelling of reach trajectories for these participants will fail, and therefore these participants were excluded.

2.1.5 - Modelling of reach trajectories. Here, we conceptualize movements as reflecting an ongoing, competitive, and graded accumulation of evidence. This simple formalization is inspired by more complex network models of decision making (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Christopoulos et al., 2015), and recordings showing competition between options in premotor cortex (Cisek & Kalaska, 2005).

Starting at the same time as the drift diffusion model (i.e., choice option onset), two independent “motor” processes accumulate the difference of accumulated evidence at each timepoint.

$$m_{Left}(t + 1) = m_{Left}(t) + a_{Left}(t) - R a_{Right}(t) \quad (1)$$

$$m_{Right}(t + 1) = m_{Right}(t) + a_{Right}(t) - R a_{Left}(t) \quad (2)$$

Where a_{Left} and a_{Right} are the independent evidence accumulators at time point t , m_{Left} and

m_{Right} are accumulated differences in accumulated evidence, and R is a free parameter (which we fit to the actual reach data, see below) that dictates the degree of competition between motor plans. An R value of 0 means that the accumulator for the left option is unaffected by the right option, and vice versa, and so these accumulators continue to increase with respect to their accumulated evidence. Alternatively, an R value of 1 means the difference between accumulated evidence for the options is added to the winner and subtracted from the loser, assuring any bias for one option comes at the equal bias away from the alternative. R can be conceptualized as the resistance in an electrical circuit, or in a neural network framework as the weight of an inhibitory connection from the source of evidence for one motor representation onto the competing motor representation.

The m_{Left} and m_{Right} values are then used as weights in a circular mean, where the angles are reach angles pointing straight toward both targets. When m_{Left} and m_{Right} are both positive, the resulting mean reach angle will be somewhere in between both targets. When one of m_{Left} and m_{Right} dominates the other, the resulting reach angle will be straight toward that target. In this way, reach trajectories are an average of motor representations weighted by their current desirability.

One important aspect of using a weighted circular mean to determine reach angles is if a weight is negative, mean angles move away from the respective target. In this way, our model already incorporates obstacle avoidance. However, research using obstacle avoidance tasks have shown that people still reach toward targets even when very undesirable obstacles are in the way (Chapman & Goodale, 2008; Fajen & Warren, 2003). To account for asymmetries in approaching targets and avoiding obstacles, any negative weights were attenuated by distance. Specifically, negative weights were multiplied by the squared distance to the obstacle divided by

total reach length (i.e., the distance from start position to the targets). As such, spatial avoidance of an obstacle decreases as the simulated hand moves farther away from the obstacle. In other words, something undesirable far away has no effect on movements, while something undesirable close by biases movements away from that location.

For our model to generalize to go-before-you-know tasks (Chapman et al., 2010), m_{Left} and m_{Right} start at insignificantly small positive values. Neural recordings show that before option presentation, potential target directions are already positively weighted (Cisek & Kalaska, 2005). Further, network models have changed this anticipatory weighting (using reinforcement learning) to explain eye gazes, hand trajectories, and motor averaging behaviour (Christopoulos et al., 2015). For our current model, adding this anticipatory start value as another free parameter fit to the data (along with R) did not improve our ability to predict reach behaviour. Therefore, instead of fitting this value, an insignificant positive starting value (eps in MATLAB) was used for all participants so both potential locations for a target were positively weighted before option onset. This value is supported by several studies, and assures that if participants begin reaching with no information (go-before-you-know), they equally approach all potential target locations.

To approximate motor noise, gaussian noise ($SD = 10^\circ$) was added to each reach angle at each time step. This was done to approximate variability in reach trajectories, seen even on single target trials (Chapman et al., 2010). Additionally, to approximate the smooth and sluggish changes in reaching movements in our behavioural data, trajectories were smoothed online with a weighted exponential moving average over the last six timepoints. Smoothing over past reach angles was performed simply because sharp and instantaneous changes in reach direction are biologically impossible. For simplicity, reach movements were modelled as having a constant velocity. Reach distance for each time step was calculated as the distance from the start position

to average reach endpoint divided by minimum movement time ($M = 101$ cm/s). This velocity was calculated for each participant individually. As such, we are purposefully unable to account for movement times and are extremely limited in our ability to account for biomechanical biases or effects.

We estimated the value of R , the competitiveness of motor representations, for each of the remaining 25 participants. R fits were obtained for each person by minimizing the negative log likelihood of two aspects of reaching behaviour we argue captures the competitiveness of motor representations. First, the difference in average correct reach trajectories for hard vs. easy, and hard vs. medium trajectories at 50% of the reaching movement was fit using Gaussian distributions. This aspect was chosen because it captures how much trajectories were forced to be straight by midway to the target (i.e., how often the competing motor representation was suppressed below a weight of zero). Second, we fit to the proportion of reaches on hard trials initially heading left, right, or center (based on ± 2 SD of reach trajectories on easy trials at 20% of the reaching movement). We argue this aspect estimates motor representation competitiveness at the onset of reaches. These values were modeled using a multinomial distribution. Fitting R to these two aspects of reaching behaviour helps estimate the degree to which motor plans compete throughout time (i.e., at reach onset and at midway to the target). R fits were performed after fitting the five drift diffusion parameters for each subject, and therefore did not influence the pre-movement decision process. Fits were computed using `fminsearchbnd` in MATLAB using 12 initial seed parameters, and simulated using 1000 reaches per condition per side per person.

Our formalization is a gross simplification intended to approximate the result of much more involved neural processing. We expect such a process implemented at the circuit level to

Figure 2.3: A) Model schematic with an example from one trial with a change of mind. B) Simulated reach trajectory from the example trial in the model schematic. The model uses a drift diffusion process, where options are sampled to derive evidence for choosing each option. Evidence is accumulated, and when a difference between accumulators reaches a threshold ($\pm B$), a movement is initiated. In parallel, the independent evidence accumulators feed into additional accumulators for motor representations of the options. Accumulated evidence boosts the related motor representation, and inhibits the competing motor representation as determined by the R parameter. When movement starts, the values of the motor representations are used as weights for a circular average between reach angles to each option. The resultant vector is the intended trajectory for that time step. In this example, initially evidence favors the dairy milk bar (blue) and eventually switches to prefer the doritos (red) which determines movement initiation. Evidence continues to accumulate during the reach, and eventually switches to again prefer the dairy milk bar shortly after movement initiation. When the evidence that signals the preferred option has switched is enough to overcome and switch the motor representation values, the reach angle begins to shift toward the dairy milk bar.

2.2 - Results

2.2.1 - Behavioural results. All p -values are greenhouse-geisser corrected where applicable. Repeated measures ANOVA showed that, as expected, participants initiated their movements faster for easy decisions relative to medium decisions, and medium decisions relative to hard decisions ($F(2,62) = 28.43, p = 1.32e-8$; all multiple comparisons $p < .001$). Further, repeated measures ANOVAs supported that participants were more accurate for easier decisions in the expected pattern, both for initial accuracy ($F(2,62) = 48.39, p = 1.80e-11$, all multiple comparisons $p < .02$), and final accuracy ($F(2,62) = 19.37, p = 9.37e-5$, all multiple comparisons $p < .002$). Average correct reach trajectories were more curved for more difficult decisions (repeated measures ANOVA, $F(2,62) = 29.82, p = 1.31e-10$), with easy correct trajectories straighter than medium and medium trajectories straighter than hard (all $ps < .017$). Although the spatial differences between conditions at 50% of reaching movement are <1 cm, all differences are statistically significant.

Repeated measures ANOVA showed a significant difference in the amount of changes of mind between decision conditions ($F(2,62) = 25.41, p = 4.22e-8$), with more changes of mind during hard decisions than medium decisions ($t(31) = 5.82, p = 2.03e-6$), but not significantly more changes of mind during medium decisions than easy decisions ($t(31) = 1.97, p = .058$). Further, there were more corrective (initial choice incorrect and final choice correct) than erroneous (initial choice correct and final choice incorrect) changes of mind ($t(31) = 5.10, p = 1.60e-5$), replicating a key finding in past research that additional decision information during a reaching movement is used to improve accuracy (Resulaj et al., 2009).

To further test between a goods-based model where a single trajectory is able to switch online (Resulaj et al., 2009), and an action-based model where trajectories can be a mixture of multiple motor plans, we removed all changes of mind from the behavioural data. Despite the removal of all trajectories identified where participants switched motor plans, we still find a difference in average reach trajectories at midpoint (repeated measures ANOVA, $F(2,62) = 5.25, p = .011$). Multiple comparisons show a significant difference in the expected direction between easy and hard decisions ($t(31) = 4.01, p = .00035$), but no difference between medium and hard ($t(31) = 1.83, p = .077$), nor easy and medium ($t(31) = 1.15, p = .26$). This result suggests that a switching mechanism does not account for an important aspect of reaching trajectories, as there is still some trajectory modulation with decision difficulty, even after removing all changes of mind.

2.2.2 - Model results. Here we directly compare data generated by model simulations to actual participant behaviour. The similarity between the model and behavioural data is reported as the proportion of variance that the model explains in the behavioural data (R^2). All R^2 values are reported first as the average proportion of variance of individual participants' behavioural

data accounted for by the model, and second as the proportion of variance in the group mean data accounted for by the group mean of model simulations. We expect the average R^2 of individuals to be smaller than that of the grand mean, as it is harder for the model to fit small samples of data from variable individuals than general patterns of decision behaviour in a population.

Our model is able to accurately account for movement initiation times ($R^2 = 0.69; 0.95$) and initial choices ($R^2 = 0.47; 0.80$). While these results are to be expected - our model was fit to these data, the drift diffusion framework has been shown to account for these behaviours in other work (Resulaj et al., 2009), and participants with particularly poor fits were removed - it is significant that these accurate fits extend to value-based decision making in a reaching task with different effectors and task geometry.

Critically, our model extends beyond movement initiation to explain several statistics of reaching and decision behaviour. Our model is able to explain the proportion of correct final choices, or what people ultimately choose ($R^2 = 0.29; 0.77$). Further, our model is able to predict the proportion of changes of mind for a given decision difficulty ($R^2 = 0.45; 0.86$), and reproduces the finding that changes of mind are more likely to correct for an initial error than switch from an initially correct movement, $t(24) = 5.85, p = 4.97e-6$ (see Figure 2.5). The similarity between changes of mind observed in our model and in the behavioural data is especially noteworthy, as our model was not fit to these characteristics, unlike previous models (Resulaj et al., 2009; van den Berg et al., 2016). In other words, the pattern of changes of mind simply falls out of our model when it was fit to align with other aspects of behaviour.

Average simulated reach trajectories are more curved for more difficult decisions as in the behavioural data (repeated measures ANOVA, $F(2,48) = 76.37, p = 1.22e-15$), with easy correct trajectories straighter than medium, and medium trajectories straighter than hard (all $ps <$

1.13e-8). Finally, we compared the pattern of average trajectories between the model and data. For left and right reaches for all three conditions (easy, medium, hard), we collapsed the lateral deviation for each averaged trajectory (100 normalized time points) into one array (1 x 600 each for model and data). In this way, the model accounts for a large proportion of variance of average reach trajectories as seen in Figure 2.6 ($R^2 = 0.97; 0.98$). While our model was partially fit to the differences between average trajectories at midline, the similarity between model and data average trajectories throughout the whole reach is especially noteworthy.

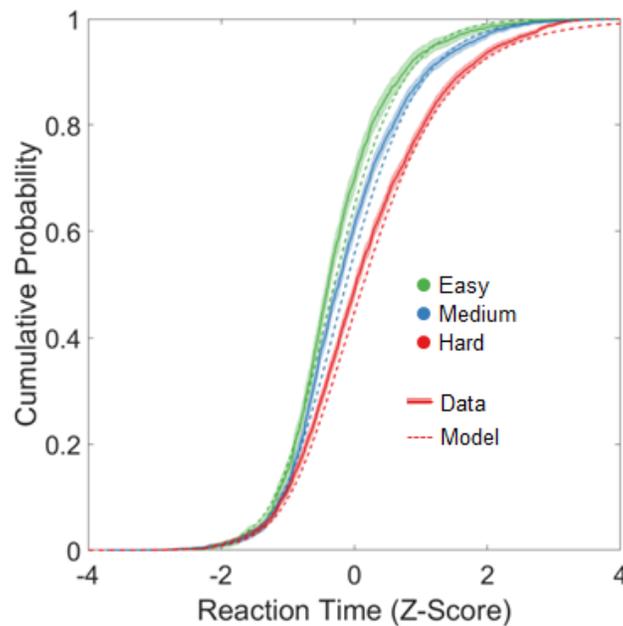


Figure 2.4: Comparison of behavioural data and model results for easy (green), medium (blue), and hard (red) conditions. Reaction times z-scored within individuals and collapsed across individuals for both model (dashed line) and data (solid line). Despite fitting only to mean reaction times, the distribution of model RTs closely align with that of the behavioural data - a powerful aspect of a drift diffusion model.

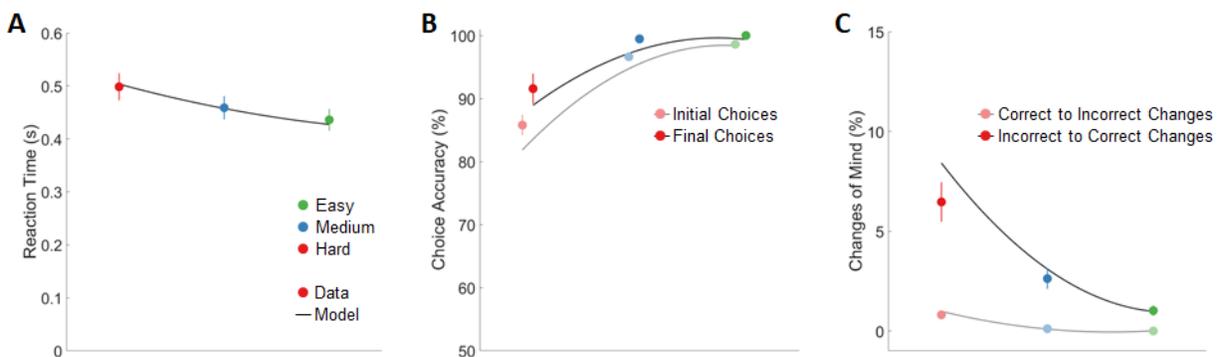


Figure 2.5: Comparison of group averaged behavioural data (filled circles) and model (black line) results for easy (green), medium (blue), and hard (red) conditions. All error bars are S.E.M. A) Group averaged reaction time ($R^2 = 0.95$). B) Group averaged choice accuracy for initial choices (light colors; $R^2 = 0.80$) and final choices (dark colors; $R^2 = 0.77$). C) Group average proportion of changes of mind ($R^2 = 0.86$) for corrective changes (dark colors) and erroneous changes (light colors).

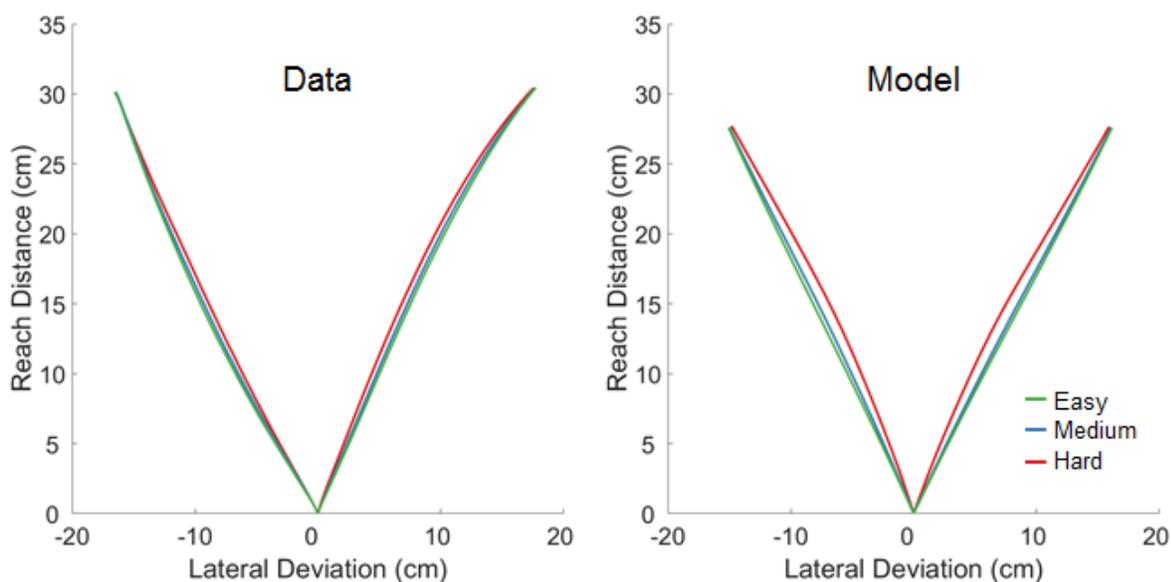


Figure 2.6: Average trajectories from correct trials (calculated as in Gallivan & Chapman, 2014) for behavioural data and model simulations ($R^2 = 0.98$).

Qualitatively, the model is also able to reproduce idiosyncratic reaching behaviour (see Figure 2.7). For instance, participant 16 displayed very straight reaches, almost all of which were correct. In contrast, participant 8 had many curved reaches, suggesting that ongoing competition

between options was resolved during movement more often than observed in participant 16. Further, multiple changes of mind within a single reach can be seen in both data and model trajectories for several participants, including participant 17. Additionally, when the left option was more preferred, participant 17 often began reaching to the right side of space before quickly moving their hand to the more preferred, left side of space. The interaction between drift diffusion parameters and the R parameter seems to capture much of the qualities seen in an individual participant's reaching movements during decision making.



Figure 2.7: Randomly selected individual reach trajectories by condition for three selected subjects. Dark trajectories represent reaches ending to the left, while light trajectories represent reaches ending to the right. Solid trajectories indicate a correct decision, and dashed trajectories

indicate an incorrect decision.

2.2.3 - Additional model analyses. Our model is highly generalizable to several other influential motor and decision making tasks. Here we use the same model and set of fitted subject parameters as representative of decision making in a general reaching study sample.

With minor changes, we can successfully replicate the pattern of results seen in Chapman et al. (2010). This “go-before-you-know” task and its hallmark spatial averaging results have been used to argue that multiple motor plans are maintained in parallel, and that evolving evidence during movement time shapes reaching behaviour. In Chapman et al. (2010), the correct reaching target was filled in at movement onset (i.e., 0 ms). Here, instead of movement initiation beginning when the drift diffusion process crosses threshold, we force our model to begin a reaching movement 100 ms after stimuli onset. Additionally, during single target trials the alternative motor plan is inactivated. With these two changes, we are able to reproduce the spatial averaging seen in go-before-you-know experiments (see Figure 2.8).

We find that the model captures the pattern observed in go-before-you-know tasks best if the model begins reaching movements at ~100 ms after stimuli onset, rather than at 0 ms as in Chapman et al. (2010). This may reflect a difference between stimuli in our task and that of Chapman et al. (2010). Above, we used snack food items in a value-based decision making task, whereas go-before-you-know tasks use low level visual information to cue a participant to the correct target. It is likely that evidence from low level visual features influences the decision system faster than evidence about personal preferences for snack foods. The non-decision times (t_{nd}) in our model are fit to value-based information, and therefore our model likely overestimates the delay between stimuli onset and low level visual information beginning to influence reaching movements. This difference between our model and the task in Chapman et

al. (2010) task may explain why our model performs better if given another 100 ms. Regardless, the similarity between average trajectories in a task far removed from pointing at snack foods is striking.

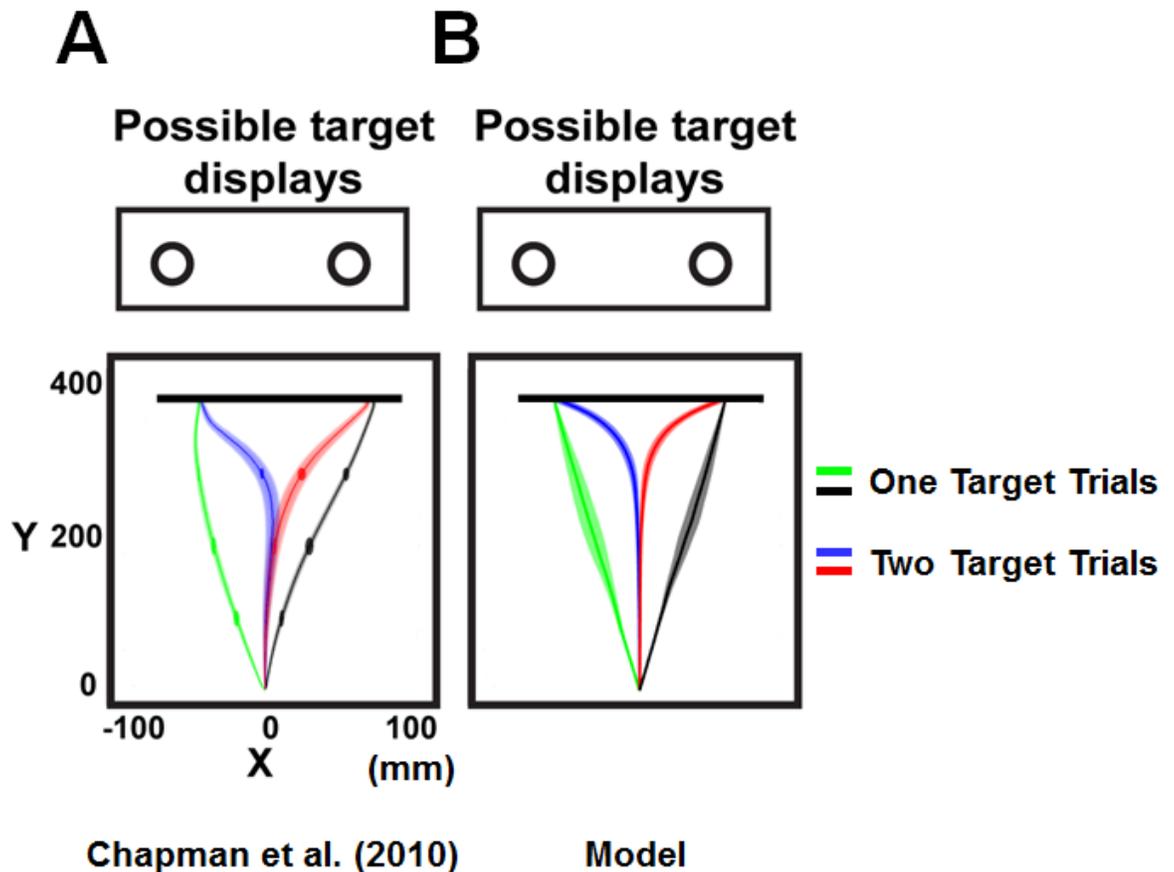


Figure 2.8: A) Average correct reach trajectory data from Chapman et al. (2010) Experiment 1, where participants were presented with one or two possible targets. When participants began a reaching movement, the correct target was filled in. B) Simulated group data from parameters fit to our 25 participants. Figure adapted from Christopoulos, V., & Schrater, P. R. (2015). Dynamic integration of value information into a common probability currency as a theory for flexible decision making. *PLoS Computational Biology*, 11(9), e1004402. Copyright 2015 Christopoulos, Schrater. Original data from Chapman, C. S., Gallivan, J. P., Wood, D. K., Milne, J. L., Culham, J. C., & Goodale, M. A. (2010). Reaching for the unknown: multiple target encoding and real-time decision-making in a rapid reach task. *Cognition*, 116(2), 168-176. Copyright 2010 Elsevier.

Additionally, our model generalizes to the obstacle avoidance literature. Almost all

studies on reaching and decision making have focused on the competition between options of varying positive relevance or value. However, a significantly overlooked, but equally important aspect of moving and deciding is the ability to avoid choosing or moving toward something that is of negative value (e.g., disliked food, potentially dangerous territory, hot stove elements). Again using the same model and parameters for each subject, but now changing the physical layout of the options and the value of the to-be-disliked option, we see behaviour that reproduces key aspects of obstacle avoidance. Reaching movements tend to avoid the negatively valued option proportional both to its negative value, and to its distance (Chapman, Gallivan, Wong, Wispinski, & Enns, 2015; Chapman & Goodale, 2008). A singular mechanism to approach valuable options in space, while steering clear of negatively valued options is behaviourally adaptive, and neurally efficient. Further, theory and experiments have shown that obstacles take more time to be repulsive for reach trajectories than targets take to be attractive (Chapman, 2010; Chapman et al., 2015). Our model is consistent with these ideas, as potential targets start at a small positive value (Chapman et al., 2010; Christopoulos et al., 2015; Gibson, 2014), and are then either boosted further by positive evidence for that action, or take time to become negatively valued if given negative evidence.

Of note, reach trajectories in this simulation make sharper angular changes than seen in obstacle avoidance experiments (Fajen & Warren, 2003). This is likely due to equal velocity time steps in our model, whereas a real reaching movement has higher velocities near the middle of the reach, and lower velocities at the start and end. Further, sharper changes than usual may be due to the absence of motor planning mechanisms in our model, which take time to plan full trajectories before movement initiation (Kawato, 1999). In contrast, our model calculates reach trajectories online without any pre-planning. While our model fits well for tasks with few objects

and simple geometries, we expect our model to deviate more from behavioural results with more complex task requirements without the addition of a motor planning component (Stewart, Gallivan, Baugh, & Flanagan, 2014).

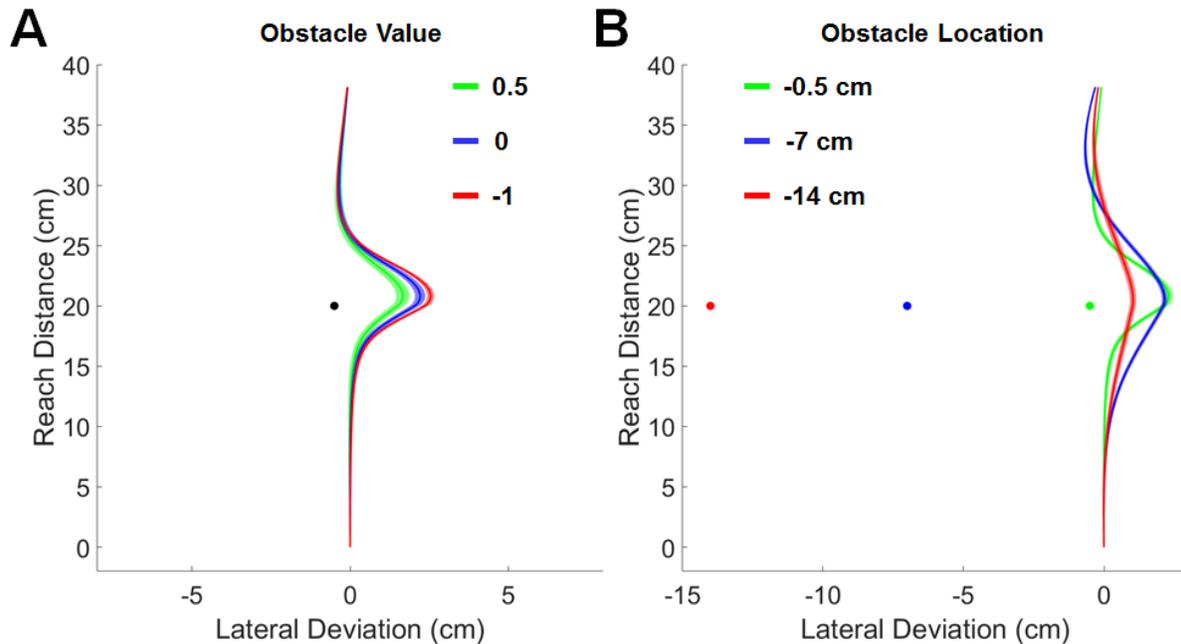


Figure 2.9: Group average reaching trajectories from model simulations. A) We place a target (valued at 1) at $[0, 40]$ and a point obstacle at $[0, 20]$. We simulate three different conditions: where the obstacle is attractive as in a medium decision (obstacle value 0.5; green trajectory), where the evidence toward reaching to the obstacle is zero (obstacle value 0; blue trajectories), and where the action for reaching to the obstacle accumulates negative evidence (obstacle value -1; red trajectories). B) We hold value constant (target value 1, obstacle value -0.25) and we instead change the physical location of the obstacle. As we move the obstacle further from the reach path, average trajectories deviate less away from the obstacle. When the obstacle is at $[-15, 20]$, we see little interference with the reach trajectory, similar to a single target trial.

2.3 - Discussion

Here we leveraged an existing model of decision making to conceptualize movements as an ongoing, graded competition between available actions. Using our new formalization, we are able to account for a diverse range of behaviour: reaction times, initial choices, changes of mind, reach curvature, and accuracy. In addition, our model is highly generalizable to tasks such as

value-based reaching, go-before-you-know tasks, and obstacle avoidance. Our results point to a new formal way of thinking about movements - not as a singular command read out by cognition (Padoa-Schioppa, 2011), nor as a noisy result of binary switching between actions (e.g., Resulaj et al., 2009), but as a graded, time-evolving competition intimately tied to ongoing decision making.

Several recent studies have proven the power of using this concept of movements as continuous decision making. By regressing multiple attributes of available options in a reaching task onto incremental changes in the angle of reach trajectories, these studies are able to infer the ongoing evolution of bias toward each of two options (Dshemuchadse, Scherbaum, & Goschke, 2013; Scherbaum, Dshemuchadse, Fischer, & Goschke, 2010; Sullivan, Hutcherson, Harris, & Rangel, 2015). This not only suggests that the hand is continuously moving with respect to fluctuating evidence, but that momentary evidence is dynamic and biased by many sources that evolve in time.

Why don't we simply wait until we have enough evidence, and then enact the most efficient, straight path to our selected target? In real world decision making, waiting has a cost. Not only does continuously integrating evidence take up neural resources, but options may pass you by while you are waiting. Additionally, the world is dynamic, and potential options and actions can change almost any time. For these reasons, it may be more efficient to begin moving once you are confident *enough*, and revise your decision if the world changes. However, the efficiency difference in switching between straight paths to options (Resulaj et al., 2009) compared to sometimes moving toward a midpoint is less clear. Certainly, if one option is the most likely to be desired, a frequent straight path toward it and a less frequent but more costly switching movement may be optimal. However, if an option is less likely to be desired all the

time, a middling trajectory to reduce frequent switching costs may be more optimal. The relative efficiency of these different mechanisms depends on a number of factors including biomechanic costs of switching, neural resources used by integrating and transmitting evidence, decision difficulty and consequences, and task parameters. While both mechanisms are not mutually exclusive, we argue that moving toward options proportional to their current desirability is the singular general mechanism that most captures human behaviour.

Like all models, ours is a simplification. The current formalization neglects movement times, motor costs, biomechanics, motor planning, and ballistic reaching movements - all which are known to influence behaviour in reaching tasks. Despite these limitations, we argue that our model, by simply fitting a single parameter after a drift diffusion model, is powerful in the range of behaviour it can explain, and in the way it conceptualizes the link between movements and decision making.

3 - Neural Signals of Decision Making

Here in Chapter 3, we present results of our own study of neural activity during decision making. Specifically, we record brain activity during a reaching decision making task, unlike most other studies which neglect having an overt and extended movement as the “output” of a decision. As outlined in Chapter 1, here we use reaching tasks as a window into naturalistic decision making, where people are faced with potential actions, and must choose between actions by completing movements. By observing neural patterns when deciding and moving, we hope to uncover the process and architecture behind a general decision making mechanism. This goal is further accomplished by using EEG, as it allows measurements of neural activity in humans because of its non-invasive nature, and is particularly adept at capturing the fast and dynamic process of decision making because of its high temporal resolution. While studying decision making in humans using EEG does not allow for high spatial resolution or signal-to-noise ratios as available in invasive monkey recordings, humans are (1) the system of specific interest, (2) able to execute controlled decisions with little training, and (3) require relatively low resources to collect data, which allows many more participants in a sample.

As mentioned in Objectives and Predictions in Chapter 1, models seeking to explain foundational human behaviour do well to constrain possible solutions with plausible neural mechanisms and observed neural patterns. While the following EEG and reaching data is a rich dataset, our analyses will focus on how neural patterns do and do not align with our model in Chapter 2, and which architectures of decision making they most support.

Our model makes three key predictions about neural signals observed during a reaching task. (1) Our model predicts an accumulation-to-bound signal (specifically a diffusion process) representing a decision to begin movement. (2) The model predicts signals observed in

sequential models of decision making (Chen & Stuphorn, 2015), where goods-based information is communicated to the action-based system with some delay. (3) Our model predicts that motor representations compete, and that this competition should be reflected in ongoing reach trajectories. While the analyses in this chapter are not yet fully complete, we will address the first two of these model predictions.

In this experiment, we ask participants to attend to two bilateral circles. Participants are asked to identify which of these circles is brighter, or which of these circles are darker (counterbalanced). Critically, this decision relies on relative luminance difference. Sometimes circles may both change in luminance, but in the same direction. Therefore, participants must sample the luminance of both circles, and compare these luminance values together. Again, we manipulate the decision difficulty using the amount of relative luminance difference. Easy decisions are called “double difference”, and have a large relative luminance difference. More difficult decisions have a lower level of relative luminance difference, and are called “single difference”. Finally, some trials have no relative luminance difference throughout, and therefore no evidence. These trials are called “no difference” trials. One set of our participants were asked to respond by keyboard, as is the norm, while the other set responded by reaching to touch the circle they thought was brighter (or darker). Unlike random dot motion tasks, participants here when responding using reaches are sampling evidence from the same physical location as they need to act toward. We argue this more accurately reflects task demands outside of the lab - sampling information at locations for potential actions.

In this experiment, we focus our EEG analyses on the centro-parietal positivity, or CPP (O’Connell et al., 2012). As outlined in the Introduction, the CPP is thought to reflect evidence accumulation in posterior parietal areas. It is modulated by decision difficulty, is independent of

the modality of decision evidence, is distinct from sensory and action signals, and is even observed when no action is required at all (O’Connell et al., 2012; Kelly & O’Connell, 2013; Murphy et al., 2015). Therefore, we look to the CPP as an index of a goods-based system in decision making, where evidence is compared distinct from motor representations. Our model in Chapter 2 predicts a drift diffusion process in a goods-based system that determines reaction times for competing motor representations in the action system. Therefore, we predict that the CPP reflects a decision to begin moving. This result has not yet been seen because other studies use verbal or keypress responses, where action initiation and final response are tightly linked in time.

Of note, studies have argued that the CPP is best observed under circumstances where stimuli change gradually. This is argued to reduce signal contamination from strong stimulus-onset signals (e.g., visually evoked potentials; VEPs). Therefore, the current task uses stimuli that slowly change in luminance over time. However, this means that evidence in the current task also changes gradually over time, and that participants therefore have relatively long reaction times. Finally, the current task is a perceptual decision task, rather than a value-based task. In other words, evidence comes from an external, rather than an internal source. Taken together, the current task is very different from the task in Chapter 2, and provides a demanding test for the flexibility of our formalization.

3.1 - Methods

3.1.1 - Experiment. 23 participants completed the keypress version of the experiment, while a different set of 19 participants took part in the reaching version of the experiment. All participants were right-handed and had normal or corrected-to-normal vision. All experiments were approved by the University of Alberta’s Research Ethics Office, and all participants

provided written consent before the experiment.

Participants sat in a dimly lit room centered, at eye level, and 50 cm away from a computer monitor with a refresh rate of 120 Hz (VIEWPiXX/EEG, VPiXX Technologies, Quebec, Canada). For participants responding using reaching movements, six wall-mounted motion tracking cameras (Optitrak V100: R2 cameras; NaturalPoint, Inc., Corvallis, Oregon) captured the participant's hand position in space at 60 Hz, synchronized to the monitor refresh rate. The position of the single reflective finger marker on the participant's right index fingernail was co-registered in space with the monitor and tabletop for each participant. Participants who responded with button presses instead had a computer keyboard (USB polling rate of 1000 Hz) placed on the table in front of them. Participants were asked to press the left and right Ctrl keys with their index fingers to respond.

Participants were directed to answer the question "Which circle is [brighter/darker]?" on every trial, counterbalanced across participants. The experimenter explained that two flashing circles would appear, and after some variable delay, the circles might change in luminance. When answering the question, participants were directed to consider the relative luminance difference between the two circles, and not if there was any absolute change in luminance for any of the circles.

For the reaching condition, participants started each trial by moving their right index finger to the start position. Participants were asked to simply touch the circle that they thought was [brighter/darker] when they had detected a relative luminance difference. For participants responding using a keyboard, participants were asked to press the key on the same side of space as the circle they identified as being relatively [brighter/darker]. All participants were notified that a third of the time there would be no relative luminance difference, and to not respond if

they did not detect one. All participants were directed to maintain fixation on a central fixation cross throughout the experiment.

Stimuli were presented on a black background with a white central fixation cross. At the start of each trial, only the fixation cross was present for a baseline epoch of 750 to 1250 ms, random and uniformly distributed. After baseline fixation, two bilateral flashing grey circles appeared 9 cm apart, and were present for five seconds. Circles were always an intermediate grey at baseline (126 in 8-bit grayscale) and flickered at ~ 17.14 Hz and 20 Hz with the stimuli being present on the screen for 1 frame per cycle ($8 \frac{1}{3}$ ms). Changes in luminance occurred between 750 and 2250 ms after the onset of both circles, randomly drawn from a uniform distribution. Regardless of condition, luminance changes always occurred at the same linear rate of 100 8-bit grayscale points over 1750 ms, and 100 points over 1000 ms, for changes and return to baseline, respectively. Once the circles had returned to baseline luminance, they remained at baseline for any remainder of the five second stimuli time. Stimuli presentation were controlled with Matlab using Psychtoolbox (Version 3, Brainard, 1997).

Participants completed trials of three critical decision types: (1) no relative luminance difference, (2) “single difference” (i.e., maximum luminance difference of 100 8-bit grayscale points), and (3) “double difference” (i.e., maximum luminance difference of 200 8-bit grayscale points; see Figure 3.1 B). For trials where the two circles were never of a different luminance, both circles remained at baseline, both increased in luminance, or both decreased in luminance. For single difference trials, one circle increased in luminance while the other remained unchanged, or one decreased while the other remained unchanged. Finally, for double difference trials, one circle increased in luminance while the other decreased. Conditions were counterbalanced for side of space (i.e., which change occurred on the left vs. right side of the

screen), and all conditions were counterbalanced for flicker frequency (~17 vs. 20 Hz). This resulted in 18 unique conditions, each with different physical stimuli pattern. All unique conditions were repeated 25 times each (for a total of 450 trials), randomized by block so that no unique condition could appear again before all others had been presented.

The computer monitor was fixed to be 35 cm forward from the start circle, and participants were seated so that the screen was 50 cm from their eyes. Circles were centered on the screen with 8 cm between them. For reaching participants, movement onset was determined using standard velocity and distance criteria (Gallivan & Chapman, 2014). Further, if participants had moved > 1 cm from the start circle before stimuli onset, the trial would be flagged as “Too Early”.

During EEG preparation, participants completed the UPPS-P impulsive behaviour scale (Lynam, 2013) and the gambling related cognitions scale (GRCS; Raylu & Oei, 2004). These data were not analyzed and are not discussed further.

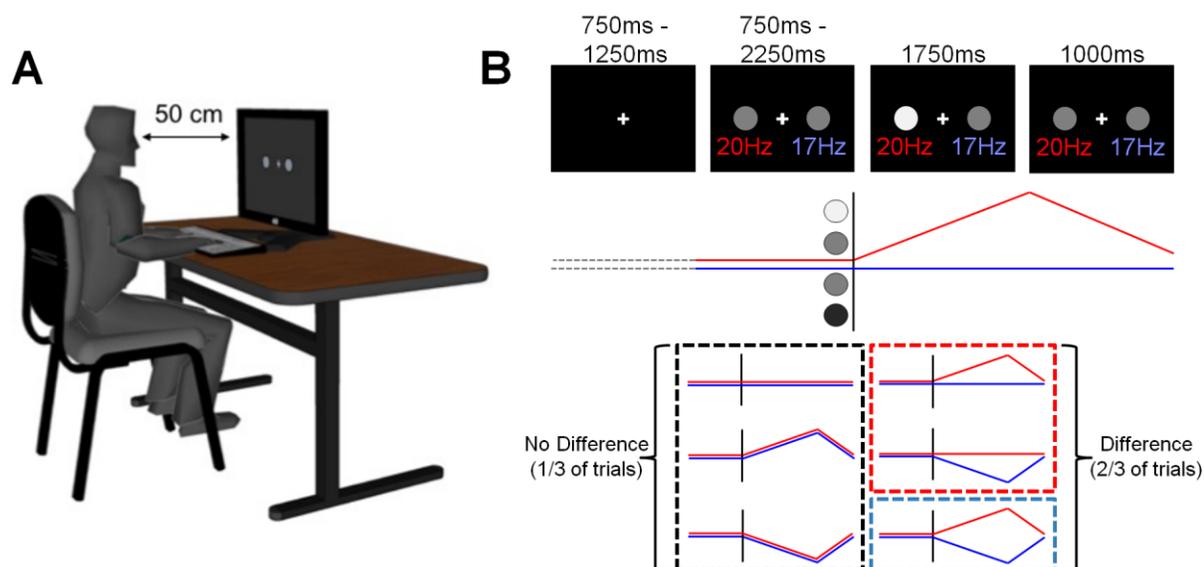


Figure 3.1: A) Experiment setup. Participants were seated 50 cm from a monitor at eye level and responded either using a keyboard, or by reaching. B) Diagram of experimental stimuli. At the start of each trial, participants saw 750-1250 ms of a fixation cross to establish a baseline for

EEG recordings. Then participants were presented with two grey circles for a variable period of time before they started to change. Changes occurred for 1750 ms, and returned from peak change back to baseline by a further 1000 ms. Three critical conditions were presented, no difference (black), single difference (red), or double difference (blue). On each trial, one circle flickered at ~17 Hz, and the other at 20 Hz. On this example trial in the diagram, the left circle flickering at 20 Hz increased in luminance while the right circle at ~17 Hz remains at baseline luminance.

3.1.2 - EEG collection and preprocessing. EEG was recorded using a wet 32-channel Ag/AgCl active electrode system (EasyCap, Herrsching, Germany), with 30 scalp channels (10/20 system) and two channels (TF9/TF0) modified for mastoid recording. Recordings were referenced online to the right mastoid, and offline re-referenced to the average of both mastoid electrodes. The ground electrode was placed at the frontal midline electrode site (Fpz). Before recording, impedance of all electrodes was reduced to below 10 k Ω . Bipolar vertical and horizontal EOG was recorded using five sintered ring electrodes (EasyCap, Herrsching, Germany), two on the left and right temples, two above and below the left eye, and one ground electrode between eyebrows. EEG was recorded at 1000 Hz using BrainVision Recorder (Brain Products, München, Germany) and amplified using an actiCHamp amplifier (Brain Products). Data were filtered online with a low-pass filter at 280 Hz. No high-pass filter was used online. Pixel-associated triggers from the VIEWPixx/EEG were used to mark the EEG data.

EEG processing was conducted using EEGLAB (Delorme & Makeig, 2004) with the ERPLAB (Lopez-Calderon & Luck, 2014) and CSD (Kayser & Tenke, 2006a, 2006b) toolboxes. Data were high-pass filtered at 0.1 Hz, low-pass filtered at 30 Hz, and baseline corrected to the 750 ms before stimuli onset. Noisy channels (where >10% of epochs were rejected because of channel noise) were deleted and were not interpolated because of subsequent ICA and CSD analyses. Epochs were manually rejected before computing and removing stereotypical ICA

components associated with common artifacts such as blinks, eye movements, or muscle clenching. After ICA rejection, epochs were manually rejected again using the same methods. Epochs were created to include 750 ms of fixation baseline and all five seconds of stimuli presentation on each trial (epoch length 5750 ms). Finally, a surface Laplacian was applied to the remaining data using the CSD toolbox in Matlab (Kayser & Tenke, 2006a; Kayser & Tenke, 2006b) in order to remove low spatial frequencies from the data. Applying a surface Laplacian before examining the CPP has been performed in past studies because of the strong anatomical predictions of the source of CPP (and decision making) activity, and to attenuate low spatial frequency activity such as movement artifacts (Kelly & O'Connell, 2013).

Participants were excluded from analysis if fewer than 50% of all trials remained after all EEG epoch rejection, or if fewer than 40 trials remained in any of the three critical conditions. This left 19 (17% rejected) participants in the keyboard response version, and 16 (16% rejected) in the reaching response version. Epoch and participant rejection left an average of 349 and 307 epochs remaining for analysis for the keyboard and reaching experiments, respectively. No channels were rejected for the keyboard participants, while an average of 0.84 channels were rejected per reaching participant. Incorrect or too early trials were removed from EEG analysis, except in the generation of the raster plots in Figures 3.6 and 3.9. For response- or movement onset-locked analyses, trials from the no difference condition were locked to the grand mean reaction time within that participant, as participants were asked not to respond on those trials.

3.1.3 - Model. As stated above, it is useful to directly apply our model in Chapter 2 to neural data to investigate if actual neural traces match those predicted by our model. While analyzing the similarity between model and data for behavioural measures in the current task can also lend or detract support for our model, the comparison between EEG data and model traces is

the main focus of Chapter 3. To date, only participants responding by keyboard have had a model applied to their data. Without a reaching component to their data, only a drift diffusion model was used. Nonetheless, we can test most of the predictions of our model architecture with a drift diffusion model determining reaction times.

Evidence in support of a decision (C) was modelled as the luminance of the option at that location (as a difference from baseline, 0, and ranging from -1 to 1). Critically, evidence in the current task changes over time as the luminance of each option changes gradually throughout 2750 ms. Changing evidence over time simply means a shifting C value, and so the means of noisy momentary evidence shift as well. This increase in evidence over time leads to an increase in the rate of evidence accumulation, which results in exponential rather than linear accumulation rates.

Fits were obtained by minimizing the negative log likelihood of responses (left, right, none), and mean reaction times, for each of the three conditions (double difference, single difference, no difference), for each side of space (left preferred, right preferred), for each person. Because participants were asked not to respond on no difference trials, reaction times for this condition were not used to fit the model. Fits of the drift diffusion model parameters were obtained by simulating 10,000 trials at 120 Hz (the screen refresh rate, and twice the motion capture sampling rate) for each condition, for each side, for each parameter iteration. Like in the experiment, each trial simulated five seconds of time where the stimuli were present on the screen. Changes in evidence began to occur at a random time between 750 to 2250 ms after stimuli onset. If the difference in accumulated evidence crossed either $+B$ or $-B$ at any time within these 5 seconds, a simulated keypress was made. Otherwise, it was assumed the simulated participant did not respond on those trials. Drift diffusion traces were brought back to baseline

after a response to simulate the cessation of the process.

To compare drift diffusion traces to the CPP data, the no difference conditions of both grand mean EEG and model traces were aligned at the first time point. As well, the model traces were scaled to the CPP by aligning the maximum value of the double difference conditions. This was simply done so that both model and EEG data were of the same absolute range to calculate R^2 , and did not change the pattern of any results. Further, absolute EEG amplitude can vary due to many factors such as scalp conductivity, skull shape, and individual variability in brain structure, and so the absolute values for this comparison are not of interest - only the relative patterns.

3.2 - Results

3.2.1 - Behavioural results. Repeated measures ANOVA returns a significant main effect of condition ($F(2,66) = 4.37, p = .017$) and response type ($F(1,33) = 9.42, p = .004$) on accuracy, with no significant interaction ($F(2,66) = 2.38, p = .10$). Multiple comparisons show that participants are more accurate when responding with a keyboard relative to reaching ($t(20) = 2.89, p = .009$). As expected, participants are more accurate in the double change condition relative to the single change condition ($p = .0001$), with accuracy in the no change condition being not significantly different from either.

For reaction time, repeated measures ANOVA returns a significant main effect of condition ($F(1,36) = 318.3, p < 2e-16$) and no main effect of response type ($F(1,36) = 1.27, p = .27$), with a significant condition-by-response type interaction ($F(1,36) = 11.30, p = .002$). This interaction seemed to be primarily driven by a marginal difference in single change RT means between response types ($p = .057$), with faster RTs for single change reaching responses relative to single change keyboard responses. Multiple comparisons show that, as expected, participants

are faster at initiating actions in the double change condition relative to the single change condition ($p = 2.2e-16$). For clarity, reaction times were not analyzed in the no change conditions, because participants were asked not to respond in this condition (and correctly did not respond on average 85% of the time).

Overall, participants are more accurate and faster at initiating movements when presented with more evidence for a decision. Of note, relatively long RTs are not unexpected in this task given the slow rate of change of the stimuli, and high accuracy for all conditions suggests participants find this task easy (as intended). Further, participants are initiating their action, be it keypress or reaching movement, before the peak of luminance difference (average RTs ~ 1125 ms, relative to an evidence peak at 1750 ms). This suggests that participants feel they have enough evidence to make a decision earlier on, and that they are not using a peak in evidence as a cue for movement.

Analysis of average correct reach trajectories for conditions where participants were asked to make a response (single and double difference) were analyzed using methods outlined in Gallivan & Chapman (2014). Of note, 18 participants had reach data sufficient for this specific analysis (some of which were not included in EEG analyses). The absolute lateral deviation at each normalized time point was used, and uncorrected t-tests of a difference between single and double difference trajectories against a mean of zero were computed for each of 100 normalized time points. In this sample, there was no significant difference between reach trajectories in conditions with different levels of evidence, despite a descriptive difference between conditions in the expected pattern (see Figure 3.2). However, this result may be due to the relatively small sample size compared with reaching studies (18, relative to a typical 30), and the small distance between targets compared to other reaching studies, including in Chapter 2 (8 cm apart, relative

to 30 cm apart in Chapter 2).

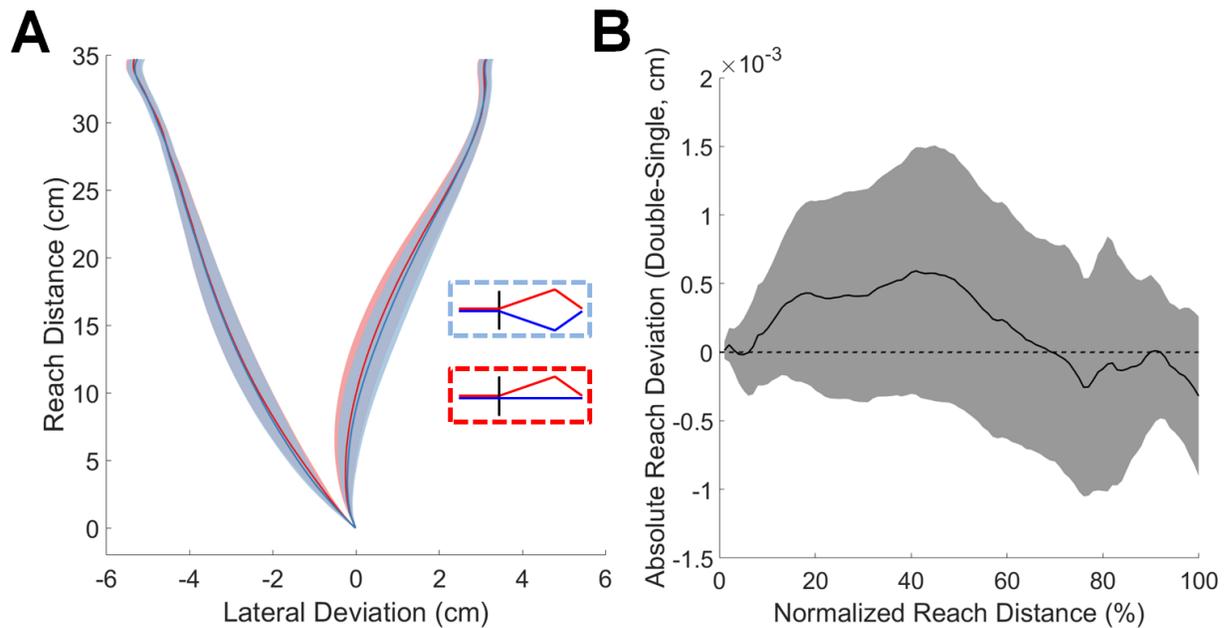


Figure 3.2: A) Average correct reach trajectories for double difference (blue) and single difference (red) conditions. B) Difference plot over 100 normalized reach distance points between double and single difference conditions. At no point was the difference between conditions significantly greater than the null hypothesis of 0. All error bars are S.E.M.

3.2.2 - EEG results. Analysis of the CPP (O’Connell et al., 2012) locked to RTs show patterns predicted by general evidence accumulation models (see Figure 3.3). Specifically, midline parietal voltage increases for several hundred milliseconds up until movement is initiated, and then returns to baseline levels. Importantly, this increase only occurs for conditions where there is decision evidence present. Uncorrected t-tests at every 1000 Hz time point show that voltages in the single difference condition are significantly greater than in the no difference condition for 67% of the 500 ms window before response in the keyboard group, and for 84% before movement initiation in the reaching group. Further, for the 500 ms preceding action initiation, voltages in the double difference condition are greater than those in the single difference condition for 22% and 5% of the time window for keypress and reach responses,

respectively. This suggests that the CPP is greater for conditions where there is more evidence. Interestingly, we also find a significant decrease in CPP voltage ~ 700 ms post-movement initiation for conditions where participants made a reach (single difference vs. no difference significant for 49% of the time window 500 to 1000 ms after movement initiation). This decrease was not present in any keypress conditions. Additionally, this decrease in the reaching group was sensitive to evidence, with the double difference condition more negative than the single condition for 18% of the same 500 to 1000 ms post-movement initiation window (compared to 0% in the keyboard group). While this effect may be solely movement related, it occurs at the later end of reaching movements (mean movement time = 822 ms). Further analyses are needed to determine if this decrease is decision-related.

While the pattern of the CPP locked to movement initiation time is consistent with general evidence accumulation models, it suggests that evidence accumulation stops at movement onset. This result is consistent with our model proposed in Chapter 2, where a drift diffusion process determines reaction times, and then stops, while individual accumulators and motor representations continue until action completion. This result is also consistent with purely goods-based models of decision making where a decision is made before a singular motor command is executed. Further, this result is inconsistent with action-based models of decision making, where competition occurs at the level of motor representations until action completion. While Resulaj et al. (2009) also posit a diffusion process that begins movement initiation, the diffusion process in their model continues past movement initiation to determine changes of mind. The pattern of the CPP here does not support the model of Resulaj et al. (2009), as the CPP does not continue past movement initiation (at least visibly, and in the same pattern as pre-movement initiation).

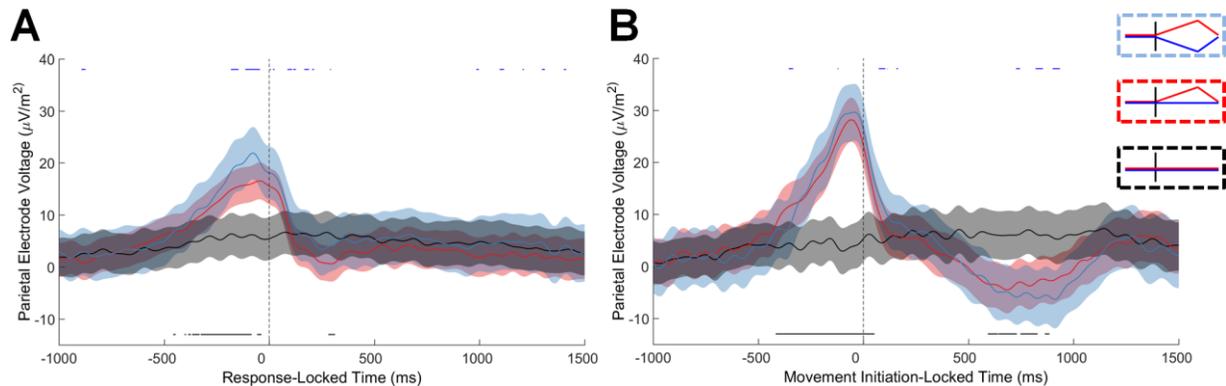


Figure 3.3: RT-locked event-related potentials (ERPs) from midline parietal electrode. Blue represents double difference conditions (e.g., easy), red represents single difference conditions (e.g., hard), and black represents the conditions with no relative luminance difference (locked to the grand mean response time for each subject). Only correct trials were included. Data were baselined to the 500 ms window just before stimuli began to change. Error bars represent S.E.M. Grand mean ERPs were filtered at 10 Hz for presentation purposes only. Black bars on the bottom of each figure show where a t-test between the no difference (black) and single difference (red) conditions reached conventional significance ($p < .05$), and blue bars across the top represent significant t-tests comparing the double difference (blue) and single difference (red) conditions. A) Keypress response group. B) Reach response group. Conditions where decision evidence (i.e., a relative luminance difference) was present show an increase in voltage for several hundred milliseconds until action initiation occurs, consistent with evidence accumulation models.

Analysis of the CPP locked to stimuli change times again show a pattern consistent with evidence accumulation models, with slower rates of accumulation in conditions with less evidence (see Figure 3.4). As with response-locked ERPs, differences waves were analyzed using uncorrected t-tests. Here, we focus on the 1000 ms time window starting 500 ms after stimuli began any luminance change to capture the expected period of evidence accumulation after accounting for neural delay and slowly changing stimuli. For keyboard responses, we see a difference between the no difference condition and the single difference condition for 9% of the time window. We also see a significant difference between the single difference and double difference conditions for 27% of this time window. For keyboard responses, all significant

differences are in the expected direction. For reach responses, we see a difference between the no difference condition and the single difference condition for 80% of the time window. We also see a significant difference between the single difference and double difference conditions for 80% of this time window. However, not all significant differences are in the expected direction. Again, we see a post-movement initiation dip in the CPP during and after reaching movements, in the opposite pattern than pre-movement initiation. While this suggests that the CPP may reflect some decision-related information during reaching movements, further evidence is needed. Taken together, these results again suggest that the CPP scales with decision evidence in the pattern expected by evidence accumulation models before an action, but breaks down during actual movements.

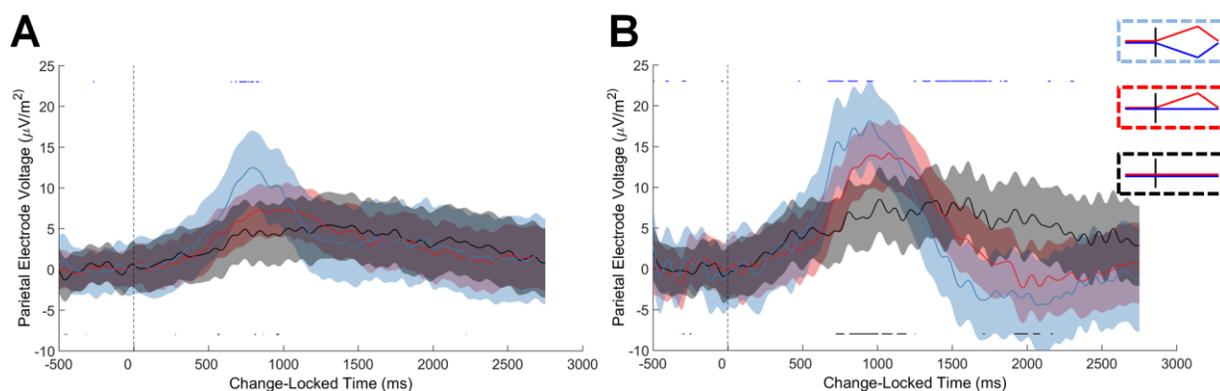


Figure 3.4: Luminance change-locked event-related potentials (ERPs) from midline parietal electrode. Blue represents double difference conditions (e.g., easy), red represents single difference conditions (e.g., hard), and black represents the conditions with no relative luminance difference. Data were baselined to the 500 ms window just before stimuli began to change. Error bars represent S.E.M. Grand mean ERPs were filtered at 10 Hz for presentation purposes only. Black bars on the bottom of each figure show where a t-test between the no difference (black) and single difference (red) conditions reached conventional significance ($p < .05$), and blue bars across the top represent significant t-tests comparing the double difference (blue) and single difference (red) conditions. A) Keypress response group. B) Reach response group. Consistent with past findings, the CPP rises proportional to the amount of presented decision evidence.

Past studies of the CPP have shown that it precedes the lateralized readiness potential (LRP) by >100 ms (Kelly & O'Connell, 2013). The LRP has long been associated with motor preparedness, and is measured by taking the difference of contra- and ipsilateral electrodes over motor cortex (Mordkoff & Gianaros, 2000). Here, we calculate the LRP for both left and right correct responses for all conditions locked to response time (no difference conditions were again locked to the grand mean RT within participants; see Figure 3.5). We analyzed the time at which both single and double difference conditions were significantly different from the no difference condition for a consecutive five time points. When looking at the CPP, we find for keyboard response participants that the difference conditions first become statistically different from the no difference condition 457 ms before response time. Analysis of the LRP showed this significant difference first occurred 171 ms before response time. Together, the difference between the CPP and LRP (286 ms) suggests that the motor system is downstream of evidence accumulation. Reach responses, which we are all made with the right hand, were not predicted to, and did not show a reliable LRP pattern, and were therefore not included in this analysis.

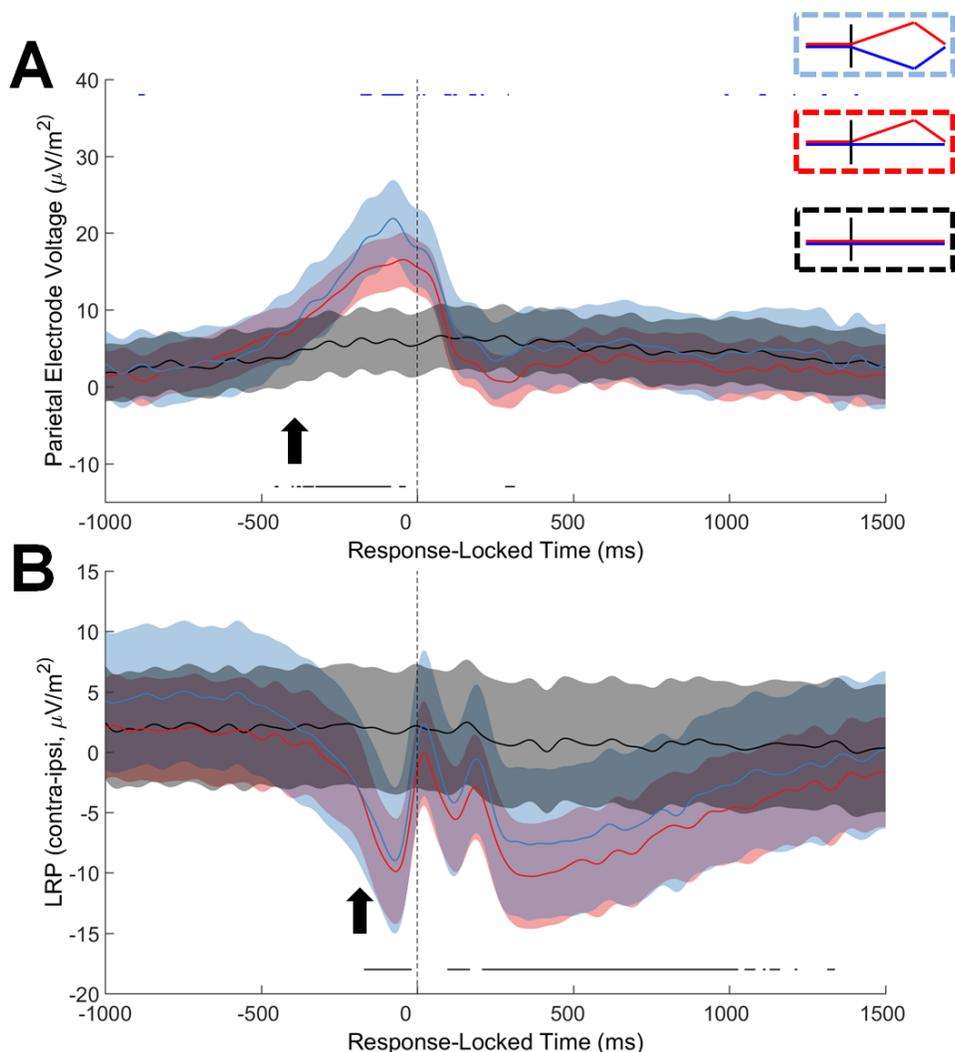


Figure 3.5: RT-locked event-related potentials (ERPs) for keypress participants. A) CPP. B) Lateralized readiness potential (LRP), which is argued to index motor preparedness. Blue represents double difference conditions (e.g., easy), red represents single difference conditions (e.g., hard), and black represents the conditions with no relative luminance difference (locked to the grand mean response time for each subject). Only correct trials were included. Data were baselined to the 500 ms window just before stimuli began to change. Error bars represent S.E.M. Grand mean ERPs were filtered at 10 Hz for presentation purposes only. Black bars on the bottom of each figure show where a t-test between the no difference (black) and single difference (red) conditions reached conventional significance ($p < .05$), and blue bars across the top represent significant t-tests comparing the double difference (blue) and single difference (red) conditions. Black arrows denote the time where both single and double difference conditions were significantly different from the no difference condition for 5 consecutive time points. Here, the difference (red and blue) conditions diverge from the no difference (black) condition faster for the CPP than the LRP, suggesting that information flows from one (goods-based system) to the other (action-based system).

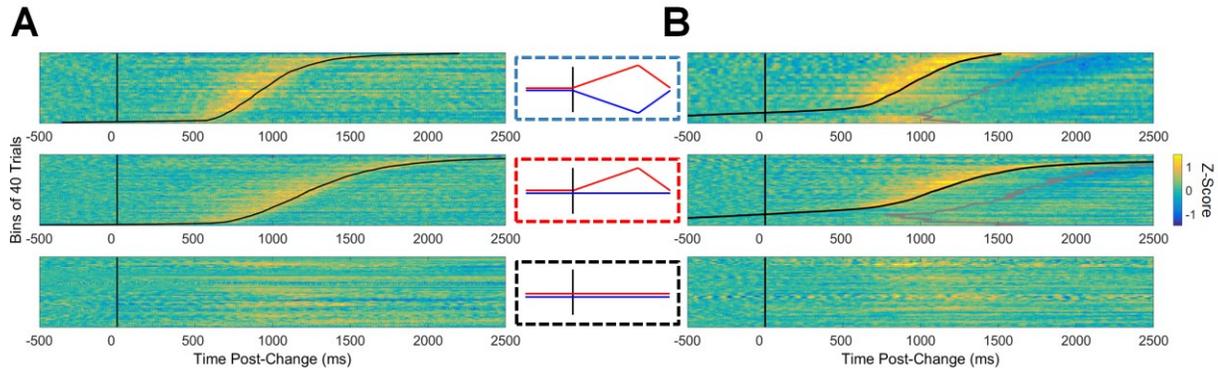


Figure 3.6: Raster plot of CPP voltage locked to the time when stimuli began to change in luminance. Trials are ordered by RT, binned into groups of 40 trials, and averaged to reduce noise. All trials are baselined to the 500 ms window just before stimuli change. Trials are z-scored within each individual and collapsed together. Raster plots for double difference, single difference, and no difference conditions are from top to bottom, respectively. A) Keyboard responses. Black curved line represents average RT for each of the 40 trial bins. CPP voltage increases until RT before returning to baseline voltage levels. B) Reaching responses. Black line represents average movement initiation time, while the grey line represents the average time when participants touched the screen to indicate their choice. CPP voltage is locked to movement onset times, and not to movement completion times.

In sum, parietal signals in our task are consistent with a decision to initiate a movement, and do not appear to continue in time. These results are consistent with our model in Chapter 2, but are inconsistent with a changes of mind model (Resulaj et al., 2009), as discussed in more detail below. Additionally, we extend the findings of O’Connell et al. (2013) to a different task with two spatially separated stimuli and find the same pattern of CPP signal that falls off with a keypress to indicate a decision.

3.2.3 - Model results. Here we directly compare data generated by model simulations to participant behaviour and EEG data in the keyboard task. Again, all R^2 values are reported first as the average R^2 of individuals, and second as the R^2 of group mean data. Only group mean R^2 are reported for EEG data because of EEG variability between individuals.

The drift diffusion model is able to accurately account for reaction times ($R^2 = 0.55; 0.90$)

but not choices ($R^2 = 0.06$; 0.00 ; see Figure 3.7). While the similarity between reaction times is reassuring for the use of a drift diffusion model, the lack of a drift diffusion model to account for correctly not responding on no difference trials is of interest and discussed further below.

Further, while the drift diffusion model was fit to these data, and has been shown to accurately account for these aspects of behaviour in other studies, the similarity between model and data (at least for reaction times) is reassuring as this is a novel task with evidence that changes over a long period of time, unlike most other tasks.

As stated above, we assume that the CPP is a direct reflection of a drift diffusion process actually implemented in the brain that determines movement onset in the current task. From this assumption, we predict that the CPP should follow the pattern of drift diffusion accumulation-to-bound. When organizing drift diffusion traces by reaction time from model simulations as is performed for the EEG data, we find that a drift diffusion model can account for 40% of the variance in the response-locked CPP ($R^2 = 0.40$; see Figure 3.8 A), but cannot account for variance in the CPP when locked to the onset of luminance changes ($R^2 = 0.00$; see Figure 3.8 B). While visually these traces are similar, especially in the raster plot (Figure 3.9), there are several dissimilarities. First, while the model predicts a peak at response, and a sharp drop off indicating the end of the process, the CPP seems to peak slightly earlier and takes more time to return to baseline. Further, in the stimuli change-locked traces, the model predicts slower accumulation than is suggested by the CPP, perhaps pointing to a mismatch between predicted and actual non-decision times. This is further supported by the notably long non-decision times (mean $t_{nd} = 352$ ms) that are output from model fitting. Finally, a difference between the peaks of single and double difference traces in the CPP may suggest an additional mechanism - specifically a collapsing bound. This is discussed further in the General Discussion.

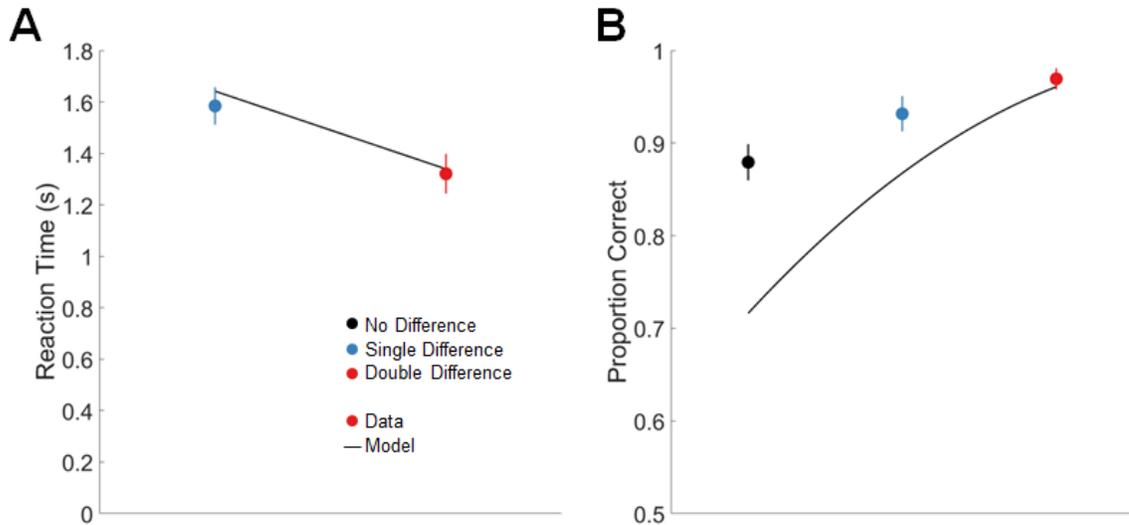


Figure 3.7: A) Reaction times for keyboard responses in our EEG task (filled circles), along with model RTs from a drift diffusion model fit to these data (black lines). B) Accuracy for keyboard responses, and predicted accuracy from model simulations.

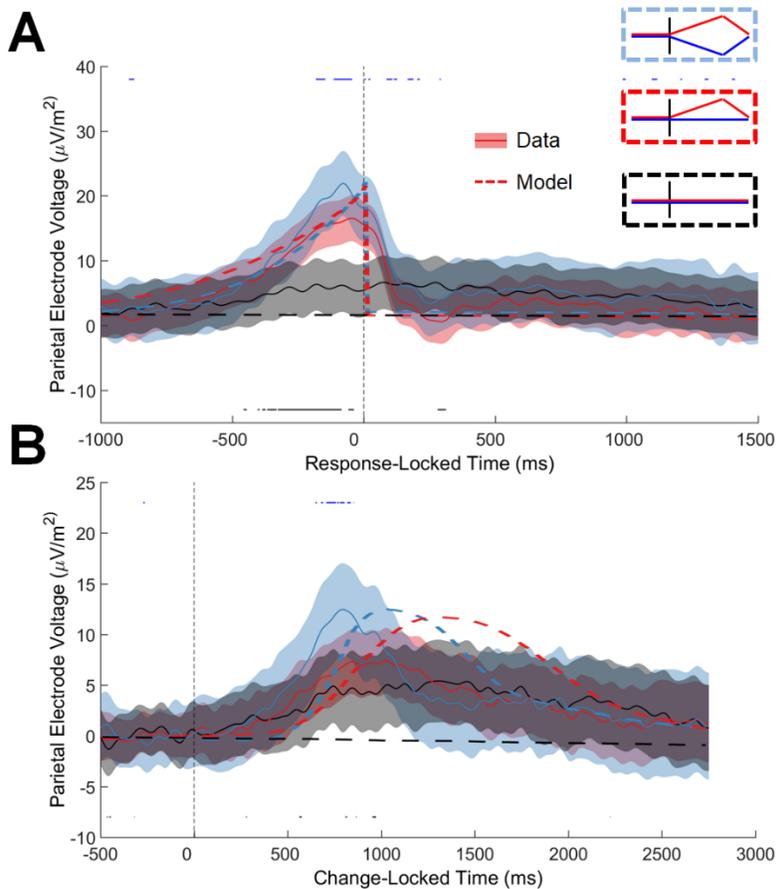


Figure 3.8: CPP data from keyboard participants (solid lines with S.E.M.) and drift diffusion traces from model simulations (dotted lines). A) Response-locked data ($R^2 = 0.40$). B) Luminance change-locked data ($R^2 = 0.00$).

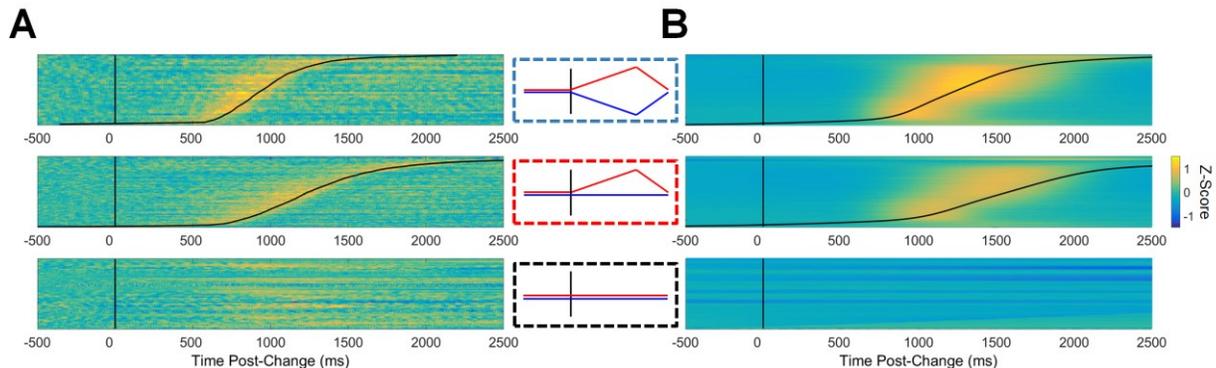


Figure 3.9: Raster plot of A) CPP from keyboard response group, and B) drift diffusion traces from model simulations. Trials were sorted by RT within condition, and averaged in order to reduce noise (bins of 40 for actual EEG data, and bins of 3000 for simulated data). All trials are baselined to the 500 ms window just before stimuli change. Trials are z-scored within each individual and collapsed together. Raster plots for double difference, single difference, and no difference conditions are from top to bottom, respectively. Black curved line represents average RT for each of the trial bins.

3.3 - Discussion

Here, we aimed to probe the architecture and process of decision making in the human brain using neural data from EEG. Specifically, we recorded neural data during a reaching decision making task in order to investigate the dynamics of goods- and action-based systems in decision making. We find that the CPP, a signal thought to reflect evidence accumulation in a goods-based system, rises to a bound, but quickly returns to baseline at movement onset and remains low throughout the movement. This result helps us distinguish between different candidate architectures of decision making. Specifically, an accumulator stopping at the onset of action supports a pure goods-based model of decision making where options are selected before a movement is planned and executed. This result does not support an action-based architecture of decision making as competition should occur before *and* during movement at the level of motor representations.

The finding that the CPP ceases at movement onset additionally helps distinguish

between formal models. Our model proposed in Chapter 2 predicts a drift diffusion process that determines reaction times before ceasing to be of use. Neural patterns in Chapter 3 support this prediction. Further, the changes of mind model (Resulaj et al., 2009) posits a diffusion process that determines reaction times, but predicts that this same process continues throughout most of the movement. This prediction is not borne out in the neural data - at least as observable in the CPP.

Our formal model, and the sequential architecture of decision making (Chen & Stuphorn, 2015) are further supported by these data. We find that the CPP, a neural signal argued to reflect goods-based competition, precedes similar patterns in the LRP, a neural signal argued to reflect motor preparation. This result generally replicates a similar result in a past CPP study (Kelly & O'Connell, 2013), but under different circumstances. In short, these results support a neural delay between goods-based information reaching the action-based system during decision making. While our estimates of this delay were much larger than that of Chen & Stuphorn (2015) and Kelly & O'Connell (2013), all three studies show a similar delay using different techniques and tasks.

Finally, we directly support our formal model by applying it to these data. We see that our model can accurately account for reaction times in this experiment, even though the task is very different from that of Chapter 2, and decision making tasks in general. However, we cannot account for accuracy in this task. Specifically, our model has an issue accounting for when people correctly decide not to move on no difference trials. Our model when applied to this task is essentially a three-alternative forced choice task, where reaching a positive or negative drift diffusion threshold determines left and right choices, but failing to reach either threshold during the trial time indicates a no movement choice. This implementation is simple, and with precedent

(Murphy et al., 2015), but comes with issues. Specifically, this implementation assumes that people accumulate noisy evidence for the whole trial period. Taken to its extreme, if people observed two unchanging stimuli for an infinite period of time, noisy evidence would eventually reach a bound and cause a decision. However, the human brain likely conserves energy as sampling information comes with metabolic cost (Drugowitsch et al., 2012). Further, people likely learn that changes only occur during a range of times in the task, and so may cease to consider actions if this time has passed. Together, this suggests that there is perhaps a different mechanism to not move, or to cease consideration of options, that we have not explored.

Finally, our formal model is supported by the degree of similarity between simulated drift diffusion traces, and actual CPP patterns. Here we assume the CPP reflects a drift diffusion process to determine reaction times, and our model simulation supports this idea. However, there are still open questions. Specifically, in our data and in others' (Kelly & O'Connell, 2013), the CPP only rises to a threshold to determine action, but does not decrease to a lower threshold. A schematic of the drift diffusion (see Chapter 1) shows that both a positive and negative threshold can determine action initiation. In our simulation, we simply flip all traces to rise to a threshold. However, it is unclear how this may happen neurally. This difference raises many interesting questions about neural circuits, and electrical interactions within the brain that propagate to be measured using EEG. However, these questions are unfortunately outside the scope of this thesis.

Another issue arises when considering a past CPP result. One study found the CPP is still present when no action is required (O'Connell et al., 2012). Specifically, participants were asked to only count the number of changes they detected, and this count had to be reported after blocks of trials. In this case, participants may still be making small actions, which were not measured in the experiment, to count their decision within a block. Further, a decision to begin moving may

still be computed, but an actual action may be suppressed downstream due to participants having no physical goals for movement. However, these explanations remain unsupported and require more study. Additionally, this issue raises an interesting question about whether one makes a decision to *not* move. This is, again, a fascinating question but beyond the scope of this thesis.

Overall, by measuring neural responses during a reaching decision making task, we are able to distinguish between models of decision making, and support our formal model proposed in Chapter 2. Our formal model makes three predictions: (1) The CPP reflects a diffusion process which determines movement initiation times, (2) goods-based information is communicated to the action-based system with some delay, and (3) competition occurs at motor representations and this competition should be reflected in ongoing reach trajectories. We find support for the first two predictions, while the third prediction has yet to be tested. We find a parietal signal that reflects theorized mechanisms which initiate movements during decision making. This signal aligns well with the specific predictions made by our model. Further, we find an expected delay between theorized goods- and action-based systems in the brain, again supporting our formal model and a general sequential model of decision making. While our neural data here cannot prove or disprove any model of decision making, it suggests our formalization in Chapter 2 may be plausibly implemented in the human brain.

4 - General Discussion

Decision making has been, and continues to be conceptualized as a relatively simple, serial process. Like a thunderdome, two response options enter, but only one may leave. This understanding of decision making has proven to be quite useful in psychology labs, where two response options are available to a person, information is presented, and then a button is pressed. However, human behaviour is much more rich and complex. We move around our world, and are presented with potential actions with every step - like a grocery aisle, where every change in our physical position changes the number of objects available to grasp. In addition, not only are we changing in our environments, but our environments are also changing around us - a quarterback may see a sudden opening in the line of defense, or a car may cut us off in traffic. In this way, movement is of special importance to decision making. We must be able to evaluate options in physical space, and alter movements quickly should our circumstances change. This process is so important for our biological success, that some even argue brains evolved specifically for this process (Llinás, 2002).

Several architectures of decision making have attempted to capture the interaction between deciding and moving. Goods-based architectures argue that one option is chosen and then communicated to the motor system. Formal models within this architecture argue that deciding can continue during movement as well, and tell the motor system to switch actions if circumstances change (Resulaj et al., 2009). Other architectures posit that decision making occurs at the level of movement plans, and so movements can reflect ongoing competition between targets (Cisek, 2007). These two architectures of decision making are extremely influential, and deservedly so. However, each suffers the opposite shortcoming. Goods-based models provide a convincing solution to what the brain does between the start of a decision and

the start of a movement. However, they cannot explain important aspects of human behaviour while moving. Conversely, action-based models provide a strong argument for how decision making occurs during movement, but fall short when explaining what came before these actions. In this way, the field is split between two halves of decision making.

In this thesis, I attempted to bridge models of decision making which focus largely on only one half of decision making (pre-movement initiation, Resulaj et al., 2009; post-movement initiation, Christopoulos et al., 2015). We put forth our own formal computational model of how decision making and movement are linked in Chapter 2, and supported this model with neural data in Chapter 3. Ultimately, the goal of this thesis was to use a computational model to redefine decision making as a single, continuous process in which deciding and moving are intimately linked.

Specifically, we formalized decision making as two parallel competitive processes occurring in both the goods-based and action-based systems. In the goods-based system, a drift diffusion process determines when people are confident enough to begin a movement. In the action-based system, motor representations compete to move us toward and away from options proportional to their ongoing desirability. Importantly, both processes are derived from the same evidence for a decision, which is sequentially sampled either from the environment (Chapter 3), or from memory (Chapter 2). Our model is one formalization within a sequential architecture of decision making, as proposed by Chen & Stuphorn (2015).

In a task where people are simply asked to touch the snack food they most prefer, we can account for reaction times, initial choices, changes of mind, reach curvature, and accuracy. In a second task, where people are asked to determine which of two circles appears brighter, we are able to account for reaction times, and neural responses at measured at parietal and motor sites.

Beyond its application to reaching tasks in perceptual and value-based decision making, our model generalizes to explain behaviour in go-before-you-know tasks and obstacle avoidance. While our formalization is certainly not how the brain actually works, the amount of behaviour it can explain despite having only implemented a single parameter after a drift diffusion model reinforces the power of this approach.

Of course, our model is not without limitations. Most notably, our formalization overestimates the degree of competition during decision making. We estimate more changes of mind on hard trials than is observed, systematically underestimate accuracy in both tasks, and slightly overestimate the curvature of average reach trajectories. This systematic overestimation of competition in our model may be because we have not implemented the best mechanism for the competition of motor representations. However, our overestimation may also be due to the exclusion of an additional mechanism, which would complicate our model. Further testing is needed to investigate this gap between our predictions and behaviour.

Another additional mechanism which may improve the performance of our model is the ability to cease to decide on an action. In our EEG task, we cannot account for people's accuracy to correctly not move on no difference trials. This is because noise within our five second time window reaches a movement initiation threshold much more often than it should for this condition. Including leaky accumulators, or a cut off for the decision process may help to bridge this gap.

One aspect of the EEG data briefly mentioned was the difference in CPP amplitude at response between single and double difference conditions. If the CPP truly represents accumulation to a fixed bound, then both of these accumulators should peak at the same amplitude. This difference may point to the presence of a collapsing bound, which would explain

why single difference trials, which tend to have longer RTs, have a lower CPP peak at response time. Another CPP study has found small but significant differences in CPP peak amplitude at response as well (O'Connell et al., 2012). A recent analysis of several datasets applied to several kinds of decision bound found that fixed thresholds better explain the majority of decision making data (Hawkins et al., 2015). One exception is in primate neurophysiology studies, where a collapsing bound tends to be a better fit to the data. In this study, the authors propose that perhaps decision bounds can change shape based on task, and that tasks where participants have an extreme amount of practice may lend themselves to collapsing bounds. However, the addition of a collapsing bound runs counter to the point above about underestimating the proportion of trials where participants correctly do not move in the EEG task. Adding a bound that collapses over time would result in even more false starts than already estimated in the model. Altogether, further research in this area is needed.

What if there are more than two options to choose between? Is the drift diffusion process just a simplification of a more general process with n available options? This question is arguably understudied, although some solutions do exist. Some propose a three-dimensional drift diffusion model for three option choices (Krajbich & Rangel, 2011). Others propose a threshold for a kind of ratio of evidence between an option and all of its alternatives, or simply reducing decision making to sequential choices between two options at a time (Bogacz et al., 2006). However, the difficulty of systematically collecting enough data where there are several potential options remains a barrier to answering this important question. Models focused on action may have a more convincing solution to this problem (e.g., our model, or that of Christopoulos et al., 2015), as several motor plans or representations can theoretically compete at once. However, this still leaves open a question about what mechanism determines movement initiation in these

cases.

Another outstanding question is whether the brain actually accumulates evidence at all. Several studies have shown signals matching evidence accumulation in the firing rates of monkey LIP neurons (Gold & Shadlen, 2007; Platt & Glimcher, 1999; Shadlen & Newsome, 2001). One study has even shown that microstimulation of LIP neurons has a causal effect on decisions (Hanks, Ditterich, & Shadlen, 2006). However, a recent study that observed the same signals in monkey LIP showed that decisions were largely unaffected if the same area was pharmacologically inactivated (Katz et al., 2016). Further, simultaneous recordings showed that area MT represented momentary motion evidence in the random dot motion task as expected, but that LIP accumulation rates did not match noisy MT fluctuations (Yates et al., 2017). These results raise two important unanswered questions. First, does the brain even accumulate evidence, or is evidence used for decision making in some other mechanism? Second, if the brain does accumulate evidence, then where? Currently we have no answers, but it does not seem like area LIP is the singular neural seat of decision making like was once thought. Of particular relevance, if there is no true evidence accumulation in LIP, then the CPP signal recorded using EEG likely does not reflect causal evidence accumulation despite all its decision-like properties. Another possibility is that the CPP does reflect evidence accumulation - just not in LIP. To speculate, whatever neural mechanism gives rise to the CPP may not necessarily be evidence accumulation itself, but may reflect a downstream readiness signal in preparation to initiate an action. Overall, the existence and possible location of an accumulation process is an area of intense study and debate in the field.

Biological mechanisms for making decisions likely evolved to be adaptive in the situations animals faced throughout millenia. These situations likely include foraging for food

and shelter, mate selection, social communication, and detection of predators in complex and natural environments. It seems likely that if these are the situations where animals have constantly needed to make decisions, that nature would choose a mechanism best optimized for these kinds of decisions. This may be the reason why humans are such notably suboptimal decision makers in new environments such as financial markets or casinos (Taleb, 2007). Understanding the situations humans evolved in, and situations we find more difficult, may help us to select better models of decision making. Further, these considerations may be able to help uncover causes and treatments for people who are affected by systematically maladaptive decision making (like seen in problem gambling).

In conclusion, while this thesis is short on definitive answers, it aims to help bias our search for the mechanism behind how we make decisions. Uncovering the mechanisms of decision making as implemented in the human brain is not only exciting for foundational science, but also promises to make our world a better place - and indeed, studying decision making already has. Studies inspired by tracking mouse trajectories have found individual differences in healthy eating biases, and include potential interventions for healthier eating (Sullivan et al., 2015). Additionally, researchers have exploited systematic biases in human decision making to “nudge” people in the United States to save millions more for retirement (Thaler & Sunstein, 2008). Further, studies of decision and movement signals in the brain have helped create brain-computer interfaces, which hold growing potential for those affected by movement disorders to operate desktop computers or prosthetic devices (Santhanam, Ryu, Yu, Afshar, & Shenoy, 2006). In sum, while we continue to refine our search for models of decision making, findings along the way will continue to shape our decisions and actions for the better.

References

- Albantakis, L., & Deco, G. (2011). Changes of mind in an attractor network of decision-making. *PLoS Computational Biology*, 7(6), e1002086.
- Ando, B., & Graziani, S. (2012). *Stochastic Resonance: Theory and Applications*. Springer Science & Business Media.
- Bogacz, R. (2007). Optimal decision-making theories: Linking neurobiology with behaviour. *Trends in Cognitive Sciences*, 11(3), 118–125.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4), 700–765.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, 12(12), 4745–4765.
- Burnham, K. P., & Anderson, D. R. (2007). *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer Science & Business Media.
- Chapman, C. S., Gallivan, J. P., & Enns, J. T. (2014). Separating value from selection frequency in rapid reaching biases to visual targets. *Visual Cognition*, 23(1-2), 249–271.
- Chapman, C. S., Gallivan, J. P., Wong, J. D., Wispinski, N. J., & Enns, J. T. (2015). The snooze of lose: Rapid reaching reveals that losses are processed more slowly than gains. *Journal of Experimental Psychology. General*, 144(4), 844–863.
- Chapman, C. S., Gallivan, J. P., Wood, D. K., Milne, J. L., Culham, J. C., & Goodale, M. A. (2010). Reaching for the unknown: multiple target encoding and real-time decision-making

- in a rapid reach task. *Cognition*, *116*(2), 168–176.
- Chapman, C. S., & Goodale, M. A. (2008). Missing in action: The effect of obstacle position and size on avoidance while reaching. *Experimental Brain Research*, *191*(1), 83–97.
- Chen, X., & Stuphorn, V. (2015). Sequential selection of economic good and action in medial frontal cortex of macaques during value-based decisions. *eLife*, *4*.
<https://doi.org/10.7554/eLife.09418>
- Christopoulos, V., Bonaiuto, J., & Andersen, R. A. (2015). A biologically plausible computational theory for value integration and action selection in decisions with competing alternatives. *PLoS Computational Biology*, *11*(3), e1004104.
- Christopoulos, V., & Schrater, P. R. (2015). Dynamic integration of value information into a common probability currency as a theory for flexible decision making. *PLoS Computational Biology*, *11*(9), e1004402.
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *362*(1485), 1585–1599.
- Cisek, P. (2012). Making decisions through a distributed consensus. *Current Opinion in Neurobiology*, *22*(6), 927–936.
- Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron*, *45*(5), 801–814.
- Cos, I., Bélanger, N., & Cisek, P. (2011). The influence of predicted arm biomechanics on decision making. *Journal of Neurophysiology*, *105*(6), 3022–3033.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial

- EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21.
- Donner, T. H., Siegel, M., Fries, P., & Engel, A. K. (2009). Buildup of choice-predictive activity in human motor cortex during perceptual decision making. *Current Biology*, *19*(18), 1581–1585.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience*, *32*(11), 3612–3628.
- Dshemuchadse, M., Scherbaum, S., & Goschke, T. (2013). How decisions emerge: Action dynamics in intertemporal decision making. *Journal of Experimental Psychology. General*, *142*(1), 93–100.
- Enns, J. T. (2004). *The Thinking Eye, the Seeing Brain: Explorations in Visual Cognition*. W. W. Norton.
- Evarts, E. V., & Tanji, J. (1976). Reflex and intended responses in motor cortex pyramidal tract neurons of monkey. *Journal of Neurophysiology*, *39*(5), 1069–1080.
- Fajen, B. R., & Warren, W. H. (2003). Behavioural dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology. Human Perception and Performance*, *29*(2), 343–362.
- Freeman, J. B., Dale, R., & Farmer, T. A. (2011). Hand in motion reveals mind in motion. *Frontiers in Psychology*, *2*. <https://doi.org/10.3389/fpsyg.2011.00059>
- Gallivan, J. P., & Chapman, C. S. (2014). Three-dimensional reach trajectories as a probe of real-time decision-making between multiple competing targets. *Frontiers in Neuroscience*, *8*, 215.

- Gallivan, J. P., Logan, L., Wolpert, D. M., & Flanagan, J. R. (2016). Parallel specification of competing sensorimotor control policies for alternative action options. *Nature Neuroscience*, *19*(2), 320–326.
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*(4771), 1416–1419.
- Gescheider, G. A. (2013). *Psychophysics: The Fundamentals*. Psychology Press.
- Gibson, J. J. (2014). *The Ecological Approach to Visual Perception: Classic Edition*. Psychology Press.
- Gold, J. I., & Heekeren, H. R. (2014). Neural mechanisms for perceptual decision making. In *Neuroeconomics* (pp. 355–372).
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.
- Grainger, J., Rey, A., & Dufau, S. (2008). Letter perception: From pixels to pandemonium. *Trends in Cognitive Sciences*, *12*(10), 381–387.
- Hanks, T. D., Ditterich, J., & Shadlen, M. N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience*, *9*(5), 682–689.
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *The Journal of Neuroscience*, *35*(6), 2476–2484.
- Henry, J. F. (2012). *The Making of Neoclassical Economics*. Routledge.
- Howell, D. C. (2011). *Statistical Methods for Psychology*.
- Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of*

Experimental Psychology, 45(3), 188–196.

- Katz, L. N., Yates, J. L., Pillow, J. W., & Huk, A. C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature*, 535(7611), 285–288.
- Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2015). Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *eLife*, 4, e04677.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6), 718–727.
- Kayser, J., & Tenke, C. E. (2006a). Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: I. Evaluation with auditory oddball tasks. *Clinical Neurophysiology*, 117(2), 348–368.
- Kayser, J., & Tenke, C. E. (2006b). Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: II. Adequacy of low-density estimates. *Clinical Neurophysiology*, 117(2), 369–380.
- Kelly, S. P., & O’Connell, R. G. (2013). Internal and external influences on the rate of sensory evidence accumulation in the human brain. *The Journal of Neuroscience*, 33(50), 19434–19441.
- Kiani, R., Cueva, C. J., Reppas, J. B., & Newsome, W. T. (2014). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Current Biology*, 24(13), 1542–1547.
- Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10), 1292–1298.
- Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National*

- Academy of Sciences*, 108(33), 13852–13857.
- Krueger, P. M., van Vugt, M. K., Simen, P., Nystrom, L., Holmes, P., & Cohen, J. D. (2017). Evidence accumulation detected in BOLD signal using slow perceptual decision making. *Journal of Neuroscience Methods*, 281, 21–32.
- Llinás, R. R. (2002). *I of the Vortex: From Neurons to Self*. MIT Press.
- Lopez-Calderon, J., & Luck, S. J. (2014). ERPLAB: An open-source toolbox for the analysis of event-related potentials. *Frontiers in Human Neuroscience*, 8, 213.
- Lynam, D. R. (2013). UPPS-P Impulsive Behaviour Scale--Short English Version. *PsycTESTS Dataset*. <https://doi.org/10.1037/t34116-000>
- McCarthy, G., & Donchin, E. (1981). A metric for thought: A comparison of P300 latency and reaction time. *Science*, 211(4477), 77–80.
- McNicol, D. (2005). *A Primer of Signal Detection Theory*. Psychology Press.
- Montague, P. R., Hyman, S. E., & Cohen, J. D. (2004). Computational roles for dopamine in behavioural control. *Nature*, 431(7010), 760–767.
- Mordkoff, J. T., & Gianaros, P. J. (2000). Detecting the onset of the lateralized readiness potential: A comparison of available methods and procedures. *Psychophysiology*, 37(3), 347–360.
- Murphy, P. R., Robertson, I. H., Harty, S., & O’Connell, R. G. (2015). Neural evidence accumulation persists after choice to inform metacognitive judgments. *eLife*, 4. <https://doi.org/10.7554/eLife.11946>
- Noorani, I., & Carpenter, R. H. S. (2016). The LATER model of reaction time and decision. *Neuroscience and Biobehavioural Reviews*, 64, 229–251.
- O’Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound

- signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1729–1735.
- Padoa-Schioppa, C. (2011). Neurobiology of economic choice: A good-based model. *Annual Review of Neuroscience*, *34*, 333–359.
- Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*(5), 376–404.
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233–238.
- Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, *117*(3), 864–901.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, *20*(4), 260–281.
- Raylu, N., & Oei, T. P. S. (2004). The gambling related cognitions scale (GRCS): Development, confirmatory factor validation and psychometric properties. *Addiction*, *99*(6), 757–769.
- Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, *461*(7261), 263–266.
- Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., & Shenoy, K. V. (2006). A high-performance brain–computer interface. *Nature*, *442*(7099), 195–198.
- Scherbaum, S., Dshemuchadse, M., Fischer, R., & Goschke, T. (2010). How decisions evolve: The temporal dynamics of action selection. *Cognition*, *115*(3), 407–416.
- Selen, L. P. J., Shadlen, M. N., & Wolpert, D. M. (2012). Deliberation in the motor system:

- Reflex gains track evolving evidence leading to a decision. *The Journal of Neuroscience*, 32(7), 2276–2286.
- Shadlen, M. N., & Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4(4), 569–579.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4), 1916–1936.
- Shadlen, M. N., & Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, 90(5), 927–939.
- Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3), 161–168.
- Smith, P. L., & Vickers, D. (1988). The accumulator model of two-choice discrimination. *Journal of Mathematical Psychology*, 32(2), 135–168.
- Song, J.-H., & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences*, 13(8), 360–366.
- Spivey, M. J., & Dale, R. (2006). Continuous dynamics in real-time cognition. *Current Directions in Psychological Science*, 15(5), 207–211.
- Stewart, B. M., Gallivan, J. P., Baugh, L. A., & Flanagan, J. R. (2014). Motor, not visual, encoding of potential reach targets. *Current Biology*, 24(19), R953–4.
- Sullivan, N., Hutcherson, C., Harris, A., & Rangel, A. (2015). Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. *Psychological Science*, 26(2), 122–134.
- Taleb, N. N. (2007). *Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets*. Penguin UK.

- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press.
- Thura, D., Beauregard-Racine, J., Fradet, C.-W., & Cisek, P. (2012). Decision making by urgency gating: Theory and experimental support. *Journal of Neurophysiology*, *108*(11), 2912–2930.
- Thura, D., & Cisek, P. (2016). On the difference between evidence accumulator models and the urgency gating model. *Journal of Neurophysiology*, *115*(1), 622–623.
- Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*(11), 1226–1235.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, *12*(8), 291–297.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592.
- van den Berg, R., Anandalingam, K., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert, D. M. (2016). A common mechanism underlies changes of mind about decisions and confidence. *eLife*, *5*, e12192.
- Wispinski, N. J., Truong, G., Handy, T. C., & Chapman, C. S. (2017). Reaching reveals that best-versus-rest processing contributes to biased decision making. *Acta Psychologica*, *176*, 32–38.
- Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current Opinion in Neurobiology*, *22*(6), 996–1003.
- Yates, J. L., Park, I. M., Katz, L. N., Pillow, J. W., & Huk, A. C. (2017). Functional dissection of signal and noise in MT and LIP during decision-making. *Nature Neuroscience*.

<https://doi.org/10.1038/nm.4611>

Zhang, J. (2012). The effects of evidence bounds on decision-making: Theoretical and empirical developments. *Frontiers in Psychology*, 3, 263.

Appendix

A.1 - Collapsing across rank difference

We opted to collapse Experiment 1 trials across conditions with the same rank difference (e.g., rank 1 vs. 3, rank 2 vs. 4) for both model and experimental data. This decision was supported by an analysis of behavioural reaction times and final choices. Additionally, the selection of a diffusion model of decision making (relative evidence) over a race model (absolute evidence) suggests that it is the *difference* in evidence between options that should be of primary interest (see section A.4).

While the collapsing of trial type by rank difference allows us greater trials per condition for statistical power and for model fitting, it is limited in that it assumes participants' subjective rankings of the food items are linear, and does not take into account research showing relative evidence is not all that matters in reaching decision making (Wisniewski, Truong, Handy, & Chapman, 2017).

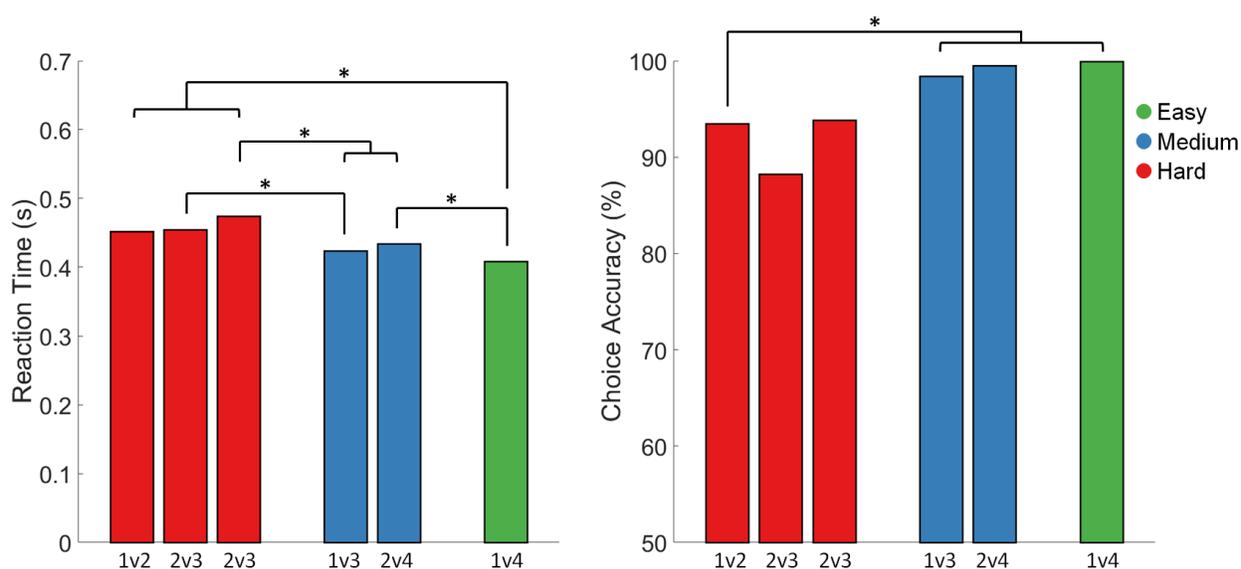


Figure A.1: Group mean reaction times and accuracy for all six unique conditions, which were subsequently collapsed to three critical conditions based on rank difference. Significance bars marked with an asterisk signify statistical significance after Bonferroni-corrected multiple comparisons (15 comparisons per dependent variable). While differences between unique

conditions across critical conditions are not all significant (7 of 11 for RT and 3 of 11 for accuracy), no unique conditions were significantly different from each other within a given critical condition (0 of 4 for RT and 0 of 4 for accuracy).

A.2 - Initial choices and biomechanical bias

Previous models linking decision making to movements have used “initial choices” to represent which bound a decision variable has crossed to initiate a movement (Resulaj et al., 2009; van den Berg et al., 2016). The natural assumption is that if evidence reaches a bound in favor of the left option, movements should begin toward the left. In these studies, initial choices have been defined by looking at the final choice and if a change of mind had occurred. If no change of mind had occurred, the initial choice is the same as the final choice. However, if a trajectory indicated a change of mind, the initial choice and final choice were defined as not matching (Resulaj et al., 2009). The definition of initial choices is critical in these models, as it is used to give an accurate estimate of a person’s decision parameters.

In our current study, we conceptualize movement as a continuous weighting of competing movement angles. This may lead to averaged trajectories up the midline, or reaches where more than one change of mind occurs. Therefore, it is ineffectual to define initial choices by tracing trajectories backward in time. We instead define initial choices by the lateral position of the hand at the start of the reach. However, in the current study several of our subjects exhibited significant biomechanical biases. This resulted in a strong pattern of hand trajectories that initially deviated to one side, only to be compensated for naturally. These biases may arise from seated posture, the participant’s arm configuration during the experiment, or by several other motor variables. Since these biases were extremely consistent and corrected very early in the reach relative to the distribution of changes of mind for our other participants or in other studies (Albantakis & Deco, 2011), we argue they almost entirely reflect non-decision making

behaviour.

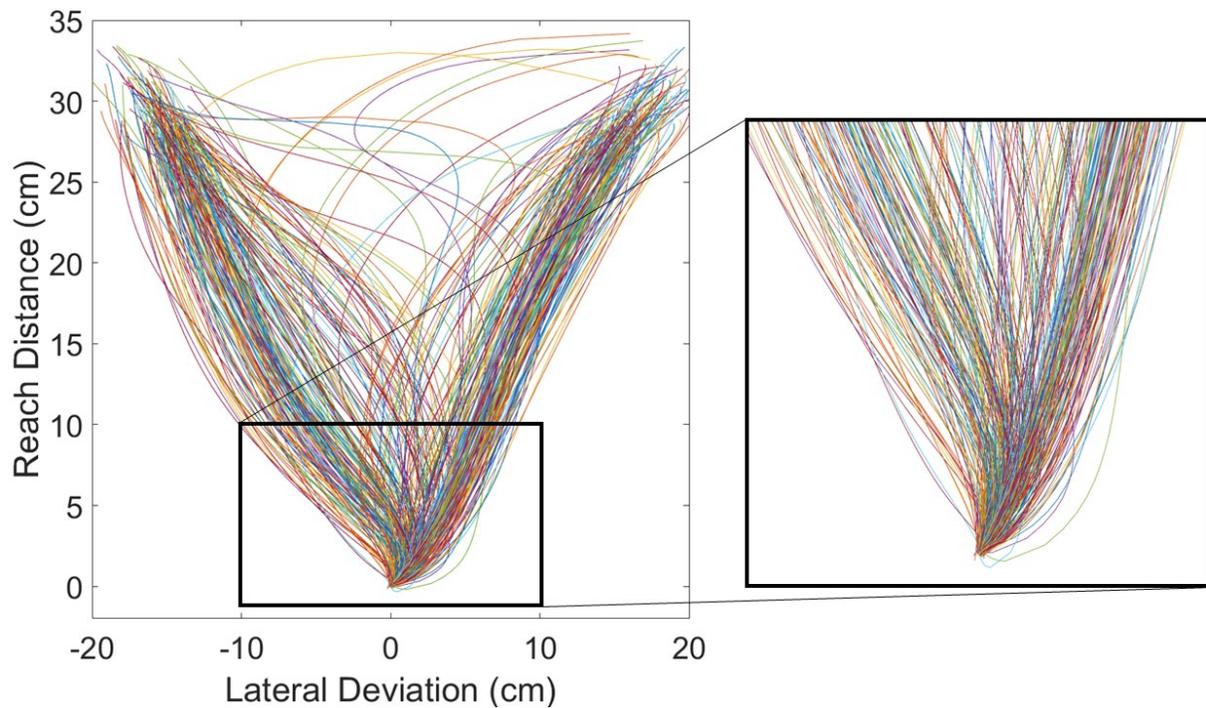


Figure A.2: All raw trajectories from participant 4. We argue the consistent and early lateral deviation rightward off the start position results almost entirely from biomechanics, and is therefore outside of the scope of our model.

As biomechanical and other motor biases are outside the scope of our current model, we aimed to minimize these biases by defining initial choices as the lateral side of space of the hand at 20% through each reaching movement (~ 6 cm) for every subject. This method takes advantage of the natural compensation of a person's hand later on in the movement, but comes at the cost of capturing fewer "true" decision-related changes of mind that occur early in the reach.

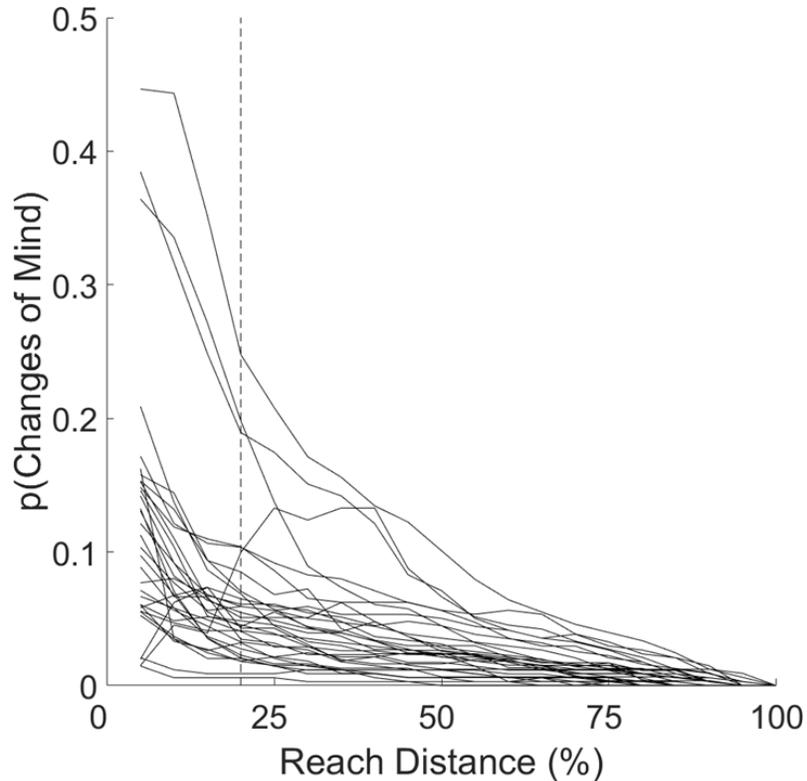


Figure A.3: Plot of changes of mind over time for all 32 participants. Changes of mind were calculated in bins of 5% of total normalized reach distance. The proportion of changes of mind early in the reach is much higher than observed in past studies (Resulaj et al., 2009; Albantakis & Deco, 2011). Dashed line represents our proposed changes of mind cutoff in order to remove changes of mind due mostly to biomechanics instead of decision processing.

A.3 - Rejecting bad drift diffusion fits

In Chapter 2, participants who had an R^2 of 0 for both RT and accuracy were rejected for further analysis. We argue that if a drift diffusion model does not provide a good fit for a subject's data, then modelling of their reaches based on a diffusion process will fail. These rejections were always for participants who showed RT and Accuracy patterns that could not be explained by the model. For example, participant 11 was faster for medium decisions than hard decisions. This is perhaps due to a mismatch between their self-reported preferences and their true preferences.

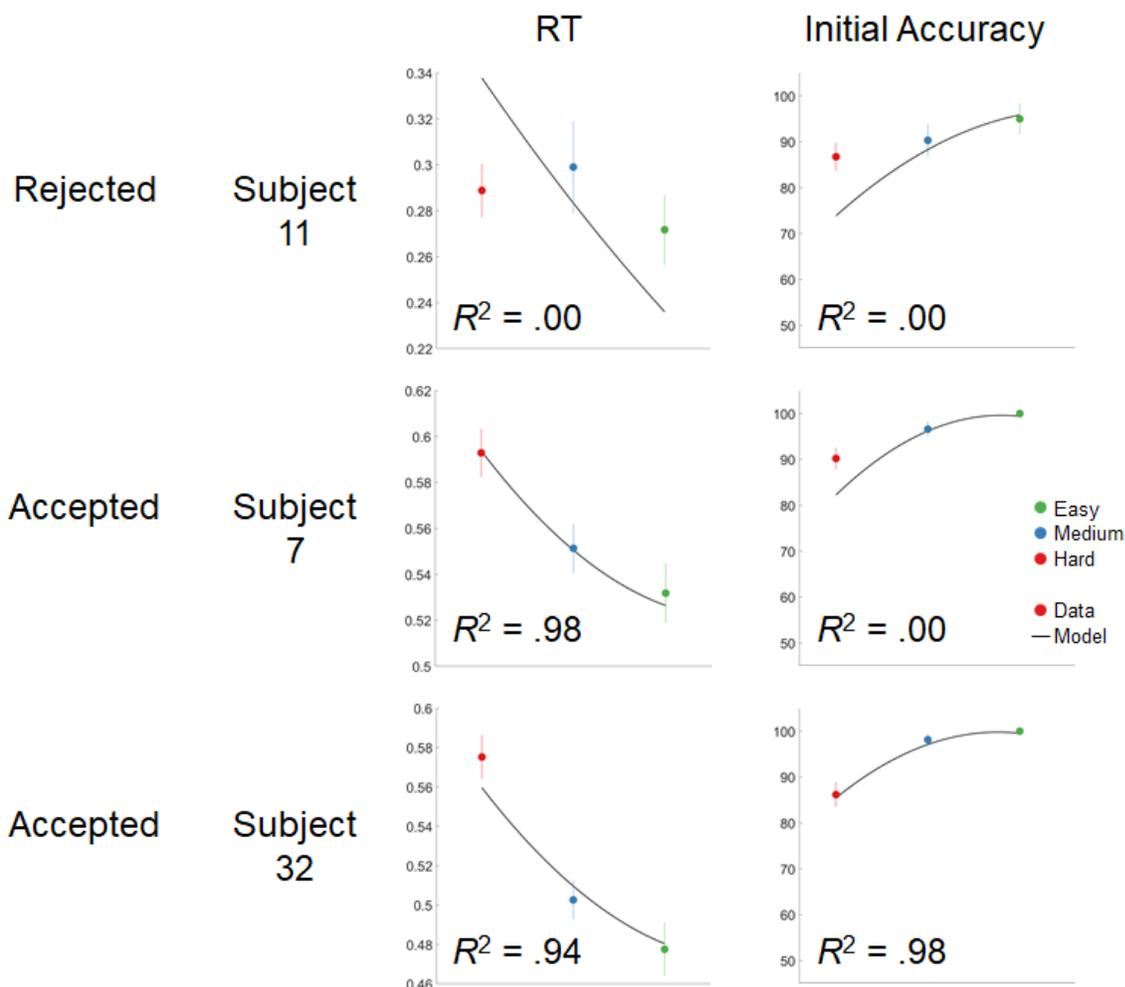


Figure A.4: Example RT and accuracy fits for three individuals from Chapter 2. Subject 11 was rejected because the drift diffusion fit could not account for reaction time, nor accuracy. In contrast, participants such as subject 7 and subject 32 were accepted for reach modelling. Of note, subject 7 did not have an adequate fit for initial accuracy, but did have an acceptable fit for reaction time (as did several other subjects).

A.4 - Drift diffusion vs. race models

Sequential sampling models in cognitive psychology and neuroscience most often take the form of drift diffusion (Ratcliff & Rouder, 1998) or race (Smith & Vickers, 1988) models. Formal models linking decision making to movement have used both drift diffusion (Resulaj et al., 2009; Burk et al., 2014), race (van den Berg et al., 2016), and others such as neurodynamical

models (Christopoulos, Bonaiuto, & Andersen, 2015). To select the most appropriate framework, we compared our behavioural data in Chapter 2 to both theorized race and diffusion models.

Normally it is difficult to distinguish between race and diffusion models because of their similarity and the kinds of tasks typically used in perceptual decision making. However, our experiment included several conditions which allowed us to distinguish between these two models. Critically while the diffusion process relies on a relative difference in accumulated evidence to reach bound, a race model's independent accumulators only consider the absolute level of evidence for each option. Here we can take conditions where the absolute value of the best option is unchanged while the alternative option is of a varying value (see Figure A.4 A). Additionally, we can take conditions where the relative value difference between options is constant, but the absolute values vary (see Figure A.4 B). While we had few trials per person in these conditions, a group-level analysis allows us to distinguish whether absolute or relative value drives reaction time.

Overall, the reaction times in the current experiment are more driven by relative value difference between options as opposed to the absolute value of the options. We find that the diffusion model provides a better fit to our data compared to a race model. Therefore we elected to implement a drift diffusion model to explain reaction times, in line with earlier (Resulaj et al., 2009) but not later (van den Berg et al., 2016) models.

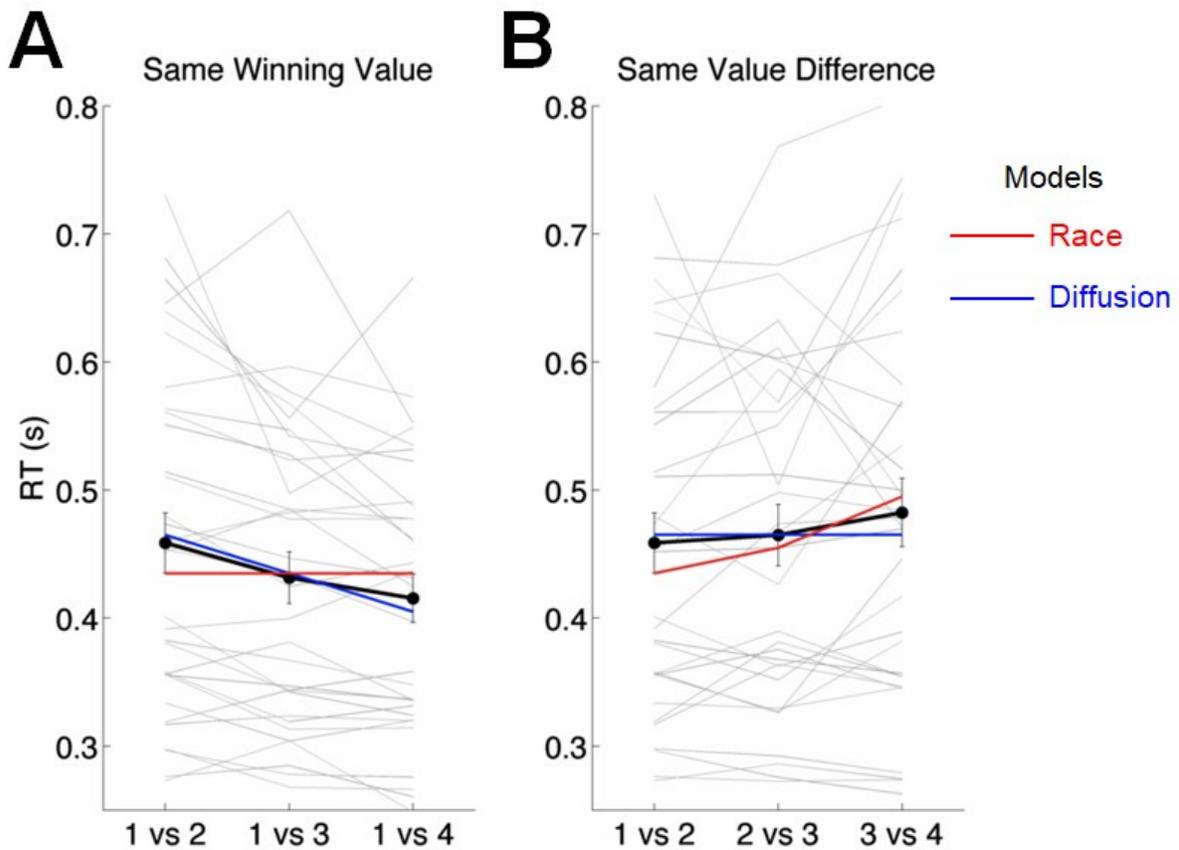


Figure A.5: Grand mean reaction times from behavioural data (black), and theorized race (red) and drift diffusion (blue) models. Gray lines are individual subject means. Error bars are S.E.M. A) Three conditions where the best ranked option is always present. A race model generally predicts the most valued option reaches threshold first regardless of its competitor, whereas a diffusion model that takes the relative difference between values predicts faster initiation times with a larger relative difference between options. B) Three conditions where the subjective value difference between options is constant. Here a diffusion model generally predicts the same reaction time because of a constant value difference between options, whereas a race model predicts faster reaction times when a higher valued option is available.

A.5 - Tables of model parameters

Fitted parameters for Chapter 2 participants

<u>Participant</u>	<u>k</u>	<u>B</u>	<u>t_{exp} (ms)</u>	<u>u_0</u>	<u>v_0</u>	<u>R</u>
1	0.782	0.018	295	0.016	-0.038	0.96
2	0.764	0.016	310	-0.009	0.178	0.94
3	0.770	0.018	279	-0.032	0.483	0.98
4	2.333	0.000	270	-0.282	0.027	1.00
5	1.511	0.009	257	-0.037	-0.018	0.93
6	0.793	0.024	431	0.018	-0.158	0.98
7	0.648	0.023	476	-0.035	0.208	0.86
8	0.557	0.028	246	-0.064	0.448	0.77
9	0.275	0.029	180	0.013	-0.012	0.93
10	0.156	0.050	149	-0.003	0.004	0.96
11	0.269	0.032	76	-0.007	0.017	0.73
12	0.599	0.021	475	-0.011	0.006	0.98
13	1.116	0.013	304	-0.043	0.367	1.00
14	0.170	0.055	101	-0.010	0.010	0.87
15	0.378	0.030	448	0.002	-0.017	0.87
16	1.168	0.026	394	0.045	-0.304	0.94
17	1.213	0.009	430	-0.052	-0.165	0.95
18	0.653	0.023	357	-0.001	0.050	1.00
19	0.252	0.011	340	-0.186	0.566	0.92
20	0.447	0.037	384	0.004	-0.079	1.00
21	0.531	0.019	404	-0.022	0.098	0.92
22	0.736	0.019	301	-0.004	0.326	0.95
23	0.420	0.040	375	0.023	-0.218	0.97
24	0.671	0.025	375	-0.008	-0.038	1.00
25	0.566	0.023	369	0.001	0.140	0.97
26	1.115	0.014	314	-0.017	0.117	1.00
27	3.880	0.000	278	-0.518	-0.106	1.00
28	0.689	0.011	462	-0.105	-0.038	0.88
29	0.323	0.032	403	-0.003	-0.027	0.83
30	2.176	0.003	328	-0.106	-0.009	0.99
31	0.543	0.027	397	0.031	-0.084	1.00
32	0.552	0.026	414	0.011	-0.023	0.98

Note. Subjects with strikethrough values were excluded from model analysis for inadequate fits between drift diffusion model predictions and behavioral data.

Fitted parameters for Chapter 3 keypress participants

<u>Participant</u>	<u>k</u>	<u>B</u>	<u>$t_{50}(ms)$</u>	<u>u_n</u>	<u>v_n</u>
1	0.167	0.363	172	-0.010	0.000
2	0.736	0.290	378	-0.001	0.000
3	0.283	0.563	439	0.004	0.012
4	0.359	0.321	368	-0.003	0.000
5	2.616	0.448	436	-0.001	0.000
6	0.738	0.290	378	-0.002	0.000
7	1.005	0.577	438	0.003	0.081
8	2.614	0.449	435	0.000	0.000
9	0.340	0.318	158	-0.002	0.000
10	0.243	0.339	157	-0.006	0.014
11	0.218	0.303	124	0.000	0.000
12	0.917	0.576	438	-0.009	0.001
13	1.485	0.576	438	-0.006	0.001
14	3.037	0.412	493	0.001	-0.001
15	1.542	0.328	388	-0.005	0.000
16	1.015	0.576	439	-0.014	0.079
17	1.716	0.577	414	0.002	0.008
18	0.356	0.295	159	0.001	0.000
19	1.079	0.576	438	0.001	0.000